CLIPS/ActiveSDN for Automated and Safe Cybersecurity Course-of-Actions Orchestration

Md Mazharul Islam, Ehab Al-Shaer, Ashutosh Dutta, Mohammed Noraden Alsaleh
{mislam7,ealshaer,adutta6,malsaleh}@uncc.edu
University of North Carolina at Charlotte

ABSTRACT

Continuous attack reports such as data breach, malware, phishing and spamming attack published daily indicate that cyber attack is inevitable in our daily life. Sometimes it takes days, even month to detect and mitigate such stealthy attacks. These require to make network systems resilient against attacks with a high assurance of defense mechanisms that can go beyond attack detection with safe mitigation. That's why we developed a flexible yet expressive policy specification language called CLIPS for Active Cyber Defence, and provably-correct policy refinement engine, ActiveSDN to enable a safe, efficient construction and execution of Course-of-Action workflow composed of investigating for analysis and mitigating for reconfiguration actions to support cyber resilience automation.

ACM Reference Format:
Md Mazharul Islam, Ehab Al-Shaer, Ashutosh Dutta, Mohammed Noraden Alsaleh. 2019. CLIPS/ActiveSDN for Automated and Safe Cybersecurity Course-of-Actions Orchestration. In *Hot Topics in the Science of Security Symposium (HotSoS), April 1–3, 2019, Nashville, TN, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3314058.3317724

1 INTRODUCTION

Recent Cybercrime statistics show that hundreds of thousands of new malware samples are produced every day and will potentially grow in the future [6]. More than 4,000 ransomware attacks occur each day in the US only [4]. People know the risks that come with clicking unknown links in emails and yet still click these links. It takes most business about 197 days to detect a breach on their network and interestingly 75% of network vulnerability are due to human misconfiguration[2]. All of these indexes demand to make system auto resilient against cyber attacks with a high assurance of defenses techniques to detect and mitigates attacks adaptively. Hence, we provide an extensible and verifiable framework (CLIPS/ActiveSDN) to create high level yet *safe* auto resilient security policies that can also be deployable into any network for immediate attack mitigation. We provide CLIPS, a policy generation language, with a verifier that ensures the safety of the policy in ActiveSDN, an OpnedayLight Software Defined Networking (SDN) controller[5], that makes sure the deployment of the policy. Using CLIPS, an event-driven security policy can be created that triggers a series of Course-of-Action (CoA). A CoA is an ordered set of cyber *actions* commands such as blocking traffic, enabling mutation, migrating a virtual machine, etc. A CLIPS policy verifier will measure the *safety* of the policy that the policy rules, whether they are executed concurrently or sequentially, do not: (1) conflict with each other, (2) cause contentions on the cyber resources, or (3) introduce violations to the high-level network mission requirements.

Now the goal of the provided framework is to take input as a set of CoAs that belong to different Active Cyber Defence (ACD) rules, pre and post-conditions for each action in a CoA, the current state of the control variables, the current state of network data plane configuration, network mission requirements. The outcome is to find a global orchestrated CoA workflow GOAL, an execution schedule of all actions, that satisfies some critical expectations. These expectations are resource integrity where no resource conflicts between concurrently executed actions, no shared object or control variable can be modified simultaneously, every action will be executed correctly satisfying its precondition, maximizing the action execution concurrency while considering the temporal action dependency within each CoA. Finally, no violations for the network mission requirements during the GOAL execution.

2 CLIPS

In figure 1 we have shown the syntax of CLIPS language. The CoA is represented as a process that executes single or multiple actions composed in different modes. It can one action, multiple consecutive actions using the sequential composition operator (;), multiple parallel actions using the parallel composition operator (||), or a conditional expression (ψ ? Λ_1 : Λ_2) that executes Λ_1 if the condition ψ evaluates to true, otherwise, it executes Λ_2. Cyber actions denoted by α are the basic building blocks, where each cyber action specifies that a command (f) (e.g., blocking or forwarding traffic flows, migrating virtual machines, and disabling/enabling services in the network) be executed by a certain actuator (u) (e.g., a firewall, an SDN switch, a virtualization server or controller) on a specific object (o) (e.g., a specific traffic flow or a virtual machine). In Figure 1, we list only a subset of the commands, actuators, objects, and control variables, and it is open for the users to define their own. Each command f can also take a set of arguments, if needed, in the format (*argument name* = *value*). Examples of these arguments include the output *port number* in the case of forward commands, and the name of the destination server in the case of migration commands.

