Evolutionary Approach to Security Games with Signaling

Adam Żychowski [1] Jacek Mańdziuk [1] Elizabeth Bondi [2] Aravind Venugopal [3] Milind Tambe [2] Balaraman Ravindran [3,4]

[1]Faculty of Mathematics and Information Science, Warsaw University of Technology [2]Center for Research on Computation and Society, Harvard University [3]Robert Bosch Centre for Data Science and AI, IIT Madras [4]Department of Computer Science and Engineering, IIT Madras

Problem definition

Security Games with Signaling

Inspiration: prevention of poaching in Africa.

2 players: **Defender** and **Attacker**
Defender's units: patrollers, drones

Drone can send one of the following signals:
- weak – sending information to patrollers about attack detection
- strong – sending information about attack and launch sound/light signals to deter the Attacker

Games on graph – each vertex is target with a set of payoffs.

Defender's strategy: assigning patrollers and drones to targets, signaling strategy.
Attacker's strategy: target to attack, signaling reaction.

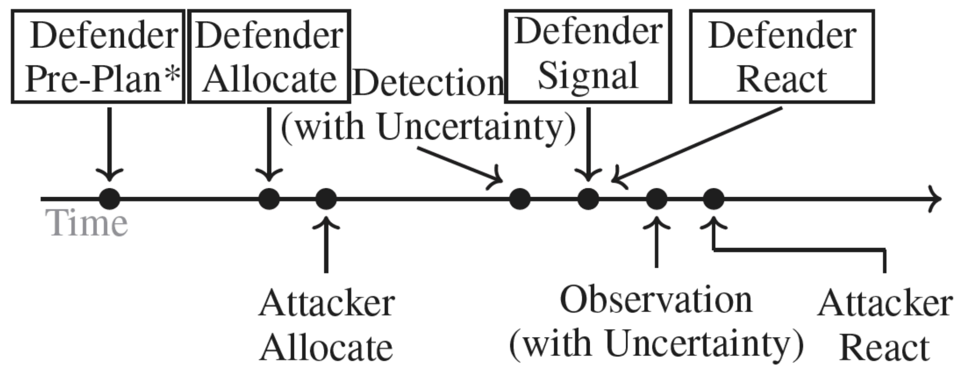

Stackelberg Equilibrium

Defender commits to his/her strategy first.
Attacker, knowing the Defender's strategy, chooses his/her strategy.
Defender always commits to a mixed strategy.

Stackelberg equilibrium: a pair of players' strategies, for which strategy change by any of players leads to his/her result deterioration.

$$(\pi_D^*, R(\pi_D^*)) \in \Pi_D \times \Pi_A$$

$$\pi_D^* = \text{argmax}_{\pi_D \in \Pi_D} U_D(\pi_D, R(\pi_D))$$

$$R(\pi_D) = \text{argmax}_{\pi_A \in \Pi_A} U_A(\pi_D, \pi_A)$$

$G \in \{D, A\}$ – players (Defender, Attacker)
Π_G – a set of player's G all mixed strategies
U_G – payoff of player G

Game uncertainties

Detection uncertainty
A drone may not detect the Attacker even if they are both located in the same target (e.g. conservation drone imagery may be imperfect, particularly given occlusions such as trees).

Observational uncertainty
The Attacker observes different signal (also no signal) according to matrix Ω due to potential occlusions or difficulties viewing the true signal.

$P[y|x]$ – probability of recognizing signal x under condition of the true signal y.

$$\Omega = \begin{bmatrix} P[n|n] & P[n|\sigma_0] & P[n|\sigma_1] \\ P[\sigma_0|n] & P[\sigma_0|\sigma_0] & P[\sigma_0|\sigma_1] \\ P[\sigma_1|n] & P[\sigma_1|\sigma_0] & P[\sigma_1|\sigma_1] \end{bmatrix}$$

Evolutionary algorithm for Security Games with Signaling (EASGS)

Solutions encoding

$$CH_j = \{(e_1^j, q_1^j), \dots, (e_i^j, q_i^j), \dots, (e_{d_j}^j, q_{d_j}^j), \mathbf{\Psi}_j^{\theta}, \mathbf{\Phi}_j^{\theta}\}$$

$e = (V_p, V_s, V_r)$ - pure strategy
V_p - a set of vertices with assigned patrollers,
V_s - a set of vertices with assigned drones,
V_r - reallocation plan, a set of vertices (connected with V_p), to which each patroller moves if no adversaries are observed.

$q_i^j \in [0, 1]$ is the probability of playing strategy e_i^j, $\sum_{i=1}^{d_j} q_i^j = 1$.

