Reminder

- Course project progress report 2: come to OH for discussions!
- HW5 due 4/4
- PRA6 due 4/16

Artificial Intelligence Methods for Social Good

Lecture 21:

Case Study: AI for Infrastructure Security

17-537 (9-unit) and 17-737 (12-unit)

Fei Fang
feifang@cmu.edu

Learning Objectives

- Describe the concept of
  - Dominant strategy
  - Stackelberg equilibrium
- Describe the Stackelberg Security Game (SSG) model
- Write down LP and MILP for solving a SSG
- For the airport protection problems, briefly describe
  - Significance/Motivation
  - Task being tackled, i.e., what is being solved/optimized
  - Model and method used to solve the problem
  - Evaluation process and criteria

Dominant Strategy

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

- Dominant Strategy
  - One strategy is always better/never worse/never worse and sometimes better than any other strategy
  - Focus on single player's strategy
  - Not always exist

s_i strictly dominates s_i' if

s_i very weakly dominates s_i' if

s_i weakly dominates s_i' if

s_i is a (strictly/very weakly/weakly) dominant strategy if it dominates $s_i', \forall s_i' \in S_i$

Dominant Strategy

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

- Dominant Strategy
  - One strategy is always better/never worse/never worse and sometimes better than any other strategy
  - Focus on single player's strategy
  - Not always exist

s_i strictly dominates s_i' if $\forall s_{-i}, u_i\,(s_i, s_{-i}) > u_i(s_i', s_{-i})$

s_i very weakly dominates s_i' if $\forall s_{-i}, u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$

s_i weakly dominates s_i' if $\forall s_{-i}, u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$
and $\exists s_{-i}, u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

s_i is a (strictly/very weakly/weakly) dominant strategy if it dominates $s_i', \forall s_i' \in S_i$

Dominant Strategy Equilibrium or Dominant Strategy Solution

- Dominant strategy equilibrium/solution
  - Every player plays a dominant strategy
  - Focus on strategy profile for all players
  - Not always exist
  - Can be found through enumeration

Q: Is there a dominant strategy equilibrium in the following game?

	Cooperate	Defect
Cooperate	-1,-1	-3,0
Defect	0,-3	-2,-2

	c	d
a	2,1	4,0
b	1,0	3,2

Power of Commitment

- NE utility=$(2,1)$
- If leader (player 1) commits to playing b, then player has to play d, leading to a utility of 3 for leader
- If leader (player 1) commits to playing a and b uniformly randomly, then player still has to play d, leading to a utility of 3.5 for leader

Player 2

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

Best Response Function

- Recall: Best response: Set of actions or strategies leading to highest expected utility given the strategies or actions of other players
  - $a_i^* \in BR(a_{-i})$ iff $\forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i, u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$

- Best Response Function
  - A mapping from a strategy of one player to a strategy of another player in the best response set
  - $f: S_1 \rightarrow S_2$ is a best response function iff $u_2\left(s_1, f(s_1)\right) \geq u_2(s_1, s_2), \forall s_1 \in S_1, s_2 \in S_2$. Or equivalently, $u_2\left(s_1, f(s_1)\right) \geq u_2(s_1, a_2), \forall s_1 \in S_1, a_2 \in A_2$

Stackelberg Equilibrium

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

- Stackelberg Equilibrium
  - Focus on strategy profile for all players
  - Follower responds according a best response function
  - $(s_1, f(s_1))$ is a Stackelberg Equilibrium iff
    - 1) f is a best response function
    - 2) $u_1(s_1, f(s_1)) \geq u_1(s_1', f(s_1')), \forall s_1' \in S_1$
  - There may exist many Stackelberg Equilibria due to different best response functions. For some best response functions, the Stackelberg Equilibrium may not exist

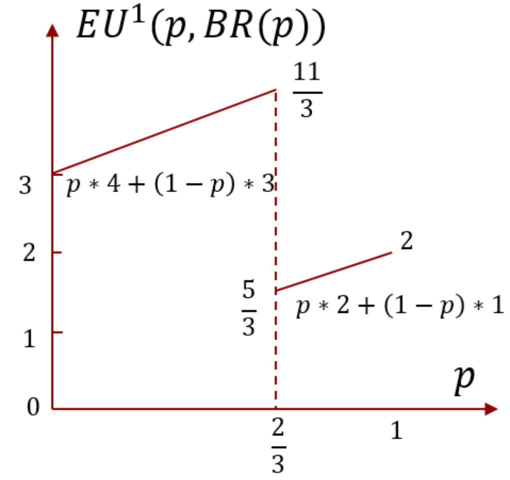

If $f\left(p = \frac{2}{3}\right) = d$, then SE is $s_1 = \left(\frac{2}{3}, \frac{1}{3}\right), s_2 = (0,1)$

If $f\left(p = \frac{2}{3}\right) = c$, then SE does not exist

Poll 1

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

- If the best response function break tie uniform randomly, does Stackelberg Equilibrium exist in this game?
  - A: Yes
  - B: No
  - C: I don't know

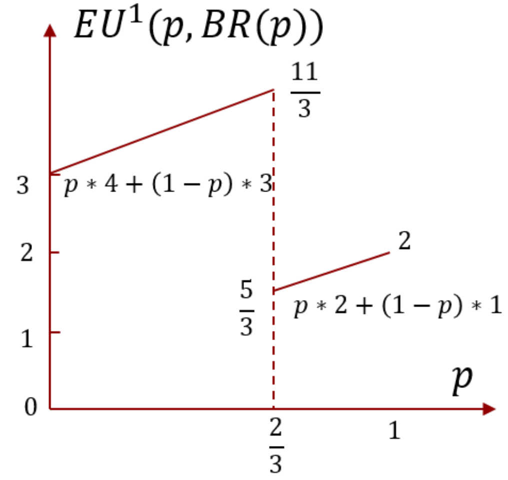

Strong Stackelberg Equilibrium

- Strong Stackelberg Equilibrium (SSE)
  - Follower breaks tie in favor of the leader
  - $(s_1, f(s_1))$ is a Strong Stackelberg Equilibrium iff
    - 1) f is a best response function
    - 2) $f(s) \in \underset{s_2 \in BR(s)}{\operatorname{argmax}} u_1(s, s_2)$
    - 3) $u_1\big(s_1, f(s_1)\big) \geq u_1\big(s_1', f(s_1')\big), \forall s_1' \in S_1$
  - SSE always exist in two-player finite games

Strong Stackelberg Equilibrium

- Remarks about Strong Stackelberg Equilibrium (SSE)
  - There may exist many SSEs but the leader's utility is the same in all these equilibria
  - Leader can induce the follower to breaks tie in favor of the leader by perturbing the strategy in the right direction
  - SSE coincide with minmax/maxmin/NE in two-player zero-sum finite games

Security Challenges


PARIS ATTACKS

Explosions in Brussels

Ansbach attack

A suicide bomb injured at least 12 in Germany's Ansbach, near Nuremberg, on July 24. This is the fourth violent incident in Germany in a week.

Wurzbuerg
July 18
A Pakistan refugee injured five people with an axe near Wurzbuerg.

Ansbach
July 24
A suicide bomb injured at least 12.

Reutlingen
July 24
A Syrian refugee killed a pregnant women.

Munich
July 22
A gunman shot nine people dead.

