Game-Theoretic Models of Patrolling Problems

Antonín Kučera

(joint work with Tomáš Brázdil, David Klaška, Tomáš Lamser, Vojtěch Řehák)

GAMENET Workshop 2019

Game-Theoretic Models of Security Problems

- How to utilize (limited) security resources to achieve the best protection of a given set of targets?
- Formalized as a game between the Defender and the Attacker.
- Solution concept: Stackelberg equilibrium.
  - The leader (Defender) commits to a strategy.
  - The follower (Attacker) chooses a best response.
  - No player is motivated to change the decision.

Patrolling Problem

- The Defender travels among a given set of vulnerable targets and aims at detecting possible attacks initiated by the Attacker.
- Completing an attack at a given target takes finite time and the Defender must visit the target to discover an ongoing attack.
- Many technical variants:
  - The number of patrolling agents and their coordination.
  - Attacker's abilities (adversarial vs. non-adversarial models).
  - Target importance.
  - Other parameters...

Non-Adversarial Patrolling

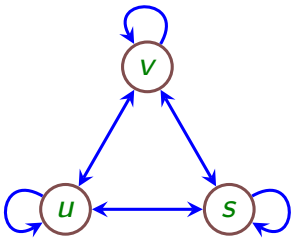

Attack length $=$ 2

- The Attacker cannot see the Defender but knows his strategy.
  - $\sigma : V^* \to \mathcal{D}(V)$, $\pi : \mathbb{N} \to V \cup \{*\}$
  - $w = v_1, v_2, v_3, \ldots$
  - $\mathit{Defend}_\pi(w)$ is either 1 or 0 depending on whether the attack scheduled by π is discovered along w or not, respectively.
  - $\mathit{Val} = \sup_\sigma \inf_\pi \mathbb{E}^\sigma[\mathit{Defend}_\pi]$.
- $\mathit{Val} = 2/3$; optimal strategies for the Defender:
  - Pick an initial target randomly, then walk around the targets.
  - Pick an initial target randomly, stay there.

Adversarial Patrolling

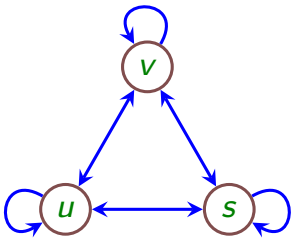

Attack length $=$ 2

- The Attacker can see the Defender and knows his strategy.
  - $\sigma : V^* \to \mathcal{D}(V)$, $\quad \pi : V^* \to V \cup \{*\}$
  - $Val \;=\; \sup_\sigma \inf_\pi \mathbb{E}^\sigma[Defend_\pi]$.
- $Val = ?$
- Optimal strategy?
- As we shall see, $Val = \frac{\sqrt{5}-1}{2}$, and an optimal strategy exists.

Adversarial Patrolling – The Model

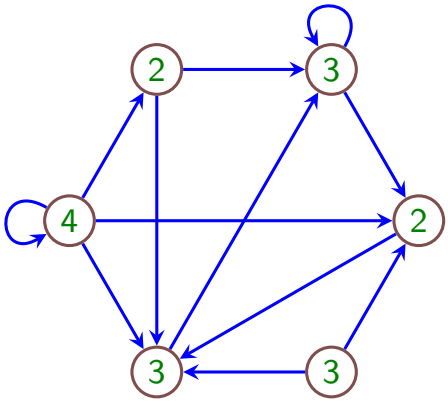

- The targets may have different weights modeling their importance.
- Traversing an edge may take more than one time unit.
- Detecting an attack may be imperfect.
- The number of patrollers may be larger than one.

Adversarial Patrolling – Basic Questions

- Do optimal Defender's strategies always exist? If so, do they require memory/randomization?
- Can we compute the value and an optimal strategy for the Defender?

Adversarial Patrolling – Optimal Strategies

Theorem 1

An optimal Defender's strategy always exists. It may require both memory and randomization.

- Let $\sigma_1, \sigma_2, \ldots$ be a sequence of Defender's strategies such that $\lim_{n\to\infty} \mathit{Val}(\sigma_n) = \mathit{Val}$.
- An optimal strategy is constructed by repeatedly selecting a subsequences of strategies that are pointwise converging for more and more points and taking the corresponding limit distributions.

Adversarial Patrolling – Suboptimal Strategies

Theorem 2

For every $\varepsilon > 0$, there exists an ε-optimal Defender's strategy with finite memory.

- Let σ be a Defender's strategy and $v_1, \ldots, v_n$ a history.
- The behaviour of $\sigma(v_1, \ldots, v_n)$ is characterized by the tuple of all

  $$\mathcal{P}(Reach^{=k}(u))$$

  where u is a target and $0 \leq k \leq d_{\max}$
- If $v_1, \ldots, v_n$ is a history whose characteristic tuple is "close enough" to the characteristic tuple of some prefix $v_1, \ldots, v_m$, then σ can start "copying" the behavior after the prefix.