$\theta \in \{\bar{s}, s^+, s^-\}$ - drones allocation states:
$\bar{s}$ - no patroller is in the drone's neighbourhood,
s^+ - a patroller is planned to visit drone's vertex in the reaction stage,
s^- - no patroller will visit drone's vertex in the reaction stage but there is at least one patroller in neighbourhood who can respond

$\mathbf{\Psi}_j^{\theta} = [\Psi_{j,1}^{\theta}, \Psi_{j,2}^{\theta}, \dots, \Psi_{j,\mathcal{N}}^{\theta}]$ - signaling strategy in case of attack detection
$\mathbf{\Phi}_j^{\theta} = [\Phi_{j,1}^{\theta}, \Phi_{j,2}^{\theta}, \dots, \Phi_{j,\mathcal{N}}^{\theta}]$ - signaling strategy in case of no attack detection

Evolutionary operators

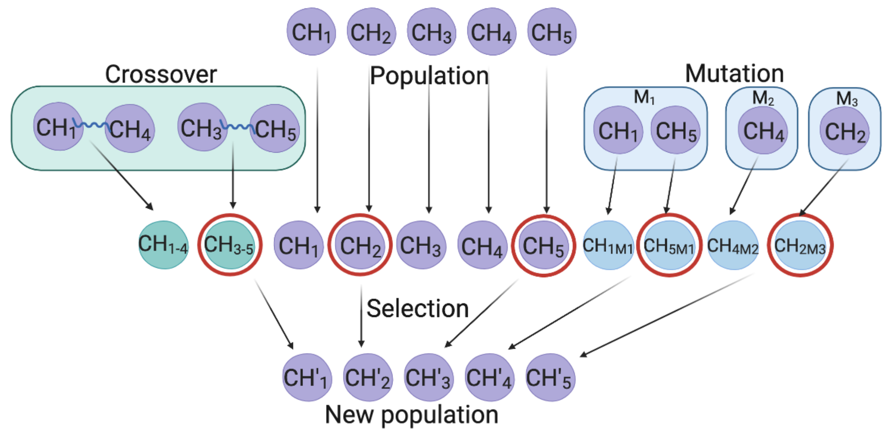

3 mutation types:
- random allocation/reallocation modification,
- random probability change,
- coverage improvement.

Crossover combines pure strategies with halved probabilities, averaging signaling probabilities.
Evaluation based on game rules (including detection and observational uncertainties).