Source: Reuters

J. Wu, 25/07/2016

REUTERS

Security Challenges

Physical Infrastructure

Transportation Networks

Cyber Systems

Environmental Resources

Endangered Wildlife

Fisheries

Security Challenges

- Improve tactics of patrol, inspection, screening etc

Game Theoretic Reasoning

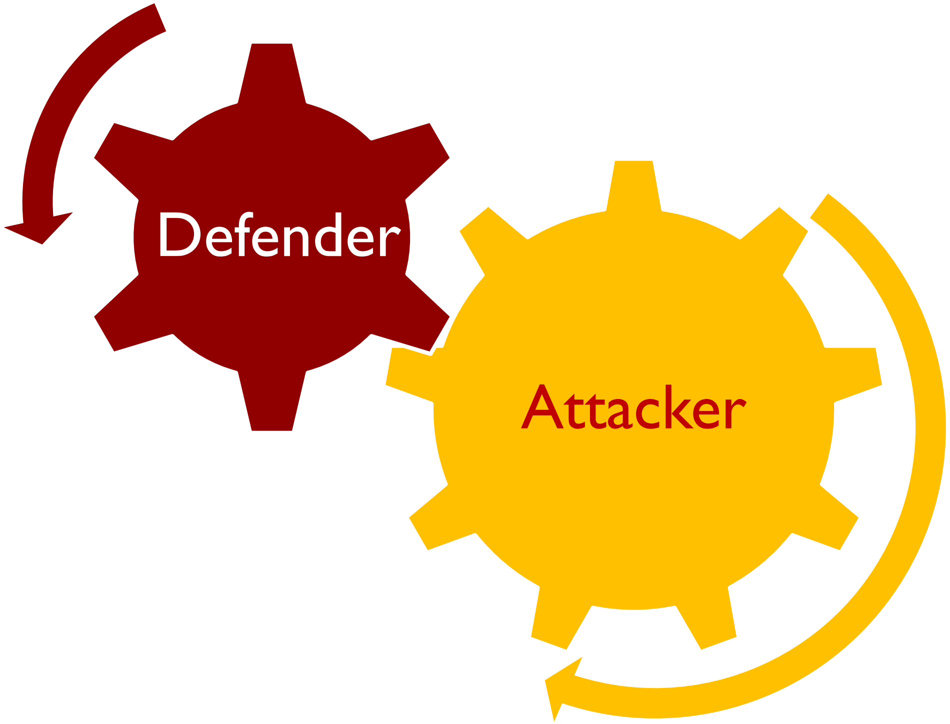

Protect Airports

- Limited resource allocation
- Adversary surveillance

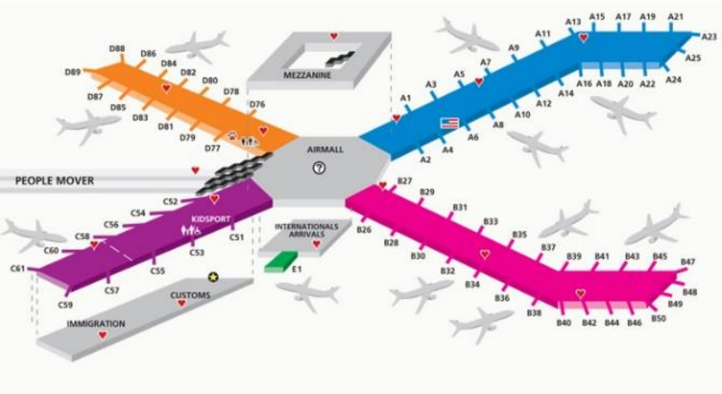

Adversary

Defender

	Target #1	Target #2
Target #1	5, -3	-1, 1
Target #2	-5, 4	2, -1

Protect Airports

- Limited resource allocation
- Adversary surveillance

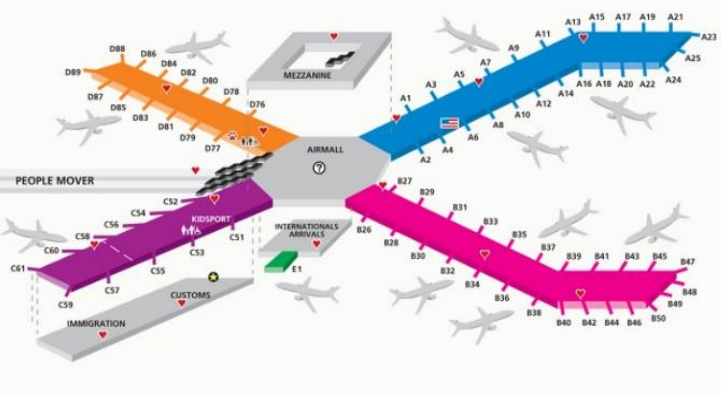

Adversary

Defender

	Target #1	Target #2
Target #1	5, -3	-1, 1
Target #2	-5, 4	2, -1

Protect Airports

- Randomization make defender unpredictable
- Stackelberg game
  - Leader: Defender; Commits to mixed strategy
  - Follower: Adversary; Conduct surveillance and best responds

Adversary

Defender

		Target #1	Target #2
55.6%	Target #1	5, -3	-1, 1
44.4%	Target #2	-5, 4	2, -1

Stackelberg Security Game (SSG)

- Leader: defender; Follower: attacker
- Defender allocate K resources to protect N targets
- Each target is associated with 4 values: $R_i^d, P_i^d, R_i^a, P_i^a$
  - If attacker attacks target i and succeeds: attacker gets R_i^a and defender gets P_i^d
  - If attacker attacks target i and fails: attacker gets $P_i^a (\leq R_i^a)$ and defender gets $R_i^d (\geq P_i^d)$

Adversary

Defender

	T1	T2	T3
T1	5, -3	-1, 1	
T2	-5, 4	2, -1	
T3			

	T1	T2	T3
R_i^d			3
P_i^d			-2
R_i^a			6
P_i^a			-2

Stackelberg Security Game (SSG)

- Leader: defender; Follower: attacker
- Defender allocate K resources to protect N targets
- Each target is associated with 4 values: $R_i^d, P_i^d, R_i^a, P_i^a$
  - If attacker attacks target i and succeeds: attacker gets R_i^a and defender gets P_i^d
  - If attacker attacks target i and fails: attacker gets $P_i^a (\leq R_i^a)$ and defender gets $R_i^d (\geq P_i^d)$

Adversary

Defender

	T1	T2	T3
T1	5, -3	-1, 1	-2, 6
T2	-5, 4	2, -1	-2, 6
T3	-5, 4	-1, 1	3, -2

↔

	T1	T2	T3
R_i^d	5	2	3
P_i^d	-5	-1	-2
R_i^a	4	1	6
P_i^a	-3	-1	-2

Poll 2

- Given a Stackelberg Security game with N targets, if we use a bimatrix to represent the payoffs, how many numbers do we need? If we use the penalty/reward for defender/attacker to represent the payoffs, how many numbers do we need?
  - A: $N^2, 4N$
  - B: N^2, N^2
  - C:$4N, 4N$
  - D:$4N, N^2$
  - E: None of the above
  - F: I don't know

Poll 3

- Let c_i be the probability the defender will protect target i in a Stackelberg security game, which of the following are the defender's expected utility when attacker attacks target i?
  - A: $c_i P_i^a + (1 - c_i) R_i^a$
  - B: $c_i R_i^d + (1 - c_i) P_i^d$
  - C: $P_i^d + c_i(R_i^d - P_i^d)$
  - D: $R_i^a + c_i(P_i^a - R_i^a)$
  - E: None of the above
  - F: I don't know

Compute SSE in SSG

$$AttEU(i) = c_i P_i^a + (1 - c_i) R_i^a$$
$$DefEU(i) = c_i R_i^d + (1 - c_i) P_i^d$$

- Strong Stackelberg Equilibrium
  - Attacker break tie in favor of defender
  - AttEU(1)=0.556*(-3)+0.444*4=0.11
  - AttEU(2)=0.556*1+0.444*(-1)=0.11
  - DefEU(1)=0.556*5+0.444*(-5)=0.56
  - DefEU(2)=0.556*(-1)+0.444*2=0.332
  - Equilibrium: DefStrat=(0.556,0.444), AttStrat=(1,0)

Adversary

Defender

		Target #1	Target #2
55.6%	**Target #1**	5, -3	-1, 1
44.4%	**Target #2**	-5, 4	2, -1

Computing SSE

$$AttEU(i) = c_i P_i^a + (1 - c_i) R_i^a$$
$$DefEU(i) = c_i R_i^d + (1 - c_i) P_i^d$$

- General-sum
  - Multiple LP
    - One LP for each target: Assume attacks target i^*
    - Choose the solution of the LP with the highest optimal value

Computing SSE

$$AttEU(i) = c_i P_i^a + (1 - c_i) R_i^a$$
$$DefEU(i) = c_i R_i^d + (1 - c_i) P_i^d$$

- General-sum
  - Multiple LP
    - One LP for each target: Assume attacks target i^*

$$\max_{\mathbf{c}} DefEU(i^*)$$
$$\text{s.t. } AttEU(i^*) \geq AttEU(i), \forall i = 1 \dots N$$
$$\sum_i c_i \leq 1$$
$$c_i \in [0,1]$$

    - Choose the solution of the LP with the highest optimal value

Computing SSE

$$AttEU(i) = c_i P_i^a + (1 - c_i) R_i^a$$
$$DefEU(i) = c_i R_i^d + (1 - c_i) P_i^d$$

- General-sum
  - MILP
    - Let $q_i \in \{0,1\}$ to indicate whether attacker attacks target i
    - Let M be a large constant, say 10^5

$$\max_{\mathbf{c},\mathbf{q},v} \sum_i DefEU(i) q_i$$
$$\text{s.t. } 0 \le v - AttEU(i) \le (1 - q_i)M, \forall i$$
$$\sum_i c_i \le 1$$
$$\sum_i q_i = 1$$
$$c_i \in [0,1], q_i \in \{0,1\}$$

Computing SSE

$$AttEU(i) = c_i P_i^a + (1 - c_i) R_i^a$$
$$DefEU(i) = c_i R_i^d + (1 - c_i) P_i^d$$

- Zero-sum
  - Single LP
  - SSE=NE=Minimax=Maximin

$$\min_{\mathbf{c},v} v$$
$$\text{s.t. } v \geq AttEU(i), \forall i = 1 \dots N$$
$$\sum_i c_i \leq 1$$
$$c_i \in [0,1]$$

ARMOR: Optimizing Security Resource Allocation [2007]

First application: Computational game theory for operational security

January 2009

•January 3rd	*Loaded 9/mm pistol*
•January 9th	*16-handguns,* *1000 rounds of ammo*
•January 10th	*Two unloaded shotguns*
•January 12th	*Loaded 22/cal rifle*
•January 17th	*Loaded 9/mm pistol*
•January 22nd	*Unloaded 9/mm pistol*

ARMOR for AIRPORT SECURITY at LAX [2008]
Congressional Subcommittee Hearings

Commendations
City of Los Angeles

Erroll Southers testimony
Congressional subcommittee

Newsweek

ARMOR…throws a digital cloak of invisibility….

