Computing Stackelberg Equilibrium

Branislav Bošanský

Artificial Intelligence Center,
Department of Computer Science,
Faculty of Electrical Engineering,
Czech Technical University in Prague

branislav.bosansky@agents.fel.cvut.cz

March 18, 2019

Stackelberg Equilibrium

Players have different roles in the Stackelberg solution concept:

Stackelberg Equilibrium

Players have different roles in the Stackelberg solution concept:

- *the leader* – publicly commits to a strategy

Stackelberg Equilibrium

Players have different roles in the Stackelberg solution concept:

- *the leader* – publicly commits to a strategy
- *the follower(s)* – play a Nash equilibrium with respect to the commitment of the leader

Stackelberg Equilibrium

Players have different roles in the Stackelberg solution concept:

- *the leader* – publicly commits to a strategy
- *the follower(s)* – play a Nash equilibrium with respect to the commitment of the leader

Stackelberg equilibrium is a strategy profile that satisfies the above conditions and maximizes the expected utility value of the leader:

Stackelberg Equilibrium

Players have different roles in the Stackelberg solution concept:

- *the leader* – publicly commits to a strategy
- *the follower(s)* – play a Nash equilibrium with respect to the commitment of the leader

Stackelberg equilibrium is a strategy profile that satisfies the above conditions and maximizes the expected utility value of the leader:

$$\arg\max_{\sigma \in \Sigma; \forall i \in \mathcal{N} \setminus \{1\} \sigma_i \in BR_i(\sigma_{-i})} u_1(\sigma)$$

There may be multiple Nash equilibria

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule
- *Strong SE* – the followers select such NE that maximizes the outcome of the leader (when the tie-braking is not specified we mean SSE),

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule
- *Strong SE* – the followers select such NE that maximizes the outcome of the leader (when the tie-braking is not specified we mean SSE),
- *Weak SE* – the followers select such NE that minimizes the outcome of the leader.

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule
- *Strong SE* – the followers select such NE that maximizes the outcome of the leader (when the tie-braking is not specified we mean SSE),
- *Weak SE* – the followers select such NE that minimizes the outcome of the leader.

Exact Weak Stackelberg equilibrium does not have to exist.

There may be multiple Nash equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule
- *Strong SE* – the followers select such NE that maximizes the outcome of the leader (when the tie-braking is not specified we mean SSE),
- *Weak SE* – the followers select such NE that minimizes the outcome of the leader.

Exact Weak Stackelberg equilibrium does not have to exist.

$1 \setminus 2$	a	b	c	d	e
U	$(2,4)$	$(6,4)$	$(9,0)$	$(1,2)$	$(7,4)$
D	$(8,4)$	$(0,4)$	$(3,6)$	$(1,5)$	$(0,0)$

There may be multiple Nash equilibria

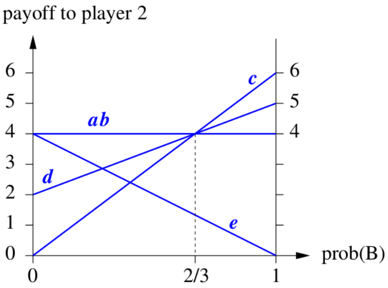

[1]Figure from [9].

There may be multiple Nash equilibria

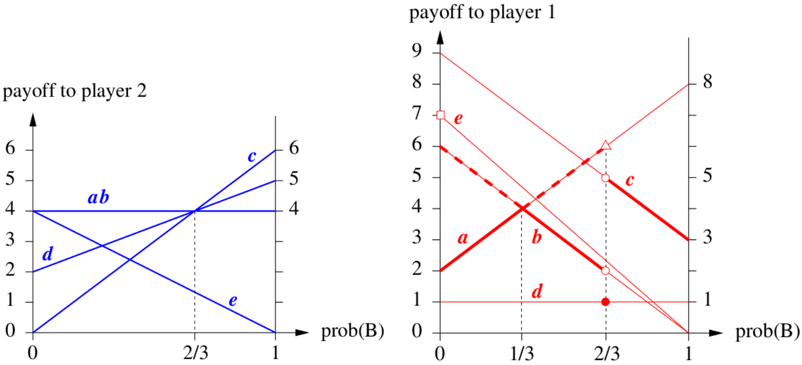

[1]

---

[1] Figure from [9].

