Optimal Cyber-Defense Strategies for Advanced Persistent Threats

Jeffrey Acquaviva[1], Mark Mahon, Bruce Einfalt[1], Tom LaPorta[2]

We introduce a novel mathematical model that treats network security as a game between cyber attackers and network administrators. The model takes the form of a zero-sum repeated game where each sub-game corresponds to a possible state of the attacker. Our formulation views state as the set of compromised edges in a graph opposed to the more traditional node-based view. Both players move independently and in continuous time allowing for the possibility of one player moving several times before the other does.

This model shows that defense-in-depth is not always a rational strategy for budget constrained network administrators. Furthermore, a defender can dissuade a rational attacker from attempting to attack a network if the defense budget is sufficiently high. This means that a network administrator does not need to make their system completely free of vulnerabilities, they only to ensure the penalties for being caught outweigh the potential rewards gained.

Motivation:

- It is important to understand how cyberattackers move through a network
- Current research on attack graphs give the probability of success for a path, but not the likelihood a path will be taken
- Likelihood must combine attacker capability, attacker reward desires, and defense strategy
- Network defenders are budget constrained and must find optimal allocation of resources to prevent attacks

Game Theory Model

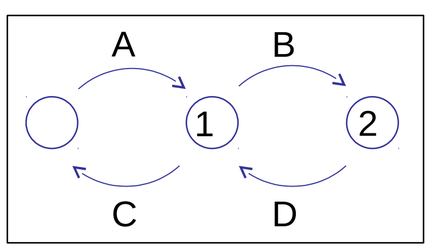

Network Architecture (input)

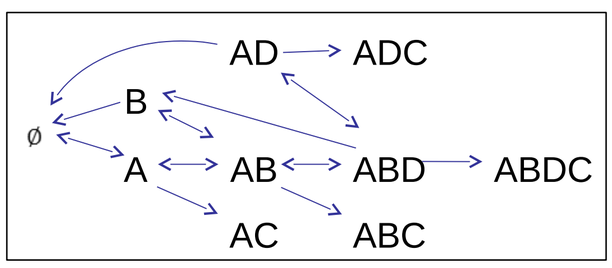

Attacker State Graph (internal representation)

Game Model Overview:

- Attacker compromises network edges (credentials)
- Attacker only wins reward when data exfiltration is possible
- Defender resets nodes to clear compromise
- Defender loses proportional to the amount exfiltrated
- Zero-sum: attacker tries to maximize data exfiltration, defender minimizes

Simulation Results

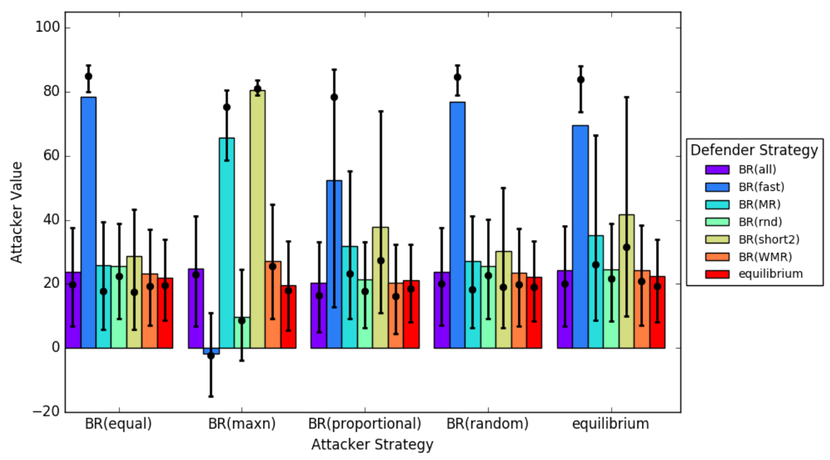

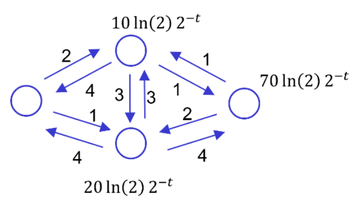

- A player's equilibrium strategy defends against their opponent learning their strategy and optimizing against it
- Simulating 10k trials shows a large variance of attacker reward
- Defender should play equilibrium since any other strategy opens them up to risk of attacker learning and optimizing