Compute optimal defender strategy

- Polynomial time solvable in games with finite actions and simple structures [Conitzer06]
- NP-Hard in general settings [Korzhyk10]
- SSE=NE for zero-sum games, SSE⊂NE for games with special properties [Yin10]

- Research Challenges
  - Massive scale games with constraints
  - Plan/reason under uncertainty
  - Repeated interaction

Challenge: Scheduling Constraints and Scalability

- Mumbai Police Checkpoints

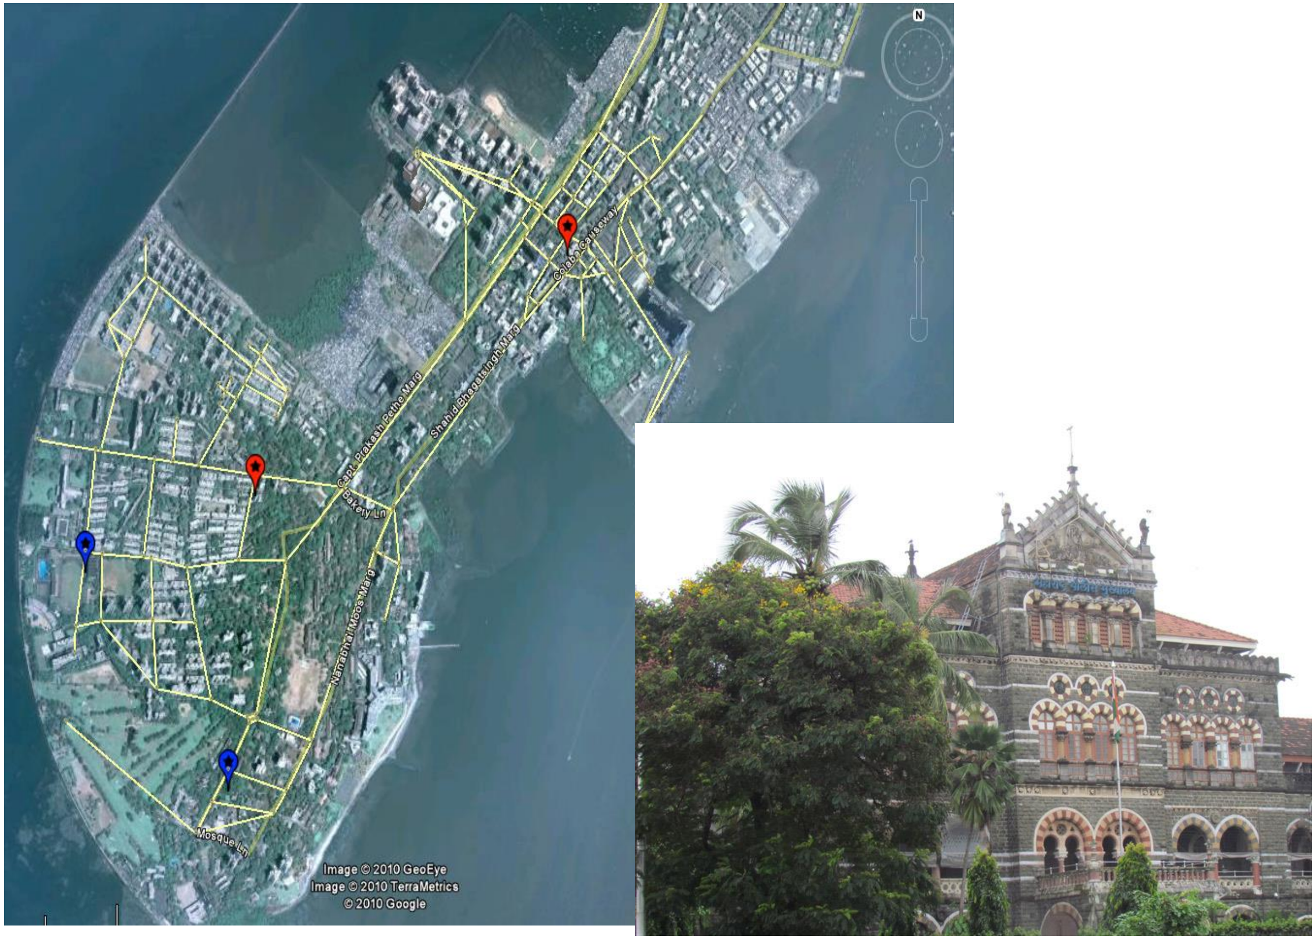

Challenge: Scheduling Constraints and Scalability

- Defender: Choose K checkpoints
- Attacker: Choose a target node (red) and a path from an entry node (green) to the target node
- Exponentially many pure strategies

Fully connected road network
20 intersections, 190 roads
5 resources, 1 target
~ 2 billion defender allocations
6.6 *quintillion* (10^{18}) attacker paths

Real Problem:
~500 intersections
~2000 roads

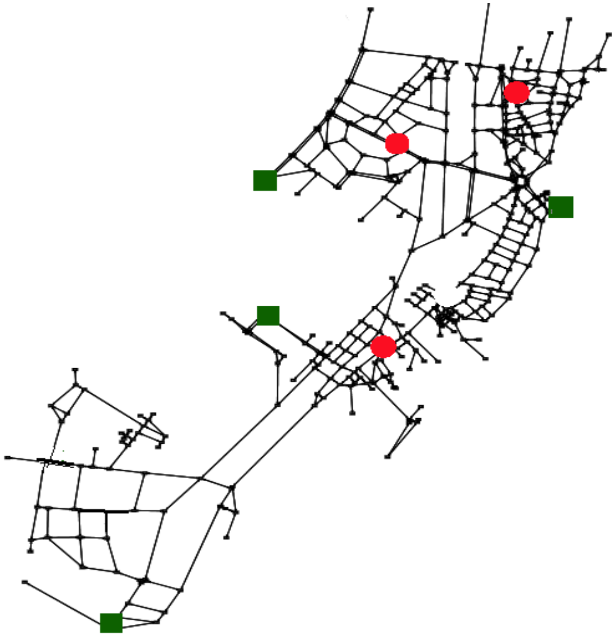

Double Oracle

- Intuition: No need to consider all possible pure strategies
- Start with a small set of pure strategies
- Iteratively add new pure strategies to be considered
- Provably converge to equilibrium

  in zero-sum games

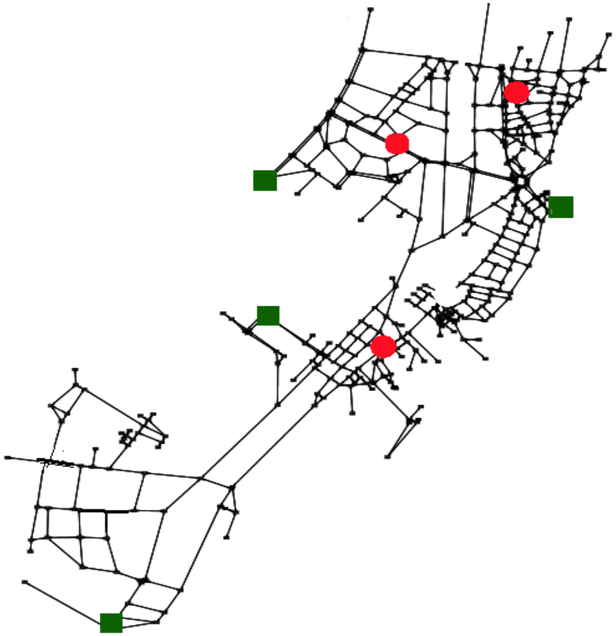

Payoff Matrix (When Zero-Sum)