Computing a Stackelberg equilibrium in NFGs

Computing a Stackelberg equilibrium in NFGs

The problem is polynomial for two-players normal-form games; 1 is the leader, 2 is the follower.

Computing a Stackelberg equilibrium in NFGs

The problem is polynomial for two-players normal-form games; 1 is the leader, 2 is the follower.

Baseline polynomial algorithm requires solving $|\mathcal{S}_2|$ linear programs:

Computing a Stackelberg equilibrium in NFGs

The problem is polynomial for two-players normal-form games; 1 is the leader, 2 is the follower.

Baseline polynomial algorithm requires solving $|\mathcal{S}_2|$ linear programs:

$$\max_{\sigma_1 \in \Sigma_1} \sum_{s_1 \in \mathcal{S}_1} \sigma_1(s_1) u_1(s_1, s_2)$$

$$\sum_{s_1 \in \mathcal{S}_1} \sigma_1(s_1) u_2(s_1, s_2) \geq \sum_{s_1 \in \mathcal{S}_1} \sigma_1(s_1) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1 \in \mathcal{S}_1} \sigma_1(s_1) = 1$$

one for each $s_2 \in \mathcal{S}_2$ assuming s_2 is the best response of the follower.

Computing a Stackelberg equilibrium in NFGs

Computing a Stackelberg equilibrium in NFGs

We can reformulate the program as a mixed-integer linear program (MILP) that is a basis for the hard cases (e.g., computing a SE in Bayesian games):

Computing a Stackelberg equilibrium in NFGs

We can reformulate the program as a mixed-integer linear program (MILP) that is a basis for the hard cases (e.g., computing a SE in Bayesian games):

$$\max_{\sigma \in \Sigma, y \in \{0,1\}^{|\mathcal{S}_2|}} \sum_{s \in \mathcal{S}} \sigma(s_1, s_2) u_1(s_1, s_2)$$

$$0 \leq \sigma(s_1, s_2) \leq y(s_2) \qquad \forall s_1, s_2 \in \mathcal{S}_{1,2}$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) = 1$$

$$\sum_{s_2 \in \mathcal{S}_2} y(s_2) = 1$$

Computing a Stackelberg equilibrium in EFGs

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

- two-player EFGs with chance (there exists a FPTAS for this case [2]),

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

- two-player EFGs with chance (there exists a FPTAS for this case [2]),
- two-player EFGs with imperfect information,

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

- two-player EFGs with chance (there exists a FPTAS for this case [2]),
- two-player EFGs with imperfect information,
- two-player EFGs with perfect information but imperfect recall (games on DAGs).

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

- two-player EFGs with chance (there exists a FPTAS for this case [2]),
- two-player EFGs with imperfect information,
- two-player EFGs with perfect information but imperfect recall (games on DAGs).

Main algorithms are based on the sequence-form LCP for computing NE:

Computing a Stackelberg equilibrium in EFGs

The problem is typically NP-hard [6, 2]:

- two-player EFGs with chance (there exists a FPTAS for this case [2]),
- two-player EFGs with imperfect information,
- two-player EFGs with perfect information but imperfect recall (games on DAGs).

