CS 590.2

(Computing) Nash Equilibria and Correlated Equilibria

Yu Cheng

Nash equilibrium

[Nash 50]

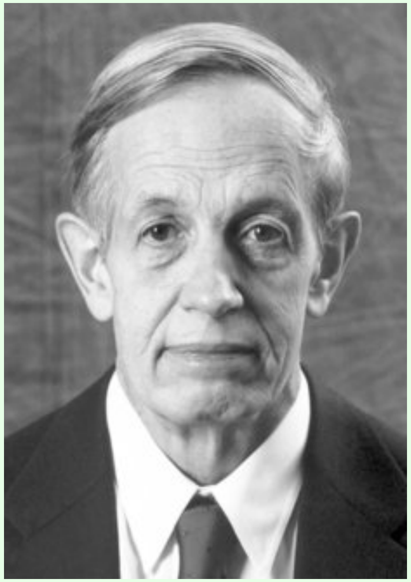

- One mixed strategy for each player
- Every player knows the mixed strategies of the other players
- No player has incentive to deviate

Nash equilibrium
[Nash 50]

- A vector of strategies (one for each player) is called a strategy profile
- A strategy profile $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ is a Nash equilibrium if each σ_i is a best response to σ_{-i}
  - That is, for any i, for any σ_i', $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$
- Note that this does not say anything about multiple agents changing their strategies at the same time
- In any (finite) game, at least one Nash equilibrium (possibly using mixed strategies) exists [Nash 50]
- (Note - singular: equilibrium, plural: equilibria)

Nash equilibria of “chicken”