Attacker Paths

Defender Allocations

$$\begin{array}{cc} & \begin{array}{ccccc} A_1 & A_2 & A_3 & A_4 & \ldots \end{array} \\ \begin{array}{c} X_1 : \\ X_2 : \\ \vdots \end{array} & \begin{bmatrix} -5 & -8 & 0 & -9 & \ldots \\ 0 & -8 & -15 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \end{array}$$

Double Oracle Algorithm

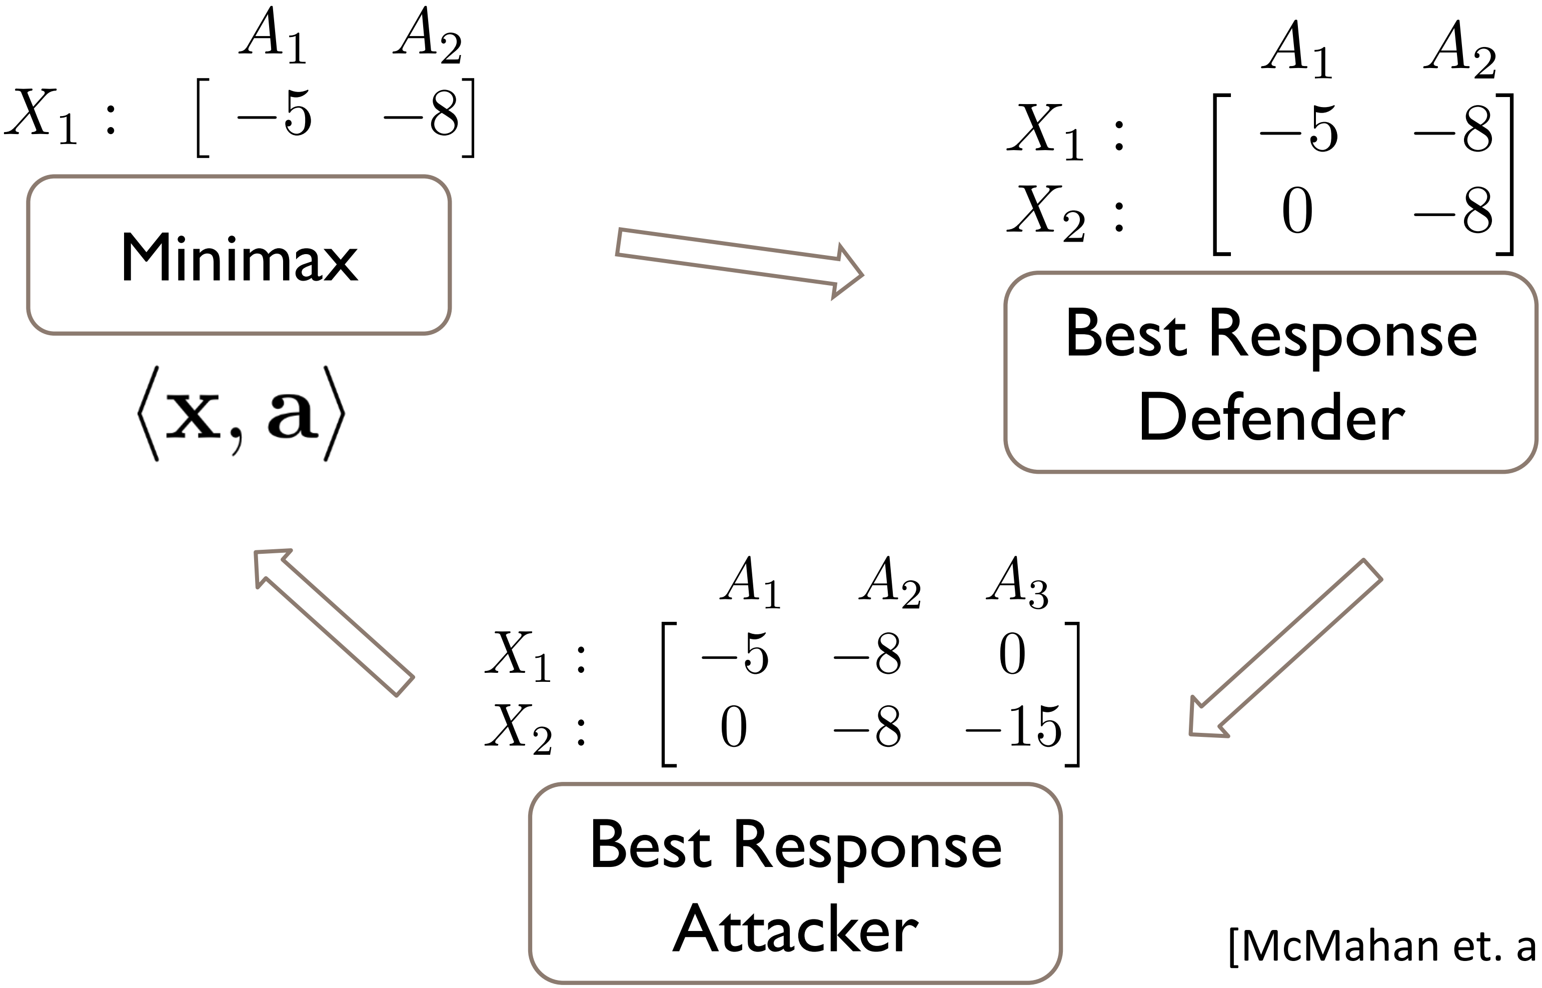

$$\begin{array}{c|cc} & A_1 & A_2 \\ X_1 : & -5 & -8 \end{array}$$

Minimax

$$\langle \mathbf{x}, \mathbf{a} \rangle$$

$$\begin{array}{c|cc} & A_1 & A_2 \\ X_1 : & -5 & -8 \\ X_2 : & 0 & -8 \end{array}$$

Best Response Defender

$$\begin{array}{c|ccc} & A_1 & A_2 & A_3 \\ X_1 : & -5 & -8 & 0 \\ X_2 : & 0 & -8 & -15 \end{array}$$

Best Response Attacker

[McMahan et. al 2003]