Main algorithms are based on the sequence-form LCP for computing NE:

$$v_{\mathsf{inf}_i(\sigma_i)} = s_{\sigma_i} + \sum_{I'_i \in \mathcal{I}_i : \mathsf{seq}_i(I'_i) = \sigma_i} v_{I'_i} + \sum_{\sigma_{-i} \in \Sigma_{-i}} g_i(\sigma_i, \sigma_{-i}) \cdot r_{-i}(\sigma_{-i}) \qquad \forall i, \sigma_i$$

$$r_i(\sigma_i) = \sum_{a \in A(I_i)} r_i(\sigma_i a) \qquad \forall i \in \mathcal{N}\ \forall I_i \in \mathcal{I}_i,\ \sigma_i = \mathsf{seq}_i(I_i)$$

$$r_i(\emptyset) = 1 \quad 0 = r_i(\sigma_i) \cdot s_{\sigma_i} \qquad \forall i \in \mathcal{N}\ \forall \sigma_i \in \Sigma_i$$

$$0 \le r_i(\sigma_i)\ ; \qquad 0 \le s_{\sigma_i} \quad \forall i \in \mathcal{N}\ \forall \sigma_i \in \Sigma_i$$

Computing a Stackelberg equilibrium in EFGs

Computing a Stackelberg equilibrium in EFGs

MILP for computing SE for two-player extensive-form game with perfect recall:

Computing a Stackelberg equilibrium in EFGs

MILP for computing SE for two-player extensive-form game with perfect recall:

$$\max_{p,r,v,s} \sum_{z \in \mathcal{Z}} p(z) u_1(z) \mathcal{C}(z)$$

$$v_{\mathsf{inf}_2(\sigma_2)} = s_{\sigma_2} + \sum_{I' \in \mathcal{I}_2 : \mathsf{seq}_2(I') = \sigma_2} v_{I'} + \sum_{\sigma_1 \in \Sigma_1} r_1(\sigma_1) g_2(\sigma_1, \sigma_2) \qquad \forall \sigma_2 \in \Sigma_2$$

$$r_i(\emptyset) = 1 \qquad r_i(\sigma_i) = \sum_{a \in A_i(I_i)} r_i(\sigma_i a) \quad \forall i \in \mathcal{N}\ \forall I_i \in \mathcal{I}_i, \sigma_i = \mathsf{seq}_i(I_i)$$

$$0 \le s_{\sigma_2} \le (1 - r_2(\sigma_2)) \cdot M \qquad \forall \sigma_2 \in \Sigma_2$$

$$0 \le p(z) \le r_2(\mathsf{seq}_2(z)) \qquad \forall z \in \mathcal{Z}$$

$$0 \le p(z) \le r_1(\mathsf{seq}_1(z)) \qquad \forall z \in \mathcal{Z}$$

$$1 = \sum_{z \in \mathcal{Z}} p(z) \mathcal{C}(z)$$

$$r_2(\sigma_2) \in \{0, 1\} \qquad \forall \sigma_2 \in \Sigma_2$$

$$0 \le r_1(\sigma_1) \le 1 \qquad \forall \sigma_1 \in \Sigma_1$$

Stackelberg and Correlated Equilibrium

Stackelberg and Correlated Equilibrium

Recall the MILP program Stackelberg equilibrium and compare it to the LP for correlated equilibrium:

Stackelberg and Correlated Equilibrium

Recall the MILP program Stackelberg equilibrium and compare it to the LP for correlated equilibrium:

- we maximize the expected utility of the leader

Stackelberg and Correlated Equilibrium

Recall the MILP program Stackelberg equilibrium and compare it to the LP for correlated equilibrium:

- we maximize the expected utility of the leader
- we restrict the joint probability distribution so that the follower plays a pure strategy

Stackelberg and Correlated Equilibrium

Recall the MILP program Stackelberg equilibrium and compare it to the LP for correlated equilibrium:

- we maximize the expected utility of the leader
- we restrict the joint probability distribution so that the follower plays a pure strategy
- there are no incentive constraints of the leader

Stackelberg and Correlated Equilibrium

Recall the MILP program Stackelberg equilibrium and compare it to the LP for correlated equilibrium:

- we maximize the expected utility of the leader
- we restrict the joint probability distribution so that the follower plays a pure strategy
- there are no incentive constraints of the leader

We can compute a Stackelberg equilibrium if we modify an algorithm for computing an optimal correlated equilibrium.

