Algorithmic Game Theory

Siddharth Barman

Indian Institute of Science

Game Theory: Study of how self-interested agents interact.

Game Theory: Study of how self-interested agents interact.

This field uses mathematical models to understand and predict outcomes of strategic interactions.

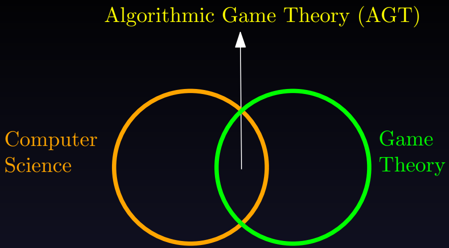

Roots: Advent of the internet

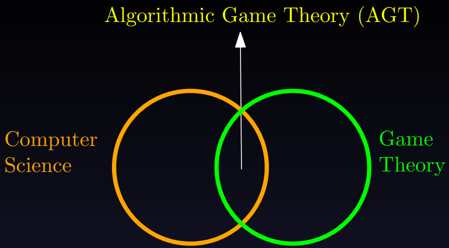

Game Theory:

- Emphasis on exact solutions/characterizations
- Largely Bayesian analysis
- Typically ignores computation/communication

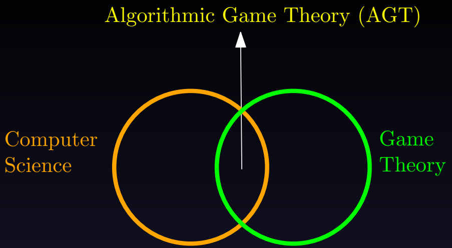

Game Theory:

- Emphasis on exact solutions/characterizations
- Largely Bayesian analysis
- Typically ignores computation/communication

Theoretical Computer Science:

- Computational complexity
- Approximation bounds
- Worst-case guarantees

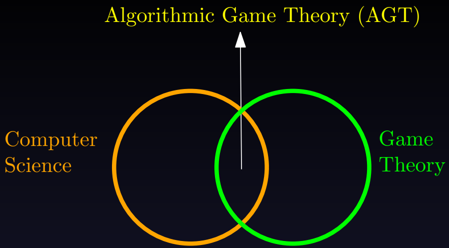

AGT Themes:

- Complexity of Equilibria
- Algorithmic Mechanism Design
  (Incentive Compatibility $\bigcap$ Efficient Algorithms)
- Computational Social Choice
- .....

Illustrative Applications:

Illustrative Applications:

- Traffic Networks: Self-interested users strategically choosing routes in a network to minimize the delay they face.

Insight: Formal explanation of Braess' paradox.
Sub-optimality of Equilibria–Price of Anarchy

Illustrative Applications:

- Traffic Networks
- Auctions: Strategic vendor auctioning goods to self-interested bidders

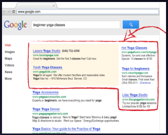

Insight: A simple auction with one extra bidder earns more revenue than the optimal auction with the original bidders (Bulow and Klemperer 1996).
Prior-Free Mechanism Design

Illustrative Applications:

- Traffic Networks
- Auctions
- Stable Matchings: Determine a stable assignment for self-interested entities that have rankings for each other

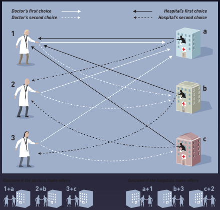

Insight: The stark effect of competition (Ashlagi et al. 2015).

Computational Complexity of Equilibria

Focus: Two-Player Games

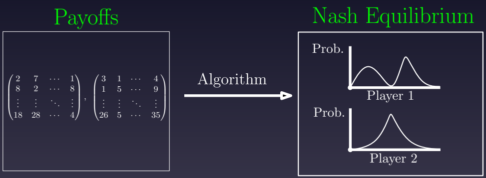

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

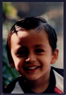

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

At (E, A) no player has an incentive to unilaterally deviate

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

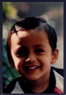

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

At (E, A) and at (NE, NA) no player has an incentive to unilaterally deviate

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

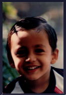

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

(E, A) and (NE, NA) are Pure Nash Equilibria of the game

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

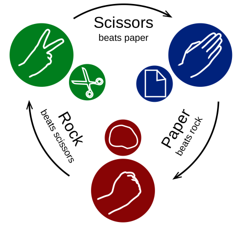

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Notation:

$$u_1(R, P) = -1$$

$$u_2(R, P) = 1$$

....

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Amongst rational players, deterministic strategies are <u>not</u> stable.

Therefore, we must consider strategies in which players randomize between actions.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Notation:

$$u_1(R, P) = -1$$

$$u_2(R, P) = 1$$

....

This is a *zero-sum* game

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0.$

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R, P, S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R, P, S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.
- Also, $u_2(\sigma, R) = u_2(\sigma, P) = u_2(\sigma, S) = u_2(\sigma, \sigma) = 0$.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\text{R, P, S}\}$.
- Expected utility of first player $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.
- Also, $u_2(\sigma, R) = u_2(\sigma, P) = u_2(\sigma, S) = u_2(\sigma, \sigma) = 0$.

