

Security and Network Effects: Centralized and Decentralized Perspectives

YEVGENIY VOROBAYCHIK

Sandia National Laboratories, Livermore, CA¹

Security, like many other complex decisions, is generally approached with a divide-and-conquer mindset. Consequences of security failures, however, can rarely be completely localized: an explosion or a fire at one building can affect neighboring structures, a debt crisis in Greece resonates throughout the tightly connected European and US financial markets, and a breach of security at one computer can facilitate access to others on the same network. It is thus crucial to view security holistically, and devise security strategies that explicitly account for interdependencies between valuable assets. Here we provide an overview of two recent approaches to security with network effects. The first approach takes a centralized perspective, attempting to compute an optimal security configuration for all interdependent assets. This approach explicitly accounts for an intelligent adversary optimally attacking one of the assets. The second approach studies the impact of decentralized decision making when local failures can propagate in complex ways through the entire system, but assumes that initial failures are random.

Categories and Subject Descriptors: I.2.11 [**Artificial Intelligence**]: Distributed artificial intelligence—*Intelligent agents*

General Terms: Algorithms, Performance, Economics, Security

Additional Key Words and Phrases: Game theory, Security, Stackelberg Games, Networks

1. MOTIVATION

Security-related decisions usually take place in a complex, dynamic, highly interdependent environment. Typical mathematical modeling of security naturally abstracts away most of the complexity. In particular, a common simplification is to ignore interdependencies between assets to secure, or parties that make security decisions. Recently, Kunreuther and Heal proposed a rather influential model that captures, in a highly stylized way, interdependencies between security decisions of different parties [Kunreuther and Heal 2003]. While Kunreuther and Heal (and follow-up work) explicitly account for interdependencies between security decisions of different players, they only consider a binary action space for each player: to secure, or not. Thus, they abstract away an important aspect of security: each decision maker is often responsible for securing a collection of interdependent assets. Each individual decision, in isolation of game theoretic interactions, is one rife with complexities, while the strategic aspect adds yet another dimension.

Below I describe two recent approaches that explicitly capture the complexity of interdependent security decisions, both at the individual (centralized) level, and accounting for strategic behavior of multiple self-interested parties.

¹Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

2. A CENTRALIZED APPROACH TO INTERDEPENDENT SECURITY

In this section we describe a centralized approach to interdependent security by Letchford and Vorobeychik [2011]. Consider a collection of valuable assets (targets), T , which must be defended from a rational attacker. The defender's options involve security configurations for each target (e.g., firewall settings), and a tradeoff between highly effective, but costly, security options (such as a highly restrictive firewall setting that prevents valuable operations from being performed), and cheaper security options which are more fragile. An attacker is endowed with a power to observe the defender's decision, and then execute a single attack against an asset which yields the greatest gain. The defender, however, has an important ally: randomization. Specifically, the defender can "commit" to a randomized strategy by which security configurations are chosen, and the attacker can only observe the stochastic security configuration, but not its ultimate realizations.

While the model described so far is similar to a typical Stackelberg security game setting [Conitzer and Sandholm 2006; Paruchuri et al. 2008; Kiekintveld et al. 2009], Letchford and Vorobeychik [2011] extend this line of work by explicitly representing interdependencies between assets as a graph, $G = (T, E)$, where assets are identified with nodes in the graph, and the set of edges E represents interdependencies between them. Moreover, each asset t has an *intrinsic* value w_t , which is lost to the defender if this asset, and only this asset, is compromised by the attacker (symmetrically, the attacker would gain an intrinsic value v_t from compromising or destroying t). Interdependencies are modeled as independent failure cascades: if an asset t fails (is successfully attacked), its network neighbors t' also fail, independently, each with probability $p_{t,t'}$.