Computing a Stackelberg equilibrium in NFGs (2)

Computing a Stackelberg equilibrium in NFGs (2)

We can reformulate the MILP program as a single LP:

Computing a Stackelberg equilibrium in NFGs (2)

We can reformulate the MILP program as a single LP:

$$\max_{\sigma \in \Sigma} \sum_{s \in \mathcal{S}} \sigma(s_1, s_2) u_1(s_1, s_2)$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) = 1$$

Computing a Stackelberg equilibrium in NFGs (2)

We can reformulate the MILP program as a single LP:

$$\max_{\sigma \in \Sigma} \sum_{s \in \mathcal{S}} \sigma(s_1, s_2) u_1(s_1, s_2)$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) = 1$$

Properties:

Computing a Stackelberg equilibrium in NFGs (2)

We can reformulate the MILP program as a single LP:

$$\max_{\sigma \in \Sigma} \sum_{s \in \mathcal{S}} \sigma(s_1, s_2) u_1(s_1, s_2)$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) = 1$$

Properties:

- the objective is the same as in the MILP case (or multiple LPs) case,

Computing a Stackelberg equilibrium in NFGs (2)

We can reformulate the MILP program as a single LP:

$$\max_{\sigma \in \Sigma} \sum_{s \in \mathcal{S}} \sigma(s_1, s_2) u_1(s_1, s_2)$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2) \geq \sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) u_2(s_1, s_2') \quad \forall s_2' \in \mathcal{S}_2$$

$$\sum_{s_1, s_2 \in \mathcal{S}_{1,2}} \sigma(s_1, s_2) = 1$$

Properties:

- the objective is the same as in the MILP case (or multiple LPs) case,
- strategy σ does not necessarily corresponds to Stackelberg equilibrium (the follower can receive multiple recommendations that are best responses).

Computing a Stackelberg equilibrium in EFGs (2)

Computing a Stackelberg equilibrium in EFGs (2)

How does it work in EFGs?

Computing a Stackelberg equilibrium in EFGs (2)

How does it work in EFGs?

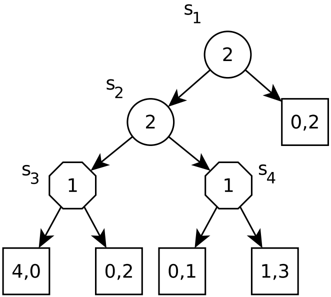

Computing a Stackelberg equilibrium in EFGs (2)

How does it work in EFGs?

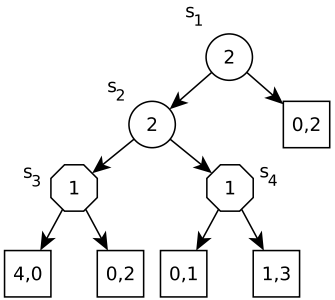

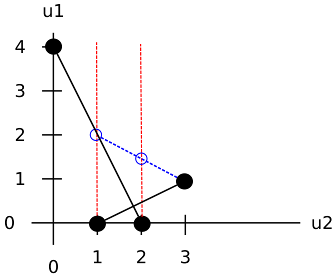

Computing a Stackelberg equilibrium in EFGs (2)

How does it work in EFGs?

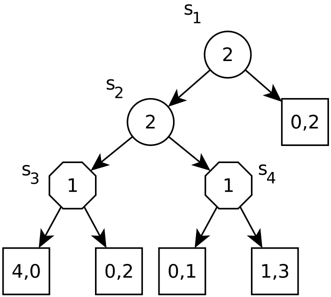

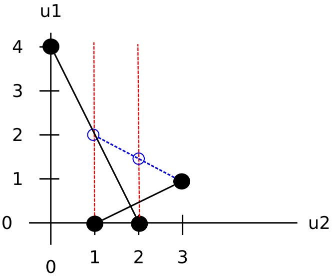

We can define a Stackelberg extension of EFCE [2] – the leader (1) controls the correlation device, (2) sends signals to the follower, (3) maximizes her expected utility.