Adversarial Patrolling – Optimal Strategies (2)

- Finite-memory strategies are sufficient for ε-optimality. The question whether they suffice for optimality is open.
- There may exist an optimal Defender's strategy σ such that the Attacker has no best response to σ.

Computing Optimal Positional Strategies

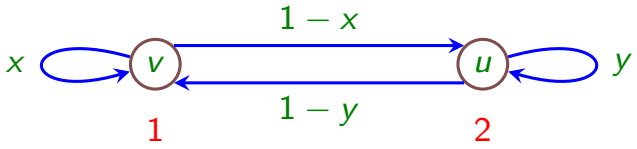

position	attack	defending probability
v	v	x
v	u	$(1-x)+x(1-x)$
u	v	$1-y$
u	u	$y+(1-y)(1-x)$

maximize α **subject to**

$$\begin{aligned} \alpha &\leq x \\ \alpha &\leq (1-x)+x(1-x) \\ \alpha &\leq 1-y \\ \alpha &\leq y+(1-y)(1-x) \\ 0 &\leq x \leq 1 \\ 0 &\leq y \leq 1 \end{aligned}$$

$\alpha = \frac{\sqrt{5}-1}{2}$

Computing Optimal Strategies

Theorem 3

Deciding whether $val = 1$ or $val \leq 1 - \frac{1}{n}$ is NP-hard.

- A trivial reduction from the Hamiltonian cycle problem.

Theorem 4

There is an exponential-time algorithm for computing ε-optimal strategies.

- A proof is non-trivial.

Computing Optimal Strategies (2)

- The computability of *Val* and an optimal Defender's strategy for a general patrolling graph is open.
- The existing works consider special topologies (circles, lines, etc.) or provide strategy-synthesis algorithms for general topologies with no optimality guarantees.

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)

Main results:

- In the Internet environment, (sub)optimal strategies for k patrollers can be computed quickly even for very large instances.
- Furthermore, the number of patrollers needed to achieve a given level of protection can be quickly determined.

Key concepts:

- Modular strategies.
- A suitable (de)composition principle.
- The use of mathematical programming is completely avoided. A certain system of non-linear equations needs to be solved.

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

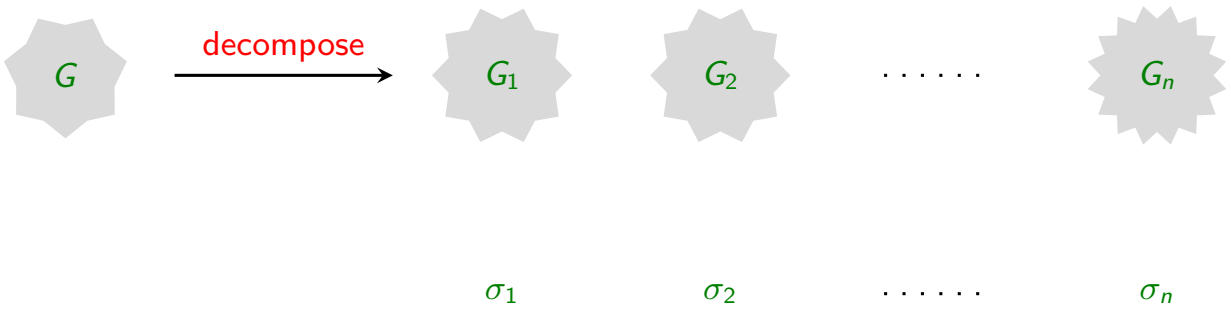

compose:

- Let η be a "suitable" distribution over $\{1, \ldots, n\}$.
- Put $\sigma = \eta[\sigma_1, \ldots, \sigma_n]$

Example 1

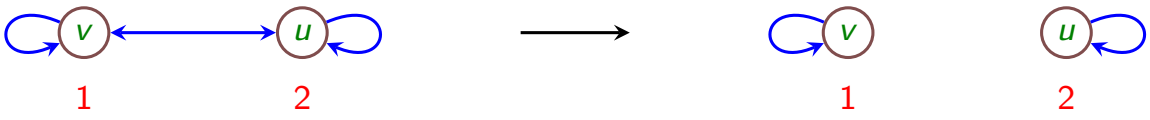

- σ_1, σ_2 are the two trivial strategies for G_1, G_2.
- How to determine the distribution η?
  - Suppose $\eta(G_1) = x$ and $\eta(G_2) = 1 - x$.
  - v are covered with prob. x, while u with prob. $1 - x^2$.
  - Setting $x = 1 - x^2$ yields $x = \frac{\sqrt{5}-1}{2}$.
- The obtained modular strategy $\eta[\sigma_1, \sigma_2]$ is optimal.

