DiscX: Should Governments Disclose or Exploit Vulnerabilities?

As more and more vulnerabilities are found in the growing space of connected devices, governments are often confronted with the problem of what to do when their cyber-warfare units discover a new vulnerability. Generally, the two options are either to disclose the vulnerability, which may be ideal if the product containing the vulnerability is manufactured or in wide use in the country, or to hold it in secret and develop an exploit. The decision to disclose or exploit often depends on several factors, such as whether an adversary will discover and disclose the vulnerability, what damage an adversary might cause if they choose to exploit, and how disclosure might protect a nation's corporations. The US and several European government use a Vulnerability Equities Process (VEP) to facilitate such decision making. NSAIL faculty jointly with military experts from the Netherlands Defense Academy were the first to develop a Repeated Cyber Warfare Game (RCWG) formulation of the problem, as well as a prototype system called DiscX. Subsequently, NU faculty have developed a reinforcement learning based SmartVEP framework that removes many assumptions made within DiscX.

DiscX is intended to augment the current decision-making procedure for "exploiting vs. disclosing" cyber-vulnerabilities with a rigorous tool that uses agency experts' inputs to help agencies such as the US Government's Equities Review Board arrive at an optimal solution.

Repeated Cyber Warfare Game (RCWG)

RCWG first considers a one stage game and then extends it to a repeated stage setting. In a stage, there are 2 players representing the 2 opposing countries, as well as a joint strategy and a joint payoff function of the two players. The payoff function depends on several factors, including development cost, exploit payoff and disclosure payoff. The one-stage game is formulated as an alternated stochastic optimization problem for computing the best response of a player.

RCWG experiments used data from authoritative national security sources, like the Rand Corporation, as actual data from governments is usually classified. Using a dataset of over 200 zero-day exploits from 2002 to 2016, values for vulnerability related parameters were obtained, such as a mean time of discovery of 200 days and an average cost of developing an exploit based on a zero-day vulnerability at $30,000.

The experiments showed that RCWG yields a much higher payoff than 3 baselines of random, always exploit, or always disclose strategies. Additionally, RCWG converges after 220 iterations, at around 20 seconds on average. It is clear that RCWG efficiently provides the optimal decision in cyber warfare strategy.

DiscX

Researchers now at NSAIL also developed the prototype DiscX system to support government decision making. This system allows for user input of vulnerability information, which runs through RCWG and gives a recommendation on the decision to make based on their knowledge of a specific situation.

https://sites.northwestern.edu/nsail/projects/discx

Sample screenshots of the DiscX system are provided for display below.

Opening screen of DiscX

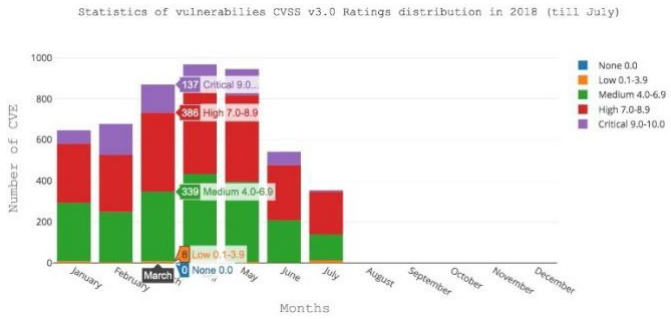

Vulnerability statistics after initial selection from opening screen

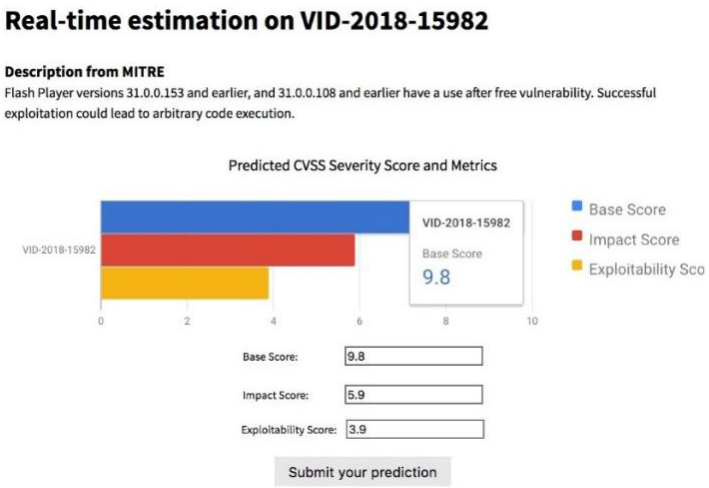

Vulnerability score information and user modification

SmartVEP

RCWG assumes a two-player game. However, there may be several adversary nations and security firms that disclose vulnerabilities. (ii) RCWG assumes that we know the payoff function, development costs/time, and more about other players. (iii) It assumes we know the opponent model. Our new Smart VEP framework removes all 3 assumptions and uses online learning to balance the public's interest (disclosure) vs. the national security interest in exploiting the vulnerability.

Smart VEP formulates this as a novel Multiplayer Adversarial Bandit (AB) problem with two types of information structures: (i) the information structure known before and (ii) after playing an action. No existing adversarial bandit algorithm considers both.

In addition, Smart VEP scales the search by two innovations: (i) only explicitly exploring the action with the largest time value for exploitation, and (ii) providing an asynchronous method for exploration and learning. We show that Smart VEP significantly outperforms several baselines.

Additional Information

References

1. Haipeng Chen, Qian Han, Sushil Jajodia, Roy Lindelauf, V.S. Subrahmanian, and Yanhai Xiong. "Disclose or Exploit? A Game Theoretic Approach to Strategic Decision Making in Cyber Warfare." Accepted for publication in *IEEE Systems Journal*. https://ieeexplore.ieee.org/document/8967205
2. Y. Zhang and V.S. Subrahmanian. SmartVEP: A Smart Vulnerability Equities Process based on Multiplayer Adversarial Bandits, *submitted,* Feb 2022.

Video

https://northwestern.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=cddad1a2-9349-4b2d-97a1-aebe0155ec8a

PARTICIPANTS

Lead: V.S. Subrahmanian

Haipeng Chen, Qian Han, Sushil Jajodia, Roy Lindelauf, Yanhai Xiong, and Youzhi Zhang.

Northwestern | BUFFETT INSTITUTE FOR GLOBAL AFFAIRS

Northwestern | McCORMICK SCHOOL OF ENGINEERING