3 ACTIVESDN

ActiveSDN is an open programming environment that enables developing or prototyping advanced ACD mechanisms and agility capabilities rapidly and safely using SDN. ActiveSDN leverages its facilitatings by providing an ActiveSDN API that gives access to cyber agility, ACD primitives and OpenFlow management functions

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
HotSoS, April 1–3, 2019, Nashville, TN, USA
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7147-6/19/04.
https://doi.org/10.1145/3314058.3317724

using OpenDayLight controller[5]. Besides, ActiveSDN provides a decision-making engine that is capable of solving computation hard problems using Constraint Satisfaction Solvers (SMT Z3)[3], Partially Observable Markov Decision Process (POMDP) [1], Game Theory, Machine Learning, etc. to optimize defense actions. ActiveSDN composes the ACD policy, ensure safe, low-level configuration changes and deploy the policy into the network. Hence, ActiveSDN makes cyber defenses techniques as a service that user can access without taking any configuration headache yet mitigate attacks immediately.

4 DECEPTION AS SERVICE

As an example of ACD, CLIPS/ActiveSDN can provide deception as a service to user for protecting network resources from a stealthy attacker. In order to determine the optimal deception parameter to protect the critical resources, the framework requires to know the current risk (proximity of the attackers towards the critical resources) imposed by the attacker and predict his/her next action. As the consequence of a defense action also depends on the attack action at the current situation, we need to integrate the uncertain attack behavior into decision-making. Therefore, our proposed framework considers the strategic reasoning between the attacker and the defender. However, besides being stealthy, the attackers may also be adaptive to deceive the defenders which make decision-making for optimal defense (deception parameter) more complex and complicated.

Though it is hard to detect the stealthy attacker, we may develop a belief over the attacker's actions using the available risk indications (e.g., IDS alarms, log files, and others). Hence, we can extend our knowledge database about the behaviors of the stealthy attackers using such risk indications as observations. Moreover, a rational attacker wants to maximize the damage by executing the optimal attack actions while following a specific strategy. As a consequence, from the knowledge database of attack behavior, we can extract patterns in adversary behaviors to predict the probabilistic distributions over the next attack actions. Therefore, we are learning the attack behavior to understand the consequences of our defense actions in the considered cyber environment using reinforcement learning.

Finally, we formulate the problem of selecting the optimal defense action by solving a POMDP model for the defender that incorporates the attack behaviors. POMDP maximizes the payoff of a defense action considering the current situation (risk) of the environment. Moreover, to enable the adaptive cyber deception, we develop an online planning tool that solves a dynamic POMDP model at each time-sequence based on the observations of the environment.

5 EVALUATION

We evaluated the performance of our orchestration synthesis framework in terms of the time required to generate the optimal GOAL that satisfies the Resource Integrity, the Action Integrity, and the CoA Concurrency properties. We evaluated the performance with respect to the total number of actions. We tested our framework for more than 60 sets of different actions, and the results are depicted in Figure 2. For up to 500 actions, the average processing time was