Variation

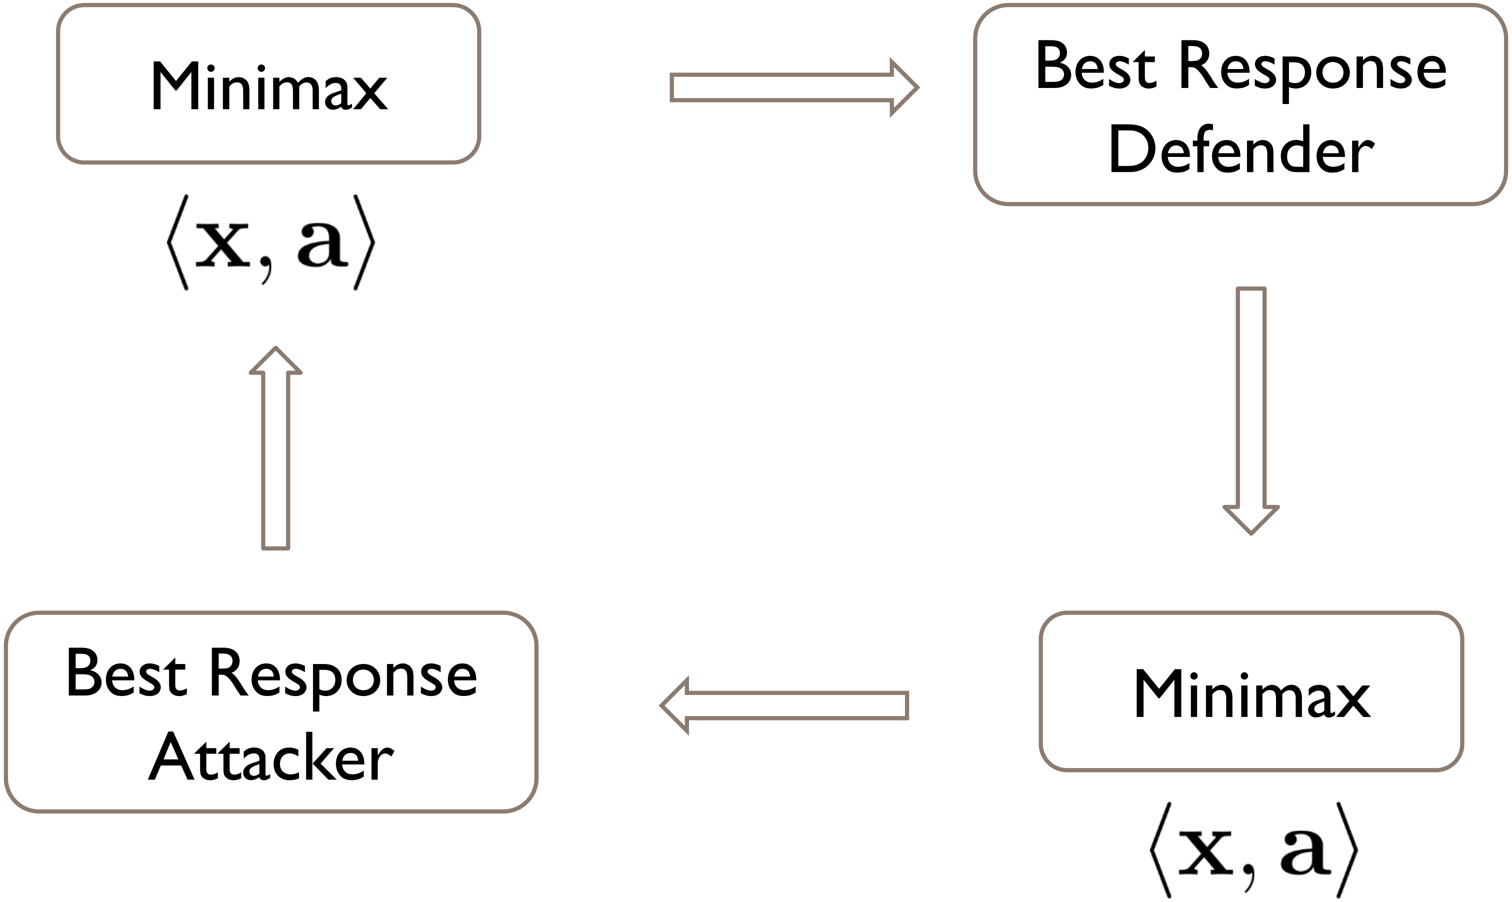

Minimax

$\langle \mathbf{x}, \mathbf{a} \rangle$

Best Response Defender

Minimax

$\langle \mathbf{x}, \mathbf{a} \rangle$

Best Response Attacker

Example

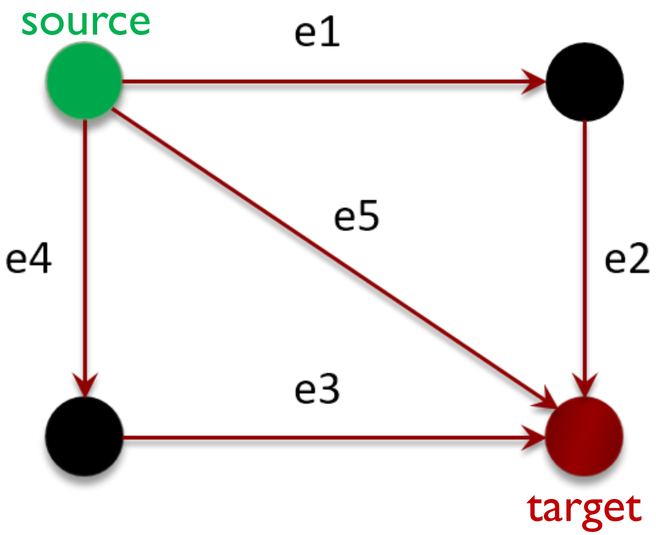

1 Defender Resource

	Defender Payoff	Attacker Payoff
Attack Successful	-T	T
Attack Failure	0	0

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e5->t	s->e4->e3->t
e1			
e2			
e3			
e4			
e5			

Example

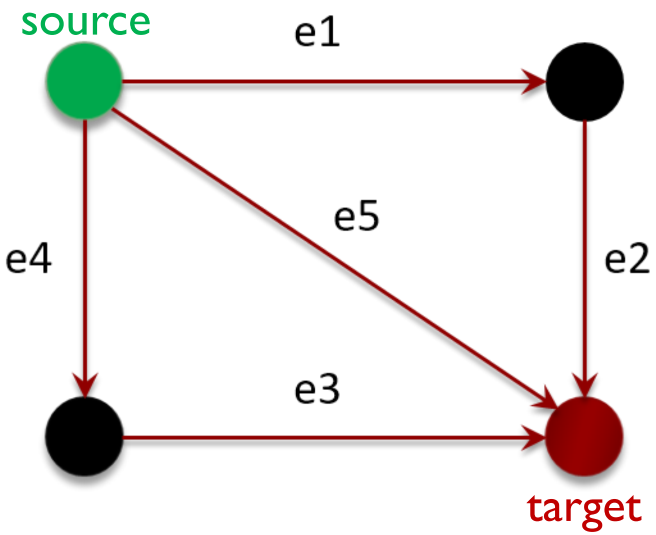

1 Defender Resource

	Defender Payoff	Attacker Payoff
Attack Successful	-T	T
Attack Failure	0	0

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e5->t	s->e4->e3->t
e1		-T, T	-T, T
e2		-T, T	-T, T
e3	-T, T	-T, T	
e4	-T, T	-T, T	
e5	-T, T		-T, T

Example

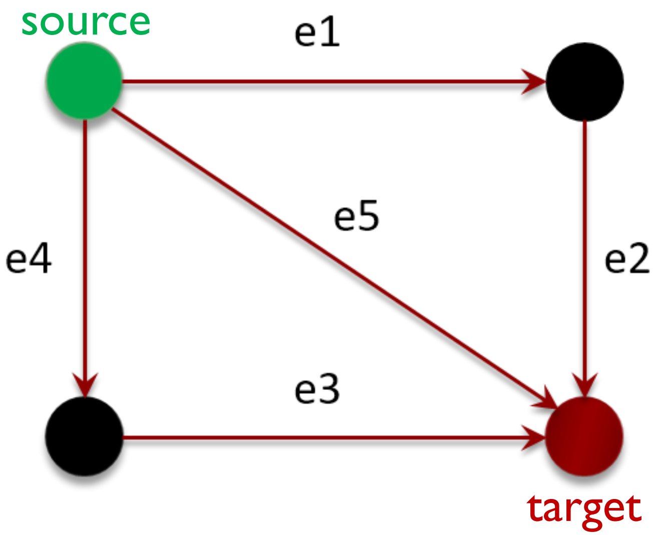

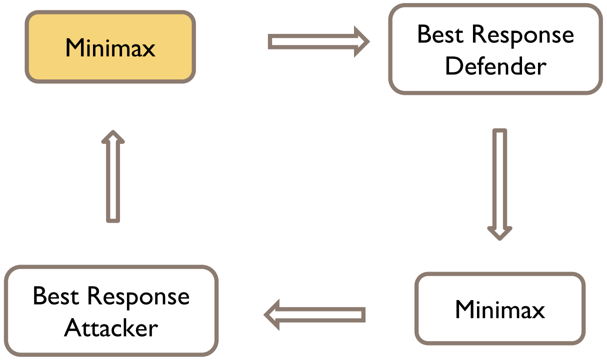
Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

	s->e1->e2->t
e1	0,0

Defender Allocations

Minimax strategy:

Defender Strategy: [1.0]

Attacker Strategy: [1.0]

Example

Minimax strategy:
Defender Strategy: [1.0]
Attacker Strategy: [1.0]