Computing a Stackelberg equilibrium in EFGs (2)

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs
- remove the incentives constraints of the leader

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs
- remove the incentives constraints of the leader
- add objective to maximize the expected value of the leader

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs
- remove the incentives constraints of the leader
- add objective to maximize the expected value of the leader
- restrict the recommendations to the follower so that only a unique action in an information set

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs
- remove the incentives constraints of the leader
- add objective to maximize the expected value of the leader
- restrict the recommendations to the follower so that only a unique action in an information set

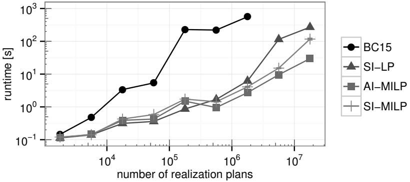

Computing a Stackelberg equilibrium in EFGs (2)

We can follow the same steps [3]:

- consider an algorithm for computing an optimal EFCE in an EFGs
- remove the incentives constraints of the leader
- add objective to maximize the expected value of the leader
- restrict the recommendations to the follower so that only a unique action in an information set

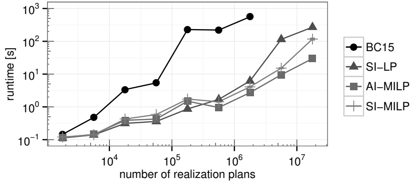

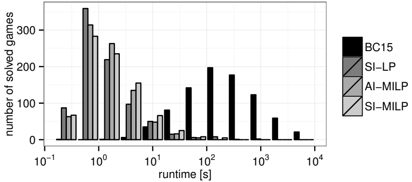

Computing a Stackelberg equilibrium in EFGs (3)

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

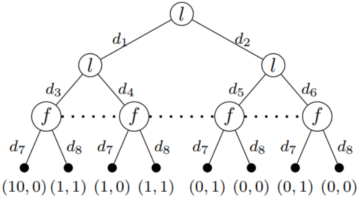

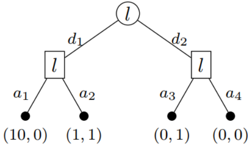

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

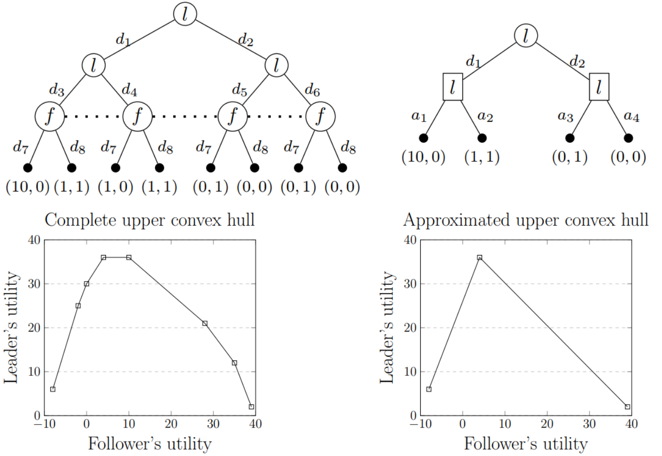

Computing a Stackelberg equilibrium in EFGs (3)

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

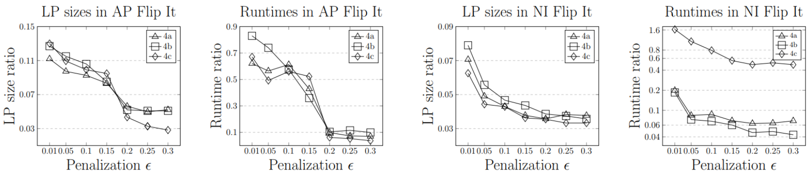

Computing a Stackelberg equilibrium in EFGs (3)

- incremental strategy generation [4]