Use Case

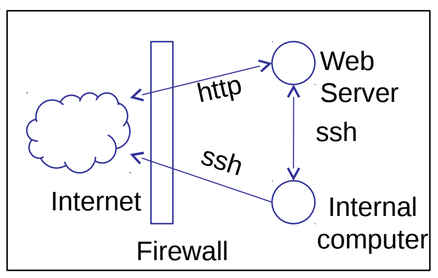

use-case network

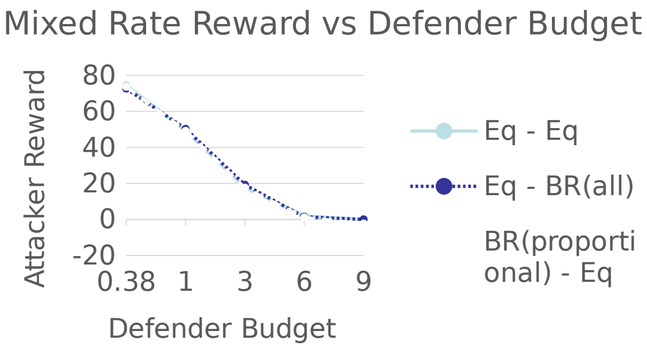

- Tested on a small model enterprise network
- Model parameters were chosen through literature survey
- If a defender's budget is large enough, a rational attacker should not attempt to compromise the network even if there are vulnerabilities

Defense-In-Depth

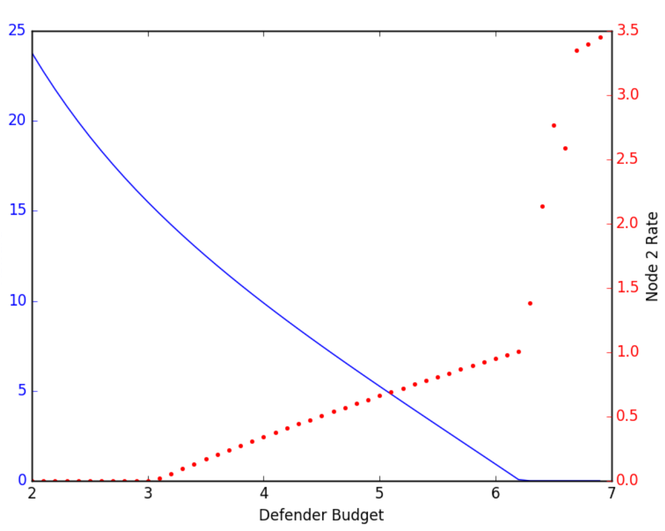

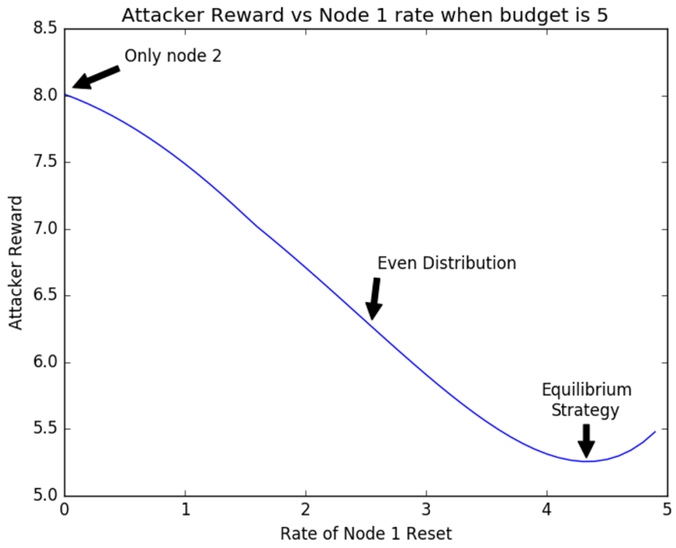

- Defense-in-Depth may not make sense when the defender is severely budget constrained
- Securing the perimeter should be the primary concern
- Perimeter defense can be aided by defense-in-depth if budget allows
- Monitoring outgoing connections is just as important as restricting incoming connections

Related Publications

1. M. Van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "Flipit: The game of "stealthy takeover"," Journal of Cryptology, vol. 26, no. 4, pp. 655–713, 2013.
2. A. Laszka, G. Horvath, M. Felegyhazi, and L. Butty´ n, "Flipthem: Modeling targeted attacks with flipit for multiple resources," in International Conference on Decision and Game Theory for Security. Springer, 2014, pp. 175–194.
3. H. Holm, K. Shahzad, M. Buschle, and M. Ekstedt, "P2 cysemol: Predictive, probabilistic cyber security modeling language," IEEE Transactions on Dependable and Secure Computing, vol. 12, no. 6, pp. 626–639, Nov 2015.

Sponsored by Applied Research Laboratory, The Pennsylvania State University[1] and Army Research Lab Cybersecurity CRA[2]