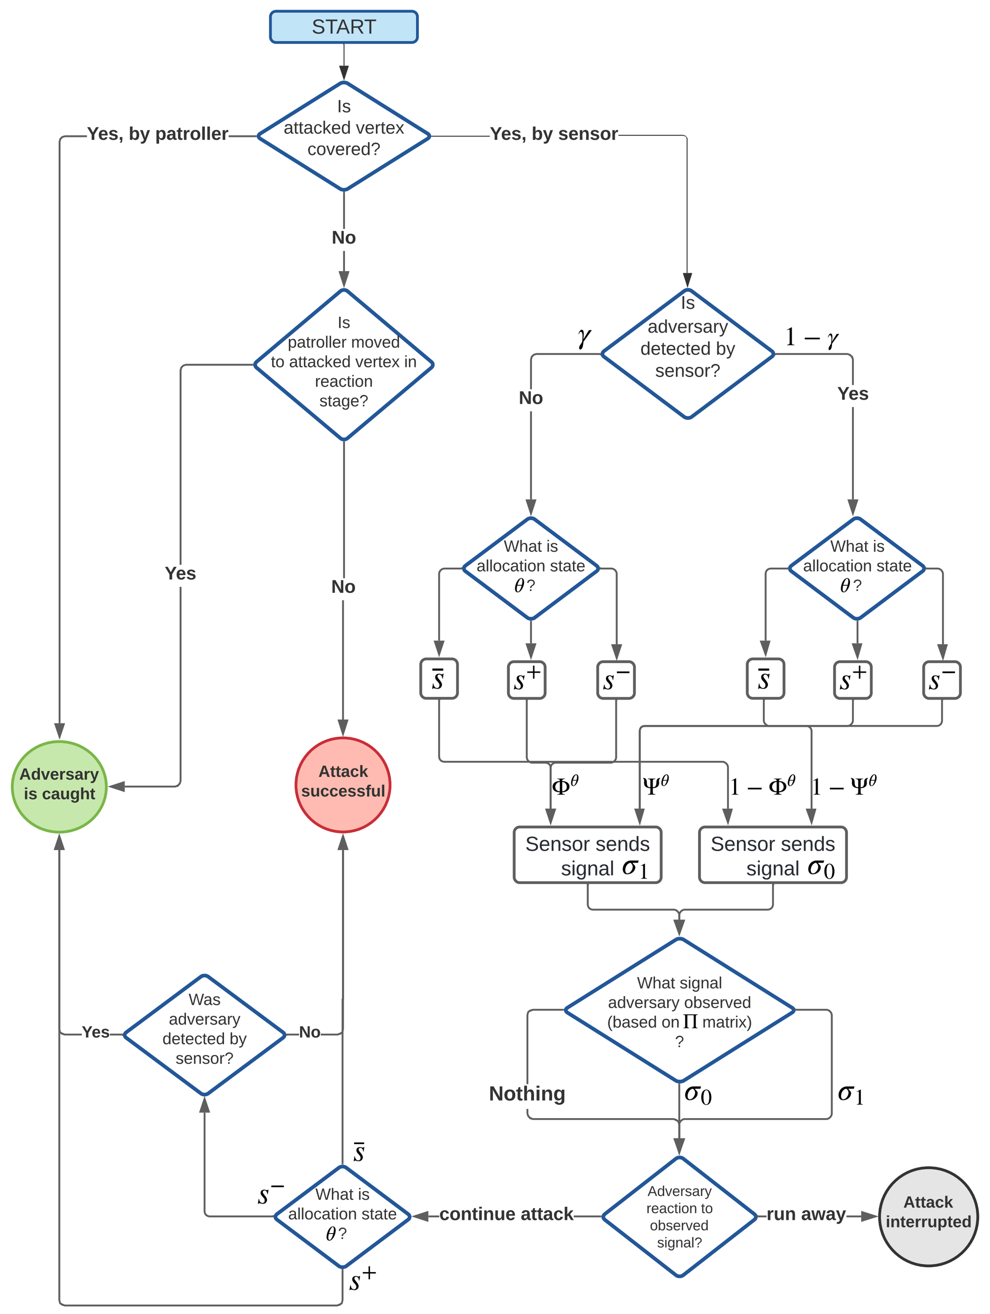

Benchmark games

342 games with different graph topologies:
- *sparse* (avg deg = 2) – 50 games
- *moderate* (avg deg = $\frac{n}{2}$) – 50 games
- *dense* (avg deg = $n-2$) – 50 games
- *locally* dense (connected cliques) – 192 games

number of vertices: $n \in [10, 100]$
number of patrollers: $k_s = \sqrt{\frac{n}{2}}$
number of drones: $k_d = \frac{2}{3}n - k_s$