An important assumption in past work on Stackelberg security games has been that security decisions are independent across assets. This assumption is clearly violated in the most general incarnation of the model just described. However, under the assumption that the cascade probabilities do not depend on security configurations, we can attain *effective independence* by simulating failure cascades initiated at each asset in T and quantifying the resulting expected utilities to the defender and attacker. Stated more generally, the requisite assumption is that security decisions targeting external threats are independent of security decisions targeting insider threats (e.g., security threats originating from other computers on a local network), a state of affairs that seems ubiquitous in network security settings. Significantly, we can subsequently use a linear programming approach to compute optimal randomized security policies that accounts for cascading failures.

Letchford and Vorobeychik perform a computational analysis of several interdependent security settings. They observe that total defense expenditures exhibit a single peak as a function of cost in graphs with a relatively homogeneous degree distribution, but two peaks in scale-free graphs. They additionally demonstrate that resilience properties of graphs have significantly different characteristics from those observed when defense decisions are not taken into account.

3. ROBUSTNESS, FRAGILITY, AND DECENTRALIZATION

The approach described in the previous section assumes that the defense decision is entirely centralized. A complementary approach in Vorobeychik *et al.* [2011]

presents a model of security decisions that accounts for interactions between security decisions of multiple players, although this alternative model presumes that the failures happen according to a fixed distribution, rather than as targeted attacks. The Vorobeychik *et al.* model also aims to gain fundamental insights about security in complex systems, rather than to provide a framework that can capture realistic interdependent security settings. As such, the complex interdependencies between security decisions arise based on a very abstract and simplified model of forest fires. The forest fire model starts with a square ($N \times N$) grid. Each cell in this grid is encoded by a binary value, where 1 indicates a presence of a tree in that cell. The grid is partitioned among a set of players, each deciding which cells, among those he owns, are to contain a tree; for every tree a player plants he pays a fixed cost c . The catch is that after a joint decision is made by all players, a lightning can strike any cell in the grid according to a predefined spacial distribution, burning down the entire connected component to which this cell belongs in the process. The goal of each player is to maximize the total *yield*, or expected number of trees he plants that survive the lightning, less total planting costs.

The complexity of this game theoretic forest fire model precludes mathematical analysis of an arbitrary instantiation. However, it is not difficult to obtain insight in the boundary cases where there is either a single player controlling the entire grid, or where each player controls a single cell. In the former case, clearly, socially optimal solution exactly matches the “equilibrium” configuration. In the latter, Vorobeychik *et al.* show that equilibrium solutions can be arbitrarily poor. Simulation-based game theoretic analysis reveals an interesting pattern when the number of players is between these extremes: there is a level of moderate decentralization where equilibrium configurations are, on average, close to socially optimal, and additionally exhibit high resilience to changes in the lightning distribution. In contrast, highly centralized solutions are extremely fragile to environment changes. The high-level reason why decentralization yields greater resilience is that individually, players’ decisions must account both for the negative events that impact them directly, as well as the spread of fire due to the selfish choices of their neighbors. Thus, players build in extra robustness into their configurations that is absent in a highly centralized decision.

REFERENCES

- CONITZER, V. AND SANDHOLM, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*. EC ’06. ACM, New York, NY, USA, 82–90.
- KIEKINTVELD, C., JAIN, M., TSAI, J., PITA, J., ORDÓÑEZ, F., AND TAMBE, M. 2009. Computing optimal randomized resource allocations for massive security games. In *In AAMAS-09*.
- KUNREUTHER, H. AND HEAL, G. 2003. Interdependent security. *Journal of Risk and Uncertainty* 26, 2-3, 231–249.
- LETCHFORD, J. AND VOROBAYCHIK, Y. 2011. Computing randomized security strategies in networked domains. In *AARM Workshop*.
- PARUCHURI, P., PEARCE, J. P., MARECKI, J., TAMBE, M., ORDONEZ, F., AND KRAUS, S. 2008. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *Seventh International Conference on Autonomous Agents and Multiagent Systems*. 895–902.
- VOROBAYCHIK, Y., MAYO, J. R., ARMSTRONG, R. C., AND RUTHRUF, J. R. 2011. Noncooperatively optimized tolerance: Decentralized strategic optimization in complex systems. *Physical Review Letters* 107, 10, 108702.