Visually Evaluating Courses of Action in a Contested Network Environment

Benjamin Blease, Jaime Peña, Kenneth Alperin, Leslie Shing, Allan Wollaber
MIT Lincoln Laboratory
Lexington, MA

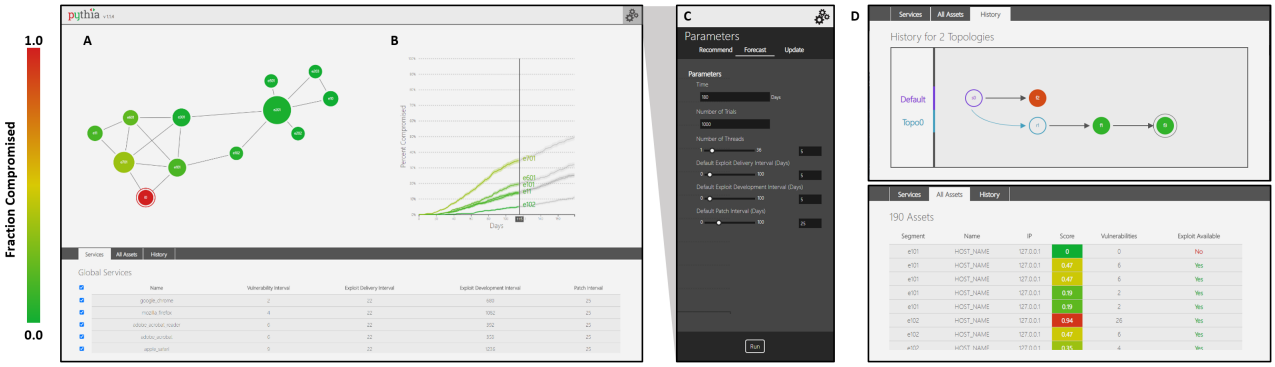

Fig. 1: Pythia components include a (A) network segmentation force diagram, (B) fractional compromise rate per segment forecast plot, (C) parameters menu for segmentation recommendations and forecasting, (D) and a dynamic user history for trading off risk comparisons between segmentation recommendations and simulation runs

***Abstract*—Current cyber situational awareness (SA) technologies primarily have aimed at better comprehension of existing elements and immediate risks or threats in networked environments. This has left a gap in cyber defenders' ability to trade off risks in proposed courses of action over an upcoming timeframe of interest under nominal or dynamic threat conditions. We present Pythia, a prototype tool that is designed to enhance cyber SA by explicitly considering attacker-defender activities in response to a selected course of action over a timeframe of interest with a visualization capability that, by adding the temporal dimension, is a novel, critical component for enhanced cyber SA. Pythia leverages a cybersecurity model utilizing High Performance Computing (HPC) simulations to provide four crucial features to improve a cyber-defenders situational awareness: (1) a high-level interactive view of the defenders current network architecture and recommended architectures, including the fraction of hosts compromised within each segment, (2) a time-dependent plot of projected fractional compromise rates per segment, (3) an interactive parameters menu in which users can fine tune parameters to be sent to the backend algorithms, and (4) an interactive history of user actions in which they can compare potential strategies to mitigate risk in their network.**

I. Introduction

Internet-connected hosts provide potential inroads to resources through software service vulnerabilities. Vulnerabilities are flaws in software an attacker may exploit for unauthorized access to a network. In most networks, hosts are connected to the internet with firewalls attempting to filter traffic and prevent compromise. Once an attacker successfully delivers an exploit for a vulnerability in a particular software service, they can leverage it to obtain a degree of lateral movement towards a target of interest in the network. Segmenting a network by restricting traffic to certain hosts with similar security and operational needs limits an attacker's mobility. Such a segmentation plan is considered a potential Course of Action (CoA) for the network defender. In this paper, we consider CoAs that mitigate the risk of attackers compromising a defended network, such as disabling software services, adding firewall rules, patching more frequently, or segmenting a network. We employ a model in which hosts in the network segments are assumed to share all software services, and each network segment is connected via software services [1]. Detecting and parameterizing vulnerabilities, software services, and exploits in order to simulate expected outcomes of attacker/defender scenarios is crucial in developing and recommending CoAs, as cyber analysts currently cannot visually trade off potential CoAs that inhibit lateral movement.

We present a dashboard prototype with the following goals: display a simplified overview of the networks segments, connections, and their current conditions, differentiate future forecasting from initial conditions, provide meaningful interactions and exploration, and allow previously considered CoAs to be compared; see Fig. 1.

II. Background

Graphical decision support tools for cyber SA most often present analysts with prioritized lists of vulnerabilities or other cyber assets to target remediation efforts. These tools provide cyber SA by displaying discovered entities and endowing them with relative risk scores, but they generally do not predict future risk status under nominal conditions. The Endsley

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited. This material is based upon work supported by the Department of the Army under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Army.
This work was supported in part by high performance computer time and resources from the DoD High Performance Computing Modernization Program.