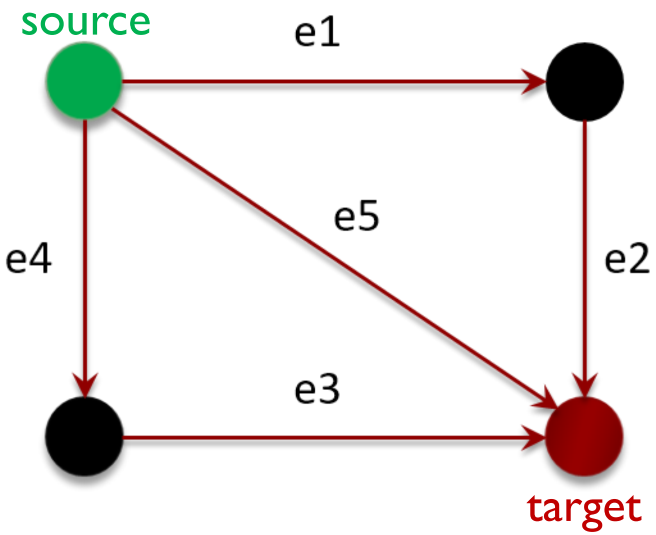

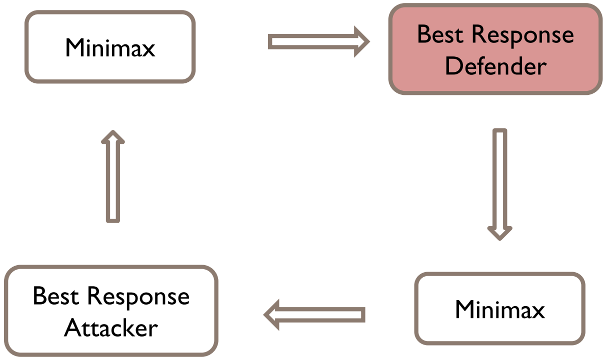

Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t
e1	0,0

Defender's best response: e1 or e2

Best response already in the table, no change

Example

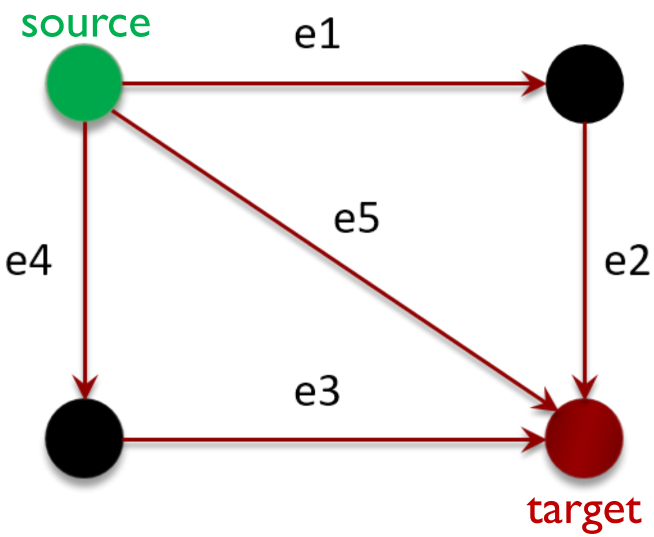

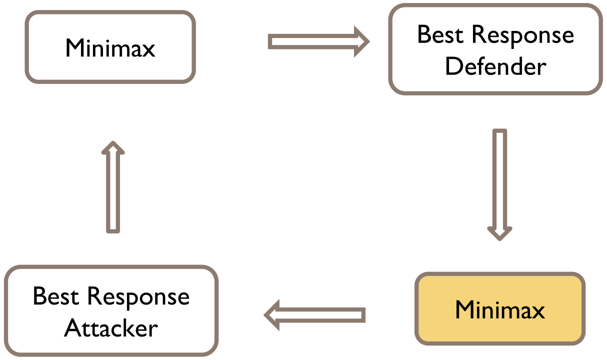

Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t
e1	0,0

Minimax strategy: no change

Example

Minimax strategy:
Defender Strategy: [1.0]
Attacker Strategy: [1.0]

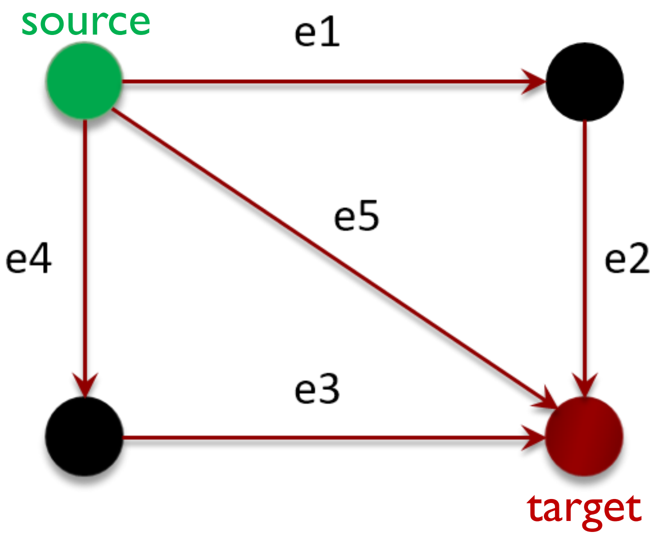

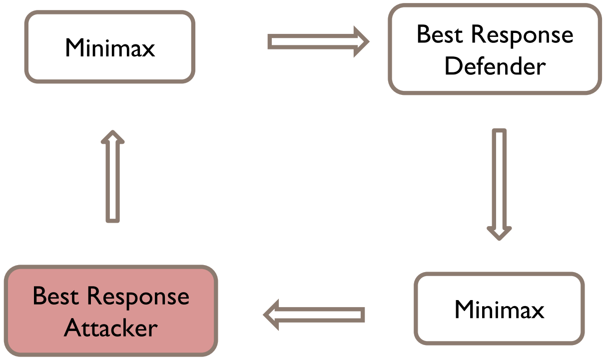

Attacker Paths

Defender Allocations

	s->e1->e2->t
e1	0,0

Attacker's best response: s->e4->e3->t or s->e5->t

Pick an arbitrary one, say s->e4->e3->t

Example

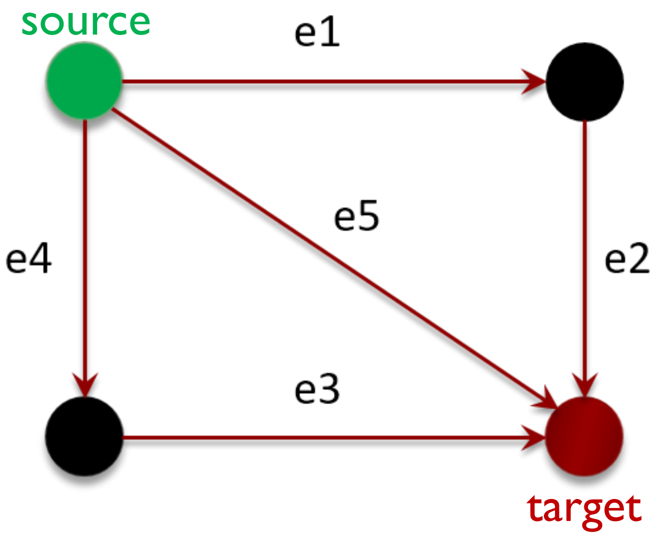

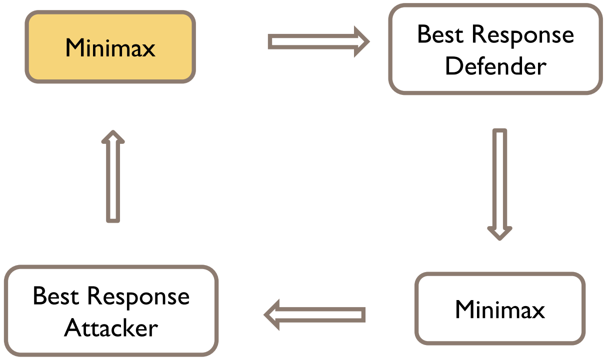

Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t
e1		-T,T

Minimax strategy:
Defender Strategy: [1.0]
Attacker Strategy: [0.0, 1.0]

Example

Minimax strategy:
Defender Strategy: [1.0]
Attacker Strategy: [0.0, 1.0]