Example 2

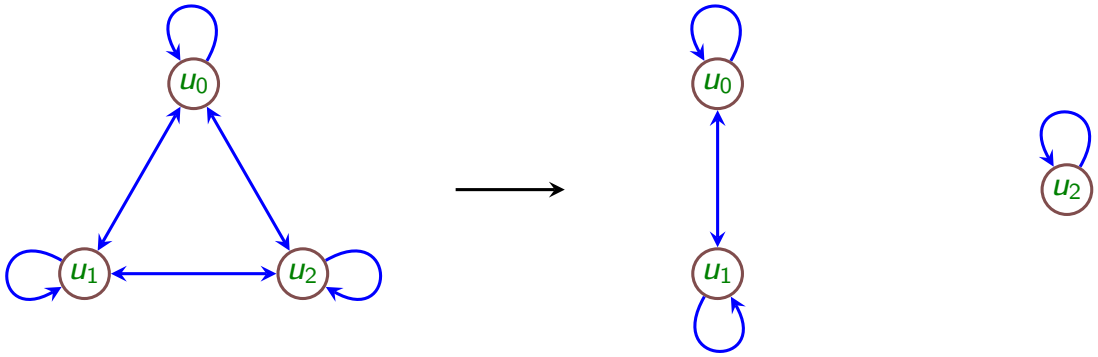

Attack length $=$ 2

- σ_1, σ_2 are the two natural strategies for G_1, G_2.
- How to determine the distribution η?
  - Suppose $\eta(G_1) = x$ and $\eta(G_2) = 1 - x$.
  - u_0, u_1 are covered with prob. x, while u_2 with prob. $1 - x^2$.
  - Setting $x = 1 - x^2$ yields $x = \frac{\sqrt{5}-1}{2}$.
- The obtained modular strategy $\eta[\sigma_1, \sigma_2]$ is optimal.

Example 3

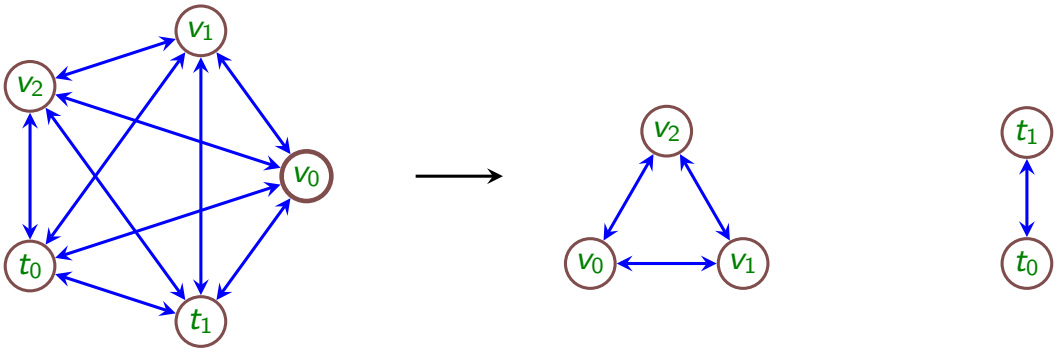

$d(t_i)$=2, $d(v_i)$=3

- σ_1, σ_2 are the two natural strategies for G_1, G_2.
- If $\eta(G_1) = x$ and $\eta(G_2) = 1 - x$, then v_0, v_1, v_2 are covered with prob. x, and t_0, t_1 with prob. $1 - x$.
- Setting $x = 1 - x$ yields $x = \frac{1}{2}$.
- The obtained modular strategy $\eta[\sigma_1, \sigma_2]$ is again optimal.

An Upper Bound on the Value

- There is 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.

A Strategy Synthesis Algorithm

- A concrete strategy synthesis algorithm is obtained 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.

Patrolling in a General Environment

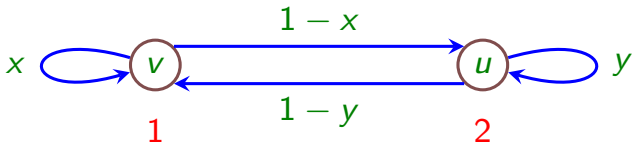

position	attack	defending probability
v	v	x
v	u	$(1-x)+x(1-x)$
u	v	$1-y$
u	u	$y+(1-y)(1-x)$

- $F(x,y) = \min\{x,\ (1-x)+x(1-x),\ 1-y,\ y+(1-y)(1-x)\}$
- Maximize $F(x,y)$ in $\langle 0,1\rangle \times \langle 0,1\rangle$.
- A local maximum can be found by a modified gradient ascent method.

Conclusions

- Some fundamental questions about patrolling problems are still unresolved.
- The main obstacle is the computational hardness of the considered problems, but this can be overcome in various ways.
- There are many challenging directions for future research:
  - Coordinating the patrollers.
  - Resilience.
  - Patrollers with bounded energy resources (drones).
  - ...