Game-Theoretic Approach to Security Problems

Antonín Kučera

Joint work with Tomáš Brázdil, David Klaška, Tomáš Lamser, and Vojtěch Řehák

(IJCAI 2018, AAMAS 2018)

Brno, September 2018

Security Problems

- One of the basic problems in operations research, heavily studied by AI community.
- How to use the (limited) security sources to achieve the best coverage of a given set of vulnerable targets?
- Many technical variants: static allocation of security resources, mobile partrollers/attackers, various levels of target importance/vulnerability, etc.
- Popular solution concept: Stackelberg equilibrium
  - The leader commits to a strategy and the follower chooses his best response so that they cannot gain anything by revising their choice.
  - The defender/attacker correspond to the leader/follower.

Adversarial Patrolling Problem

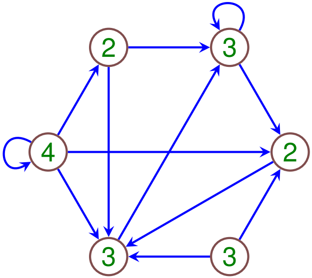

- Defender's strategy: $\sigma : V^{+} \to \Delta(V)$
- Attacker's strategy: $\pi : V^{+} \to V \cup \{*\}$ (must be "prefix free")
- $\mathcal{P}^{\sigma,\pi}(DRuns)$
- $val = \sup_{\sigma} \inf_{\pi} \mathcal{P}^{\sigma,\pi}(DRuns)$
- Optimal Defender's strategy exists.

Patrolling in a General Environment

- Deciding whether $val = 1$ or $val \leq 1 - \frac{1}{n}$ is NP-hard.
- There is an exponential-time algorithm for computing ε-optimal strategies.
- Existing strategy synthesis algorithms are mostly based on (non)linear programming and often compute only positional strategies for games with hundreds of vertices.

Patrolling in the Internet Environment

- The graph is fully connected.
- The number of targets can reach millions/billions.
- The Defender's are software processes run by a central authority (they are fully coordinated).
- The targets have different importance
- Intrusion detection is not perfect.

Patrolling in the Internet Environment (2)

- In the Internet patrolling, we can compute (sub)optimal strategies for k Defenders quickly for VERY large instances.
- Furthermore, we can quickly determine the number of Defenders needed to achieve a given level of protection.

Patrolling in the Internet Environment (2)

- In the Internet patrolling, we can compute (sub)optimal strategies for k Defenders quickly for VERY large instances.
- Furthermore, we can quickly determine the number of Defenders needed to achieve a given level of protection.

Key new concepts:

- Modular strategies.
- A suitable (de)composition principle.
- The use of mathematical programming is completely avoided. We need to solve a certain system of non-linear equations.

Modular Strategies

- A Defender's strategy σ is modular if $\sigma(h)$ depends only on $|h| \bmod c$ where c is a suitable integer. Hence, a modular strategy can be seen as a function with domain $\mathbb{N}$.
- In particular, modular strategies are independent of the current Defender's position (the currently visited vertex/vertices). Hence, modular strategies do not subsume positional strategies.
- Intuitively, modular strategies appear weak. This intuition is incorrect.

The (De)Composition Principle

- Suppose there is only one Defender.
- Let $G_1, \ldots, G_\ell$ be fully connected patrolling graphs.
- Suppose we already computed a modular Defender's strategy σ_i for every G_i.
- Let η be a "suitable" distribution over $\{1, \ldots, n\}$.
- We can compose the modular strategies $\sigma_1, \ldots, \sigma_n$ into a modular strategy σ for $G_1, \ldots, G_n$ as follows:

  $\sigma(\ell) =$ *A "ν-combination" of* $\sigma_1(\ell), \ldots, \sigma_n(\ell)$

- For k Defenders, we first need to "assign" them to $G_1, \ldots, G_n$, i.e., choose $k_1, \ldots, k_n$ such that $\sum_{i=1}^{k} k_i = k$, and solve G_i for k_i Defenders.
- We can give a lower bound on val_σ based on $val(\sigma_1), \ldots, val(\sigma_n)$.

Example 1

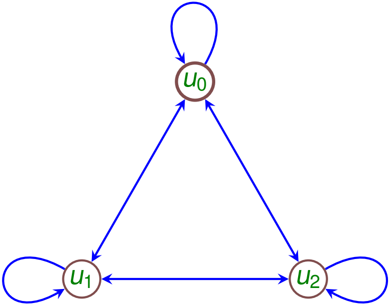

Attack length $=$ 2

Example 1

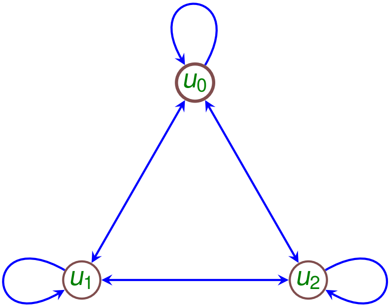

Attack length $= \ 2$

$\sigma(h) = \mu_\ell, \ \ \ell = |h| \, mod \, 2$

$\mu_0(u_0) = 0,$
$\mu_0(u_1) = \kappa,$
$\mu_0(u_2) = 1 - \kappa$

$\mu_1(u_0) = \kappa,$
$\mu_1(u_1) = 0,$
$\mu_1(u_2) = 1 - \kappa$

$\kappa = (\sqrt{5} - 1)/2 = 0.618\ldots$

Example 2

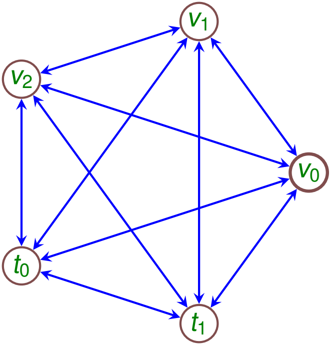

$d(t_i)=2, d(v_i)=3$

Example 2

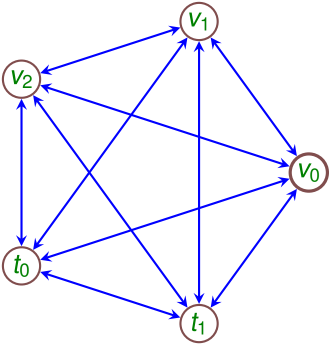

$d(t_i) = 2, d(v_i) = 3$

$\sigma(h)$ selects uniformly between $v_{|h|+1 \; mod \, 3}$ and $t_{|h|+1 \; mod \, 2}$

$val^{\sigma} = 1/2$

$val = 1/2$

An Upper Bound on the Value

- We give an upper bound on the achievable value which can be computed "quickly" for a given patrolling problem.
- This bound is not tight in general, but can serve as a "yardstick" for measuring the quality of constructed strategies.

Our Algorithm

- We design a concrete strategy synthesis algorithm by designing a suitable decomposition tactic.
- Computing appropriate “mixing ratios” for the modular strategies constructed for the subgames requires solving a system of non-linear equations, which is done by Maple.
- The algorithm can solve instances with billions of vertices and thousands of Defenders in seconds.
- The value of the produced strategies matches the principal bound in some well-defined cases.
- If the intrusion times are taken from a fixed finite set of eligible values, then the values of the constructed strategies approach the upper bound very quickly as the number of targets increases.

Future Work

- What is precise complexity of the patrolling problem in the Internet environment?
- Can we compute (a symbolic representaion of) optimal strategies for all instances?
- Can we solve other types of games compositionaly?