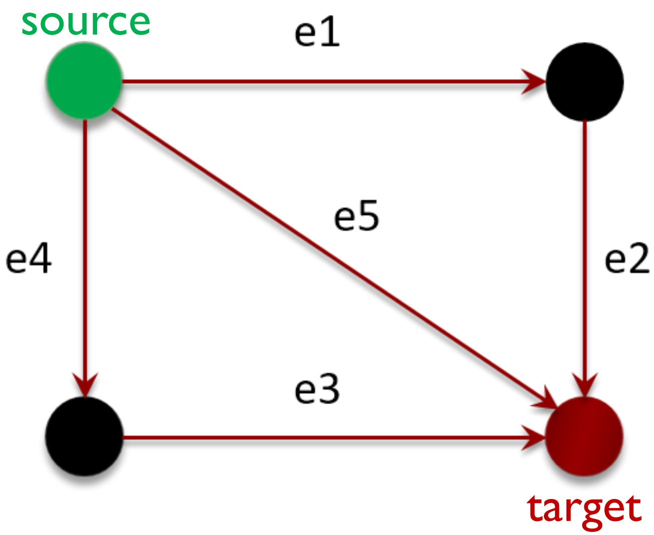

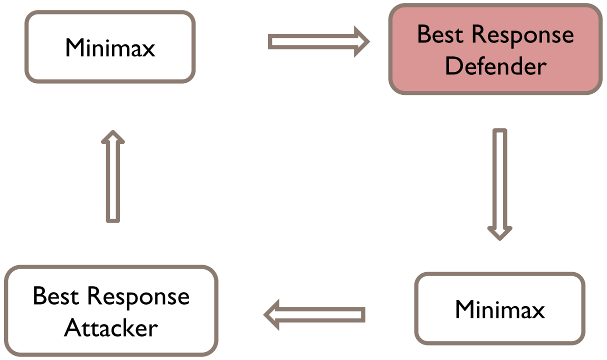

Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t
e1		-T,T

Defender's best response: e3 or e4

Pick e3

Example

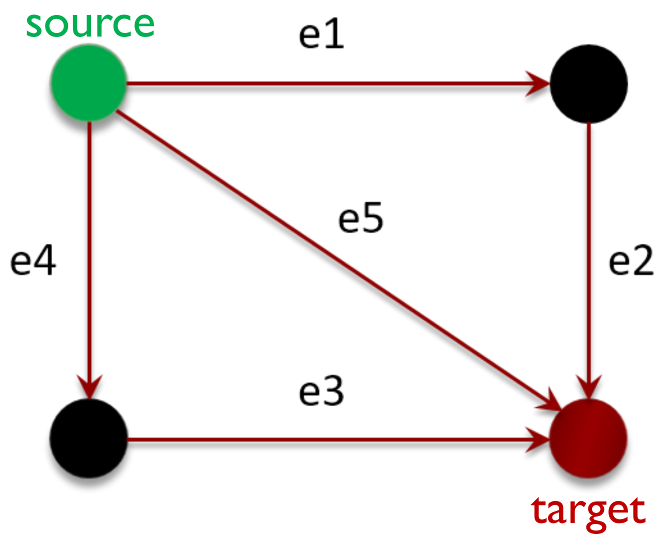

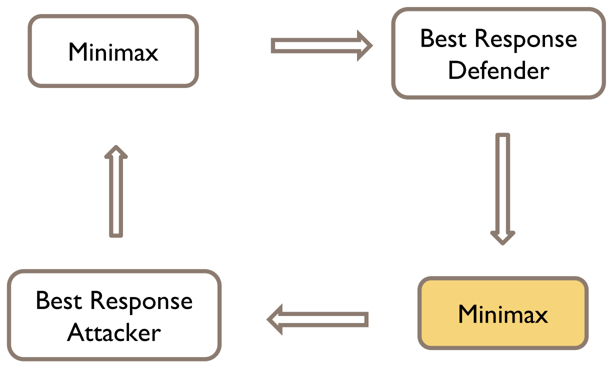
Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t
e1		-T,T
e3	-T,T	

Minimax strategy:

Defender Strategy: [0.5, 0.5]

Attacker Strategy: [0.5, 0.5]

Example

Minimax strategy:
Defender Strategy: [0.5, 0.5]
Attacker Strategy: [0.5, 0.5]

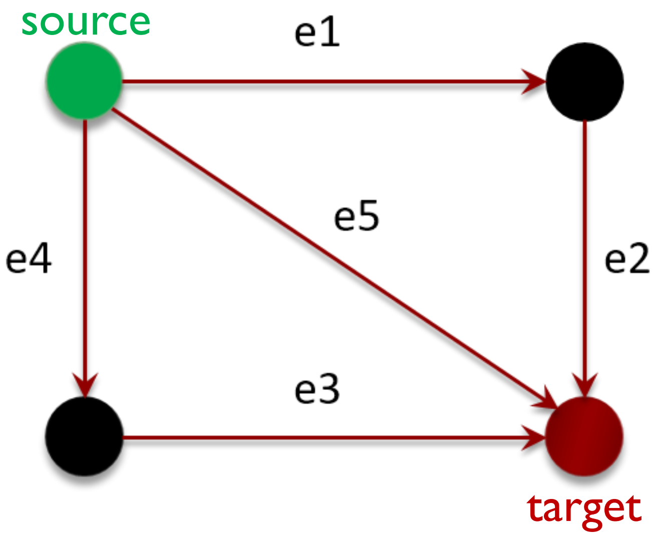

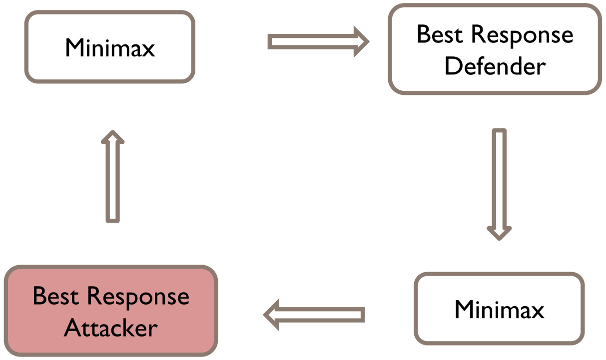
Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t
e1	0,0	-T,T
e3	-T,T	

Attacker's best response: s->e5->t

Example

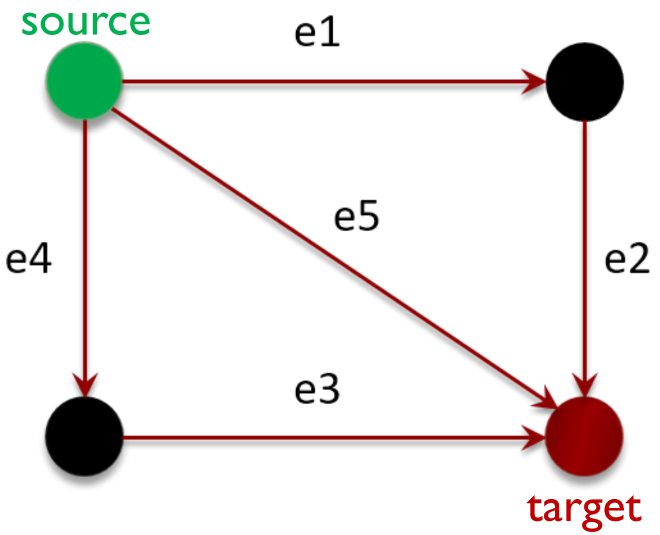

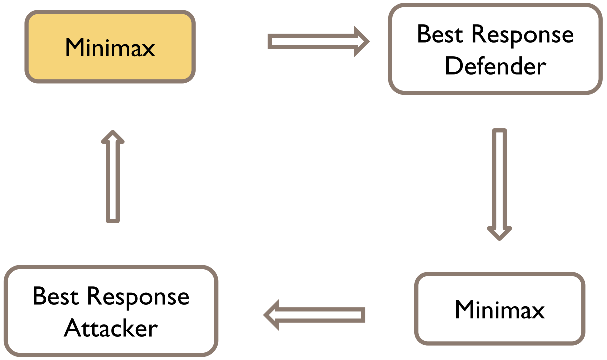
Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t	s->e5->t
e1		-T,T	-T,T
e3	-T,T		-T,T

Minimax strategy:
Defender Strategy: arbitrary, say [1.0, 0.0]
Attacker Strategy: [0.0, 0.0, 1.0]

Example

Minimax strategy:
Defender Strategy: [1.0, 0.0]
Attacker Strategy: [0.0, 0.0, 1.0]

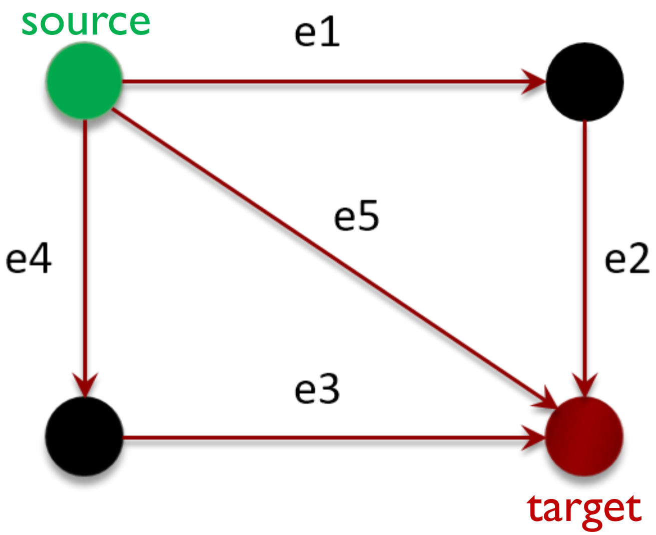

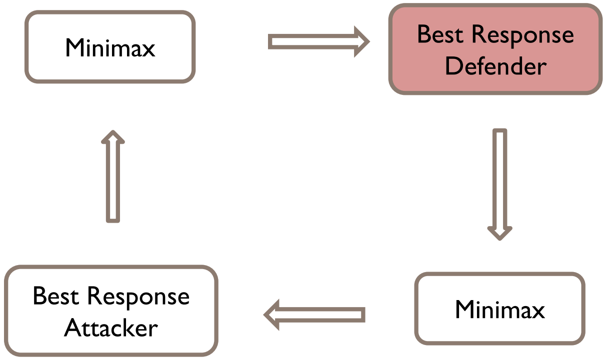

Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t	s->e4->e3->t	s->e5->t
e1		-T,T	-T,T
e3	-T,T		-T,T