$$
\begin{aligned}
\textit{Policy } \Pi &::= \{\langle event\rangle \twoheadrightarrow \Lambda\} \\
\textit{CoA } \Lambda &::= \alpha \mid \Lambda\,;\,\Lambda \mid \Lambda \parallel \Lambda \mid \psi\ ?\ \Lambda : \Lambda \\
\textit{Action } \alpha &::= [\boldsymbol{A}\ \psi]\ \boldsymbol{DO}\ f(\{b = \langle number\rangle\})\ \boldsymbol{BY}\ u\ \boldsymbol{ON}\ o\ [\boldsymbol{G}\ \psi] \\
\textit{Expression } \psi &::= v['] \mid v['] \bowtie \langle num\rangle \mid v['] \bowtie \psi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \\
\textit{Operator } \bowtie &::= == \mid > \mid < \mid \leq \mid \geq \mid + \mid - \mid \times \mid / \\
\textit{Command } f &::= \pi \mid \tau \\
\textit{Configuration } \pi &::= block \mid forward \mid limit \mid inspect \mid encrypt \mid \\
&\quad\ enable \mid disable \mid migrate \mid reroute \mid \langle other\rangle \\
\textit{Investigation } \tau &::= SNMPGet \mid LogAudit \mid SplunkActions \mid \\
&\quad\ MITRE\text{-}ATT\&CK/CWE/CAPEC\text{-}InvActions \mid \langle other\rangle \\
\textit{Argument } b &::= portno \mid threshold \mid ccuid \mid \langle other\rangle \\
\textit{Actuator } u &::= \langle list\ of\ unique\ actuators\rangle \\
\textit{Object } o &::= \langle list\ of\ unique\ objects\rangle \\
\textit{Variable } v &::= capacity \mid bandwidth \mid \langle other\rangle
\end{aligned}
$$

Figure 1: ACD Policy Language Syntax.

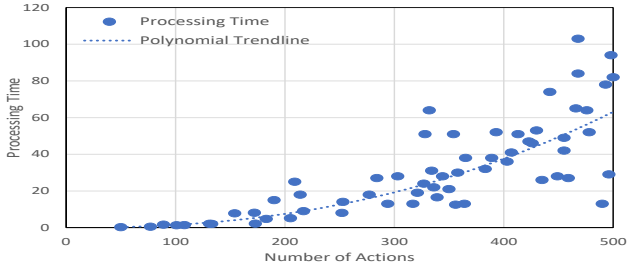


Figure 2: The impact of number of actions.

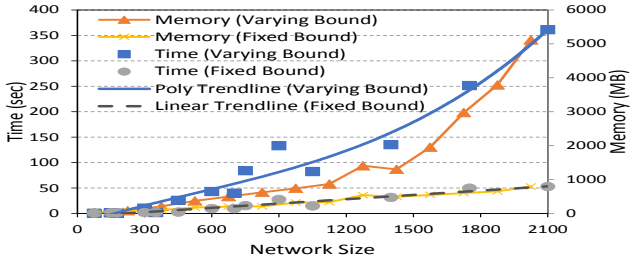


Figure 3: The impact of network size.

within 60 seconds, and it exhibits a polynomial growth rate with the time complexity grows reasonably according to the number of actions. Figure 3 shows that in the case of a fixed bound, the time and space requirements are linear with respect to the network size. However, in the case of varying bound, the performance is affected by both: the network size and the bound. The growth rate with respect to the network size is best described by a quadratic polynomial function.

REFERENCES

[1] Darius Braziunas. 2003. Pomdp solution methods. *University of Toronto* (2003).

[2] Data Breach Discovery 2018. Survey Finds Breach Discovery Takes an Average 197 Days. (2018). https://securityboulevard.com/2018/07/survey-finds-breach-discovery-takes-an-average-197-days/
[3] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
[4] FBI 2019. Ransomware Prevention and Response for CISOs. (2019). https://www.fbi.gov/file-repository/ransomware-prevention-and-response-for-cisos.pdf/view
[5] Jan Medved, Robert Varga, Anton Tkacik, and Ken Gray. 2014. Opendaylight: Towards a model-driven sdn controller architecture. In *Proceeding of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks 2014*. IEEE, 1–6.
[6] PandaLabs 2015. 27% of all recorded malware appeared in 2015. (2015). https://www.pandasecurity.com/mediacenter/press-releases/all-recorded-malware-appeared-in-2015/