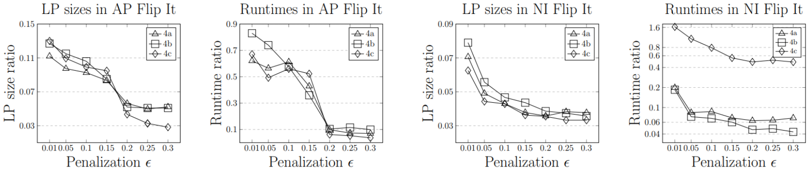

Instance \ ϵ	0.01	0.05	0.1	0.15	0.2	0.25	0.3
4a All-Points	0%	0%	0%	0.9%	2.42%	3.01%	3.23%
4a No-Info	0%	0.35%	0.72%	1.29%	1.77%	2.5%	2.55%
4b All-Points	0%	0%	0%	0.56%	0.8%	2.39%	2.48%
4b No-Info	0%	0.16%	0.67%	1.27%	2.15%	2.42%	2.86%
4c All-Points	0%	0%	0.033%	0.79%	3.47%	4.8%	6.38%
4c No-Info	0%	0.24%	0.89%	1.75%	1.75%	1.75%	1.75%

Computing a Stackelberg equilibrium in EFGs (4)

Computing a Stackelberg equilibrium in EFGs (4)

- Using Finite State Machines for Computing SE (under review for EC '19)

Computing a Stackelberg equilibrium in EFGs (4)

- Using Finite State Machines for Computing SE (under review for EC '19)
  - we can restrict the set of pure strategies that we consider for the follower

Computing a Stackelberg equilibrium in EFGs (4)

- Using Finite State Machines for Computing SE (under review for EC '19)
  - we can restrict the set of pure strategies that we consider for the follower
  - in an EFG, these restrictions can be described using (for example) Finite State Machines

Computing a Stackelberg equilibrium in EFGs (4)

- Using Finite State Machines for Computing SE (under review for EC '19)
  - we can restrict the set of pure strategies that we consider for the follower
  - in an EFG, these restrictions can be described using (for example) Finite State Machines

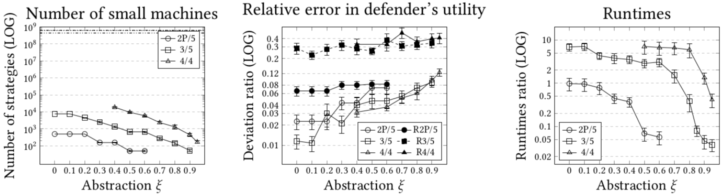

References I

(besides the books)

[1] B. Bošanský and J. Čermák, "Sequence-Form Algorithm for Computing Stackelberg Equilibria in Extensive-Form Games," in *AAAI Conference on Artificial Intelligence*, 2015.

[2] B. Bošanský, S. Branzei, K. A. Hansen, P. B. Miltersen, and T. B. Sørensen, "Computation of stackelberg equilibria of finite sequential games," in *Proceedings of Web and Internet Economics: 11th International Conference (WINE)*, pp. 201–215, 2015.

[3] J. Čermák, B. Bošanský, K. Durkota, V. Lisý, and C. Kiekintveld, "Using correlated strategies for computing stackelberg equilibria in extensive-form games," in *Proceedings of AAAI Conference on Artificial Intelligence (to appear)*, 2016.

[4] J. Černý, B. Bošanský and C. Kiekintveld, "Incremental Strategy Generation for Stackelberg Equilibria in Extensive-Form Games," in *ACM Conference on Economic Computation (EC)*, 2018.

References II

[5] V. Conitzer and D. Korzhyk, "Commitment to Correlated Strategies," in *Proceedings of AAAI Conference on Artificial Intelligence*, 2011.

[6] J. Letchford and V. Conitzer, "Computing optimal strategies to commit to in extensive-form games," in *Proceedings of the 11th ACM conference on Electronic commerce*, (New York, NY, USA), pp. 83–92, ACM, 2010.

[7] J. Letchford, L. MacDermed, V. Conitzer, R. Parr, and C. L. Isbell, "Computing optimal strategies to commit to in stochastic games," in *AAAI*, 2012.

[8] von Stengel, Bernhard and Zamir, Shmuel, "Leadership with commitment to mixed strategies," 2004.

[9] von Stengel, Bernhard and Zamir, Shmuel, "Leadership games with convex strategy sets," 2010.