Defender's best response: e5

Example

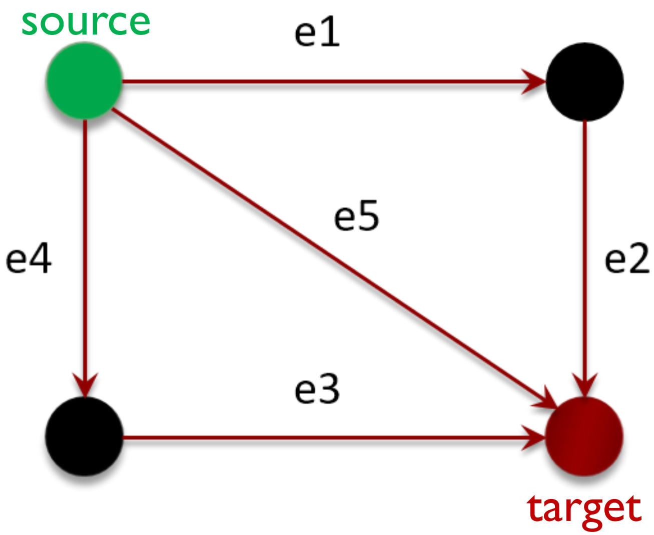

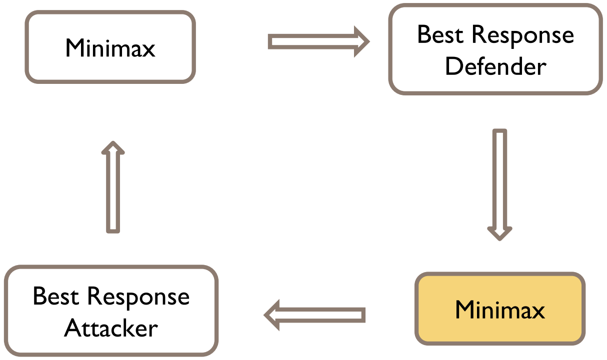
Minimax

Best Response Defender

Minimax

Best Response Attacker

Attacker Paths

		s->e1->e2->t	s->e4->e3->t	s->e5->t
Defender Allocations	e1		-T,T	-T,T
	e3	-T,T		-T,T
	e5	-T,T		-T,T

Defender Strategy: [1/3, 1/3, 1/3]
Attacker Strategy: [1/3, 1/3, 1/3]

No new best responses will be added in the next iteration. Terminate.

Poll 4

- Assume the following table is the game matrix (zero-sum). At some point in the process of the double oracle algorithm, a smaller game is being considered, with row 1, 2 and column 3,4. What action should be added in the next iteration?
- A_1
- A_2
- X_1
- X_2

Attacker Paths

Defender Allocations

$$\begin{array}{c|cccc} & A_1 & A_2 & A_3 & A_4 \\ X_1 : & -5 & -8 & 0 & -9 \\ X_2 : & 0 & -8 & -15 & 0 \end{array}$$

Poll 4

- Assume the following table is the game matrix (zero-sum). At some point in the process of the double oracle algorithm, a smaller game is being considered, with row 1, 2 and column 3,4. What action should be added in the next iteration?
- A_1
- A_2
- X_1
- X_2
- None

The minimax strategy of this smaller game is Def: (5/8, 3/8), Att: (3/8,5/8). Expected utility for attacker of taking each of the action is 5*5/8, 8, 15*3/8, 9*5/8

Attacker Paths

Defender Allocations

$$\begin{array}{c|cccc} & A_1 & A_2 & A_3 & A_4 \\ \hline X_1: & -5 & -8 & 0 & -9 \\ X_2: & 0 & -8 & -15 & 0 \end{array}$$

Warm Start

- Initialize with some subset of pure strategies (e.g., for defender, K edges in the min-cut)

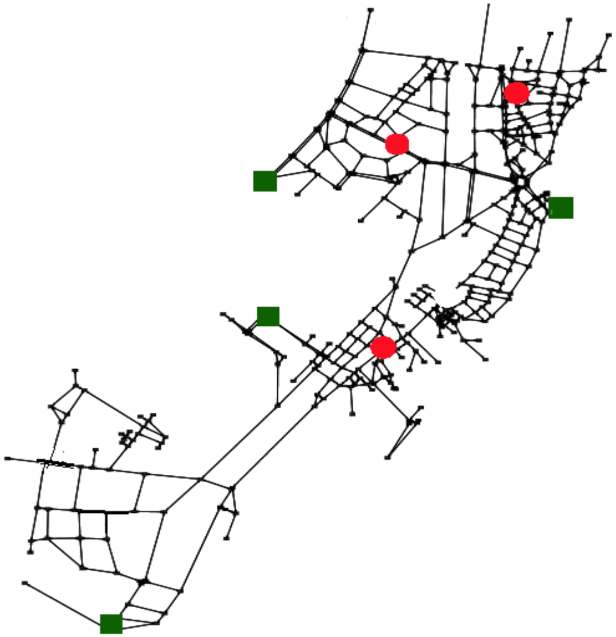

Better Responses

- No need to find the best response
- If you find a better response but not sure if it is the best response, it is OK to add it and move on
- If you cannot find a better response, it means the best response is already in the current support
- Impact on computation time varies

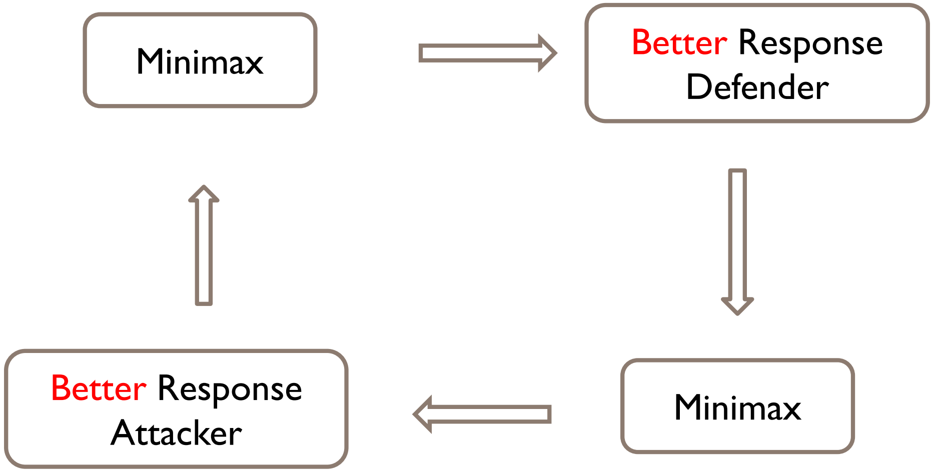

Minimax

Better Response Defender

Minimax

Better Response Attacker

Column Generation: Using One Oracle Only

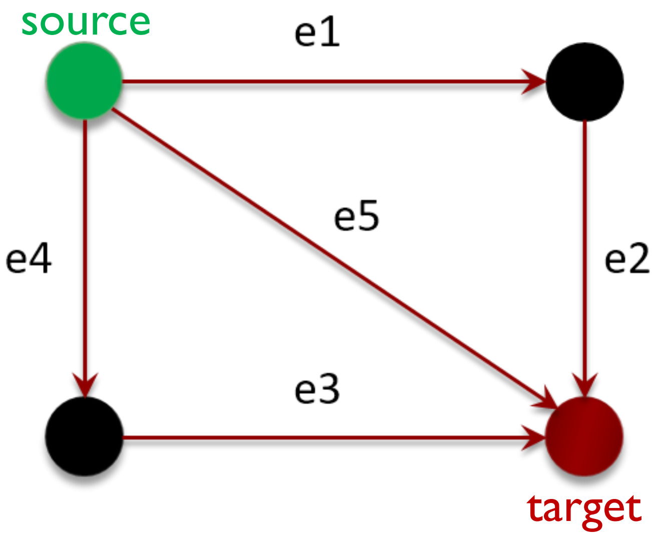

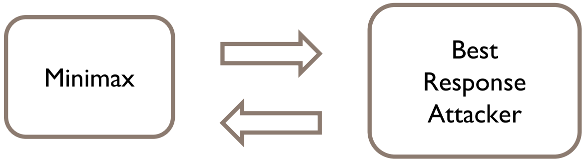

Minimax

Best Response Attacker

Attacker Paths

Defender Allocations

	s->e1->e2->t
e1	
e2	
e3	-T,T
e4	-T,T
e5	-T,T

Additional Resources and References

Additional Resources

- *Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport*
- A Double Oracle Algorithm for Zero-Sum Security Games on Graphs

References

- Conitzer, Vincent, and Tuomas Sandholm. "Computing the optimal strategy to commit to." In *Proceedings of the 7th ACM conference on Electronic commerce*, pp. 82-90. 2006.
- McMahan, H. Brendan, Geoffrey J. Gordon, and Avrim Blum. "Planning in the presence of cost functions controlled by an adversary." In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 536-543. 2003.

Backup Slides