At (σ, σ) players have no incentive to unilaterally deviate

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R, P, S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.
- Also, $u_2(\sigma, R) = u_2(\sigma, P) = u_2(\sigma, S) = u_2(\sigma, \sigma) = 0$.

(σ, σ) is a Nash equilibrium of the game

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Two Players
- Actions of the players: $\{1, 2, \ldots, n\}$

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Two Players
- Actions of the players: $\{1, 2, \ldots, n\}$
- Action profiles: $(i, j) \in [n] \times [n]$
- Player p's utility: $u_p : [n] \times [n] \to \mathbb{R}$

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Two Players
- Actions of the players: $\{1, 2, \ldots, n\}$
- Action profiles: $(i, j) \in [n] \times [n]$
- Player p's utility: $u_p : [n] \times [n] \to \mathbb{R}$

In Rock-Paper-Scissors, $n = 3$.

$u_1(R, P) = -1$, $u_2(R, P) = 1$,...

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Two Players
- Actions of the players: $\{1, 2, \ldots, n\}$
- Action profiles: $(i, j) \in [n] \times [n]$
- Player p's utility: $u_p : [n] \times [n] \to \mathbb{R}$

Prob. distributions (σ_1, σ_2) denote a Nash equilibrium iff

$$u_1(\sigma_1, \sigma_2) \geq u_1(a, \sigma_2) \qquad \text{for all } a \in [n]$$
$$u_2(\sigma_1, \sigma_2) \geq u_2(\sigma_1, b) \qquad \text{for all } b \in [n]$$

Fundamental Results

Guaranteed Existence of Nash Equilibria

- In two-player zero-sum games [von Neumann 1928]
- In finite games [Nash 1950]

John von Neumann

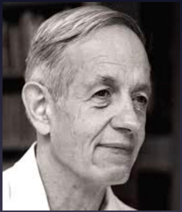

John Nash

Two-Player Zero-Sum Games

Recall Rock-Paper-Scissors:

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Two-Player Zero-Sum Games

In general, for any action a_1 and a_2

$$u_1(a_1, a_2) + u_2(a_1, a_2) = 0$$

Two-Player Zero-Sum Games

In general, for any action a_1 and a_2

$$u_2(a_1, a_2) = -u_1(a_1, a_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\max_{\sigma_1 \in \Delta(n)} \min_{\sigma_2 \in \Delta(n)} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value $=$ smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value $=$ smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

Maximin value $=$ Minimax value

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\max_{\sigma_1} \min_{\sigma_2} \; u_1(\sigma_1, \sigma_2) = \min_{\sigma_1} \max_{\sigma_1} \; u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Minimax Theorem $\Rightarrow$ Existence of Nash Eq. in zero-sum games

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

$$u_1(\sigma_1^*, \sigma_2^*) \geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2^*) &\geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \end{aligned}$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} \; u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} \; u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min\nolimits_{\sigma_2} \; u_1(\sigma_1^*, \sigma_2) = \max\nolimits_{\sigma_1} \; u_1(\sigma_1, \sigma_2^*)$$

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2^*) &\geq \min_{\sigma_2} \; u_1(\sigma_1^*, \sigma_2) \\ &= \max_{\sigma_1} \; u_1(\sigma_1, \sigma_2^*) \\ &\geq u_1(a_1, \sigma_2^*) \qquad \text{for all } a_1 \in [n] \end{aligned}$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Minimax Theorem $\Rightarrow$ Existence of Nash Eq. in zero-sum games $\square$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Linear Program

$$\begin{aligned} \max \quad & \mu \\ \text{s.t.} \quad & \mu \leq u_1(\sigma_1, b) \qquad \text{for all } b \in [n] \\ & \sigma_1 \in \Delta(n) \end{aligned}$$

Fundamental Results

Guaranteed Existence of Nash Equilibria

- In two-player zero-sum games [von Neumann 1928]
- In finite games [Nash 1950]

John von Neumann

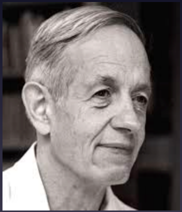

John Nash

von Neumann

- Poly. time
  [Dantzig'51, Khachiyan'79]

- Distributed algorithms
  [Robinson'51, Blackwell'56,
  Hannan'57...]

von Neumann

- Poly. time [Dantzig'51, Khachiyan'79]
- Distributed algorithms [Robinson'51, Blackwell'56, Hannan'57...]

Nash

- Existence uses fixed pt thm/Sperner's Lemma
- Intense effort for computation [Kuhn'61, Lemke-Howson'64, Scarf'67,...]

No efficient algo for 60+ years.

Equilibrium from a computational perspective

Equilibrium from a computational perspective

“If your laptop cannot find the equilibrium, how can the market?”
Kamal Jain

Equilibrium from a computational perspective

NP-completeness is the standard framework to establish the computational intractability of problems [Cook'71, Karp'72]

Equilibrium from a computational perspective

NP-completeness is the standard framework to establish the computational intractability of problems [Cook'71, Karp'72]

NP does not seem to be the right complexity class to capture equilibria

Equilibrium from a computational perspective

What in Nash's theorem is hard to simulate computationally?