model is a generalized description of SA that includes three levels: (1) perception of elements in the environment, (2) comprehension of the current situation, and (3) prediction of future status [2]. For example, vulnerability discovery systems provide levels 1 and 2 cyber SA. Gray et al. [3] have designed an inventive visualization of large networks using a Radial Reingold-Tilford Tree, although it does not incorporate time-dependence for level 2 or 3 SA. O'Hare et al. [4] present a more complete, but older, SA tool that competently addresses levels 1 and 2 simultaneously. The tool provides a top-level view of the attack paths in the network and allows analysts to explore potential near-term attacks for hosts and network structures. Additionally, Chu et al. [5] present a similar tool with a network view containing hosts and their fitness. Our application differs in that, instead of recommending a "worst" host or vulnerability to remediate or shortest attack path to remove, time-dependent impact is displayed as a result of executing potential CoAs in order to evaluate its efficacy. In the case of a segmentation CoA, the optimal CoA may be implemented by a software defined networking solution. We have developed a tool that addresses all three levels of the Endsley model through the use of visualizations. By interacting with previously developed technologies through a backend, we have built an intuitive, easily deployable dashboard that utilizes visualizations to provide high-level views (level 1), initial risk assessments (level 2), and forecasts (level 3). Because the implementation of CoAs will likely have some impact on network usability, a visual history provides users with a means to weigh trade offs between CoAs.

III. System Description

Pythia is designed to provide new, unique features including comparative user histories of CoA forecasts and threat modeling, while improving cyber analysts' precision for typical workflows, such as network monitoring using high-level visualizations. In addition, it provides a platform for further development of cyber SA visualizations and tools.

Computations that require high performance computing, such as simulation and learning algorithms, are handled by Pythia's backend [6], [7], and results are displayed by the dashboard. This poster focuses on the dashboard component of Pythia–a web application written in ReactJS. The dashboard provides an intuitive user experience that utilizes visualizations to largely convey top-level information, while reserving fine-grain control for traditional tables and forms.

The network graph in Fig. 1A is a proposed segmentation CoA. The evaluated network of interest is represented by three categories of values: (1) segments, (2) connections, and (3) services. Each segment is a collection of hosts with similar service needs. Each connection includes a list of services, which are each composed of a name, vulnerability, exploit delivery, exploit development, and patch interval. Each interval is a parameter for the LR9 attacker/defender model [8]. The forecast plot in Fig. 1B displays the forecasted fraction of hosts that are compromised per segment (percent) over time (days) for the improved network topology. Users may compare this

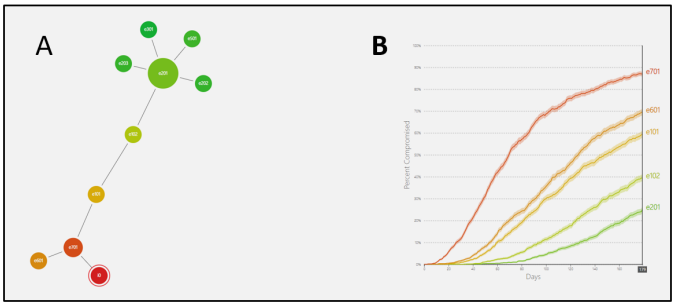

Fig. 2: Initial forecast for a user's current network topology

suggested topology with the default (user-specific) network topology displayed in Fig. 2A. The forecast in Fig. 2B displays a higher fraction of compromised hosts on average when compared to the recommended configuration in Fig. 1B.

The parameters menu in Fig. 1C provides control over recommendation and forecast simulation settings. In addition, users may update their current network topology to reflect any changes in hosts or software configurations.

The History Tab in Fig. 1D provides a clickable, graphical history of previous recommendations and forecasts in a tree structure similar to version control tools.

IV. Future Work

Continued development of Pythia will include various feature improvements and operator validation. In particular, the history component will be extended to enable a comparison of two CoAs' details. In addition, an extension featuring software defined perimeter as a new CoA is planned. Techniques to visualize large networks, such as selecting subsets of a larger network or implementing a network creator, will also be explored.

References

[1] N. Wagner, C. Ş. Şahin, J. Pena, J. Riordan, and S. Neumayer, "Capturing the security effects of network segmentation via a continuous-time markov chain model," in *Proceedings of the 50th Annual Simulation Symposium*, p. 17, Society for Computer Simulation International, 2017.

[2] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," in *Situational Awareness*, pp. 9–42, Routledge, 2017.

[3] C. C. Gray, P. D. Ritsos, and J. C. Roberts, "Contextual network navigation to provide situational awareness for network administrators," in *2015 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 1–8, IEEE, 2015.

[4] S. O'Hare, S. Noel, and K. Prole, "A graph-theoretic visualization approach to network risk analysis," in *International Workshop on Visualization for Computer Security*, pp. 60–67, Springer, 2008.

[5] 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*, pp. 22–33, ACM, 2010.

[6] A. Wollaber, J. Peña, B. Blease, L. Shing, K. Alperin, S. Vilvovsky, P. Trepagnier, N. Wagner, and L. Leonard, "Proactive cyber situation awareness via high performance computing," in *2019 IEEE High Performance Extreme Computing Conference*, October 2019.

[7] K. Alperin, A. Wollaber, D. Ross, P. Trepagnier, and L. Leonard, "Risk prioritization by leveraging latent vulnerability features in a contested environment," in *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, AISec'19, (New York, NY, USA), p. 4957, Association for Computing Machinery, 2019.

[8] R. P. Lippmann and J. F. Riordan, "Threat-based risk assessment for enterprise networks," *Lincoln Lab. J*, vol. 22, no. 1, pp. 33–45, 2016.