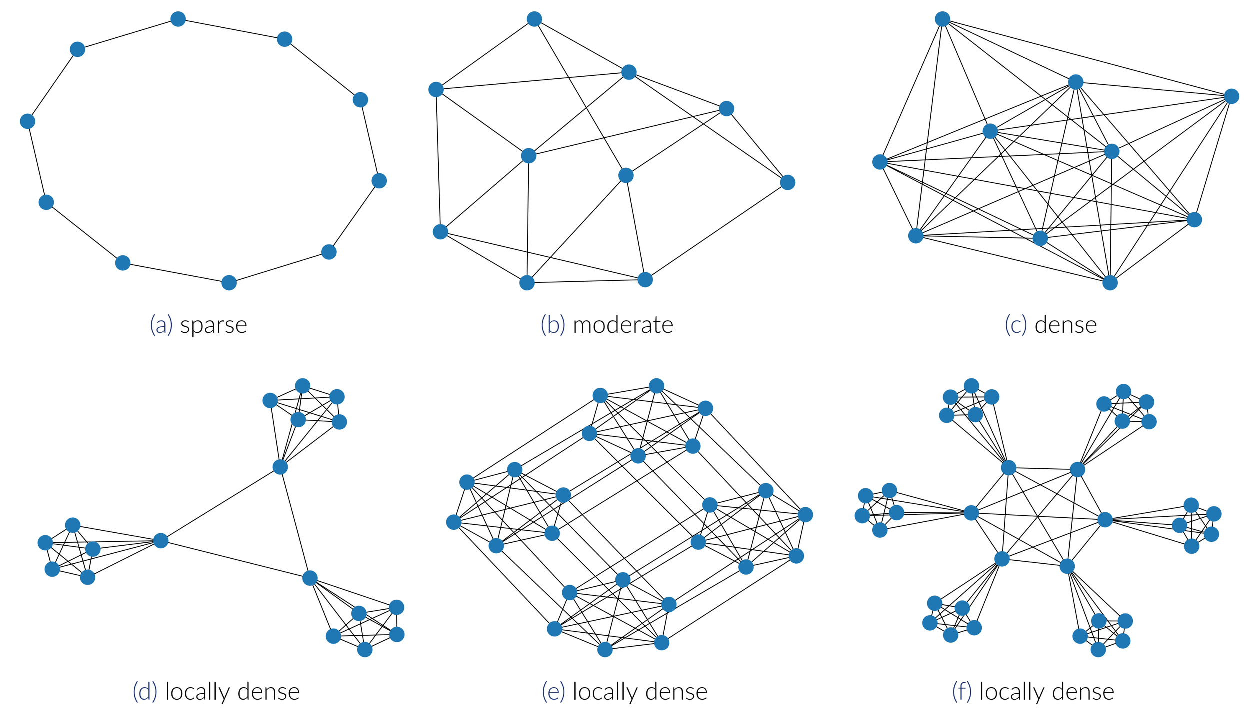

Results

EASGS obtained the best result for 200 out of 342 games.

	SBP	SBP+W	m-CombSGPO	EASGS
sparse	-86.68 (84%)	**-86.01 (92%)**	-419.86 (0%)	-91.32 (6%)
moderate	-75.01 (2%)	-72.75 (36%)	-255.73 (0%)	**-69.92 (62%)**
dense	-58.72 (2%)	-57.98 (34%)	-149.14 (0%)	**-51.47 (64%)**
locally-dense	-60.68 (4%)	-57.80 (26%)	-340.65 (0%)	**-54.36 (70%)**

Table 1. Averaged Defender's payoff across all benchmark games.

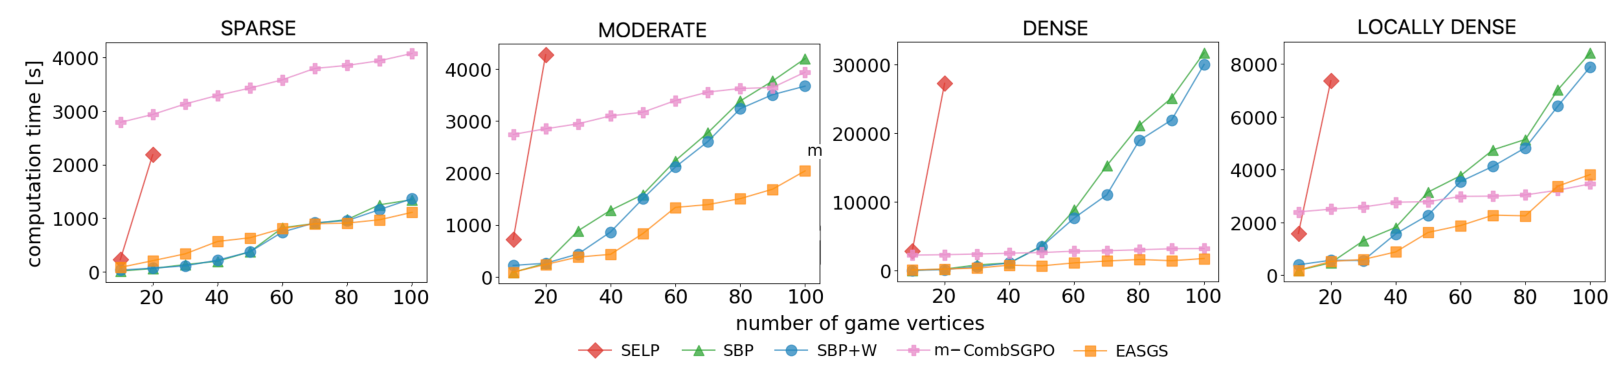

Figure 1. Time scalability.

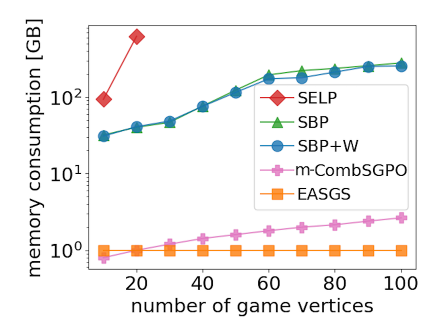

Figure 2. Memory consumption.

Conclusions

- new evolutionary method for Security Games with Signaling
- results close to optimal
- much better time and memory scalability than competitive methods
- viable alternative to exact method and state-of-the-art heuristics

a.zychowski@mini.pw.edu.pl