What in Nash's theorem is hard to simulate computationally?
The non-constructive step in Nash's theorem?

Parity Lemma

If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

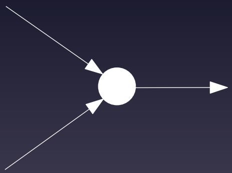

Sperner's Lemma: Color the boundary using three colors in a legal way. No matter how the internal nodes are colored, there exists a tri-chromatic triangle.

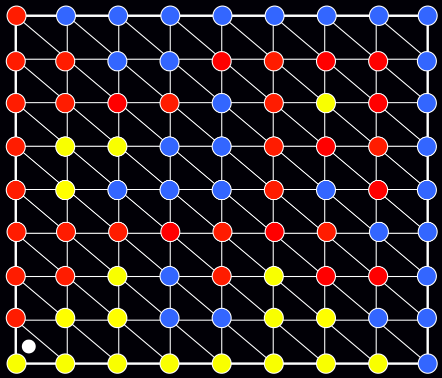

Image Source: Notes by Costis Daskalakis

Sperner's Lemma: Color the boundary using three colors in a legal way. No matter how the internal nodes are colored, there exists a tri-chromatic triangle.

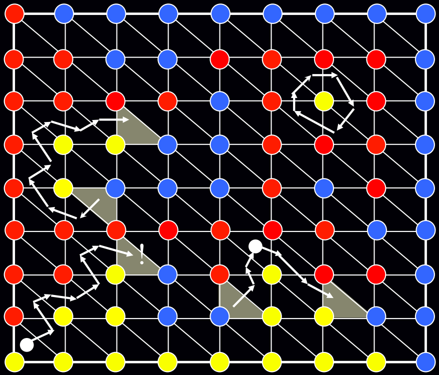

Directed graph over triangles
An unbalanced initial triangle

Sperner's Lemma: Color the boundary using three colors in a legal way. No matter how the internal nodes are colored, there exists a tri-chromatic triangle.

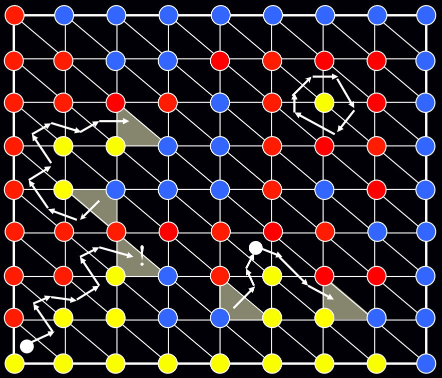

Covering the space of triangles with directed paths, cycles, and isolated vertices

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

Input: Induced graph G with an unbalanced node u

Output: Find another unbalanced node

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

Input: Induced graph G with an unbalanced node u
Output: Find another unbalanced node

G is exponentially large, implicitly defined.

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

Input: Induced graph G with an unbalanced node u

Output: Find another unbalanced node

G is exponentially large, implicitly defined.

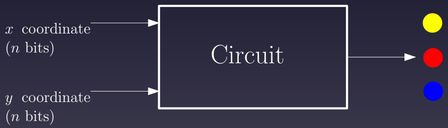

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

$\textsc{Sperner} \in \mathrm{PPAD}$ (complete, even 2d [Chen & Deng'05])

$\textsc{Brouwer} \in \mathrm{PPAD}$ [Papadimitriou'94]

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

$\textsc{Sperner} \in \mathrm{PPAD}$ (complete, even 2d [Chen & Deng'05])

$\textsc{Brouwer} \in \mathrm{PPAD}$ [Papadimitriou'94]

$\textsc{Nash}$ is PPAD-complete [DGP'05, CD'06]

$\textsc{Nash}$ is as hard as any other problem in this complexity class

Complexity class PPAD [Papadimitriou'94]

The class of all problems with guaranteed solution by the parity argument:

> If a directed graph has an unbalanced node (a node with in-degree $\neq$ out-degree), then it must have another.

$\textsc{Sperner} \in \mathrm{PPAD}$ (complete, even 2d [Chen & Deng'05])

$\textsc{Brouwer} \in \mathrm{PPAD}$ [Papadimitriou'94]

$\textsc{Nash}$ is PPAD-complete [DGP'05, CD'06]

$\textsc{Nash}$ is as hard as any other problem in this complexity class

Bypassing complexity barriers: approximation, structure,...

Selected References:

- *Algorithmic Game Theory.* Nisan et al.
- *Algorithmic Game Theory.* CACM (2010) review article by Roughgarden.
- *Computing Equilibria: A Computational Complexity Perspective.* Roughgarden, Economic Theory, 2010.

Selected References:

- *Algorithmic Game Theory.* Nisan et al.
- *Algorithmic Game Theory.* CACM (2010) review article by Roughgarden.
- *Computing Equilibria: A Computational Complexity Perspective.* Roughgarden, Economic Theory, 2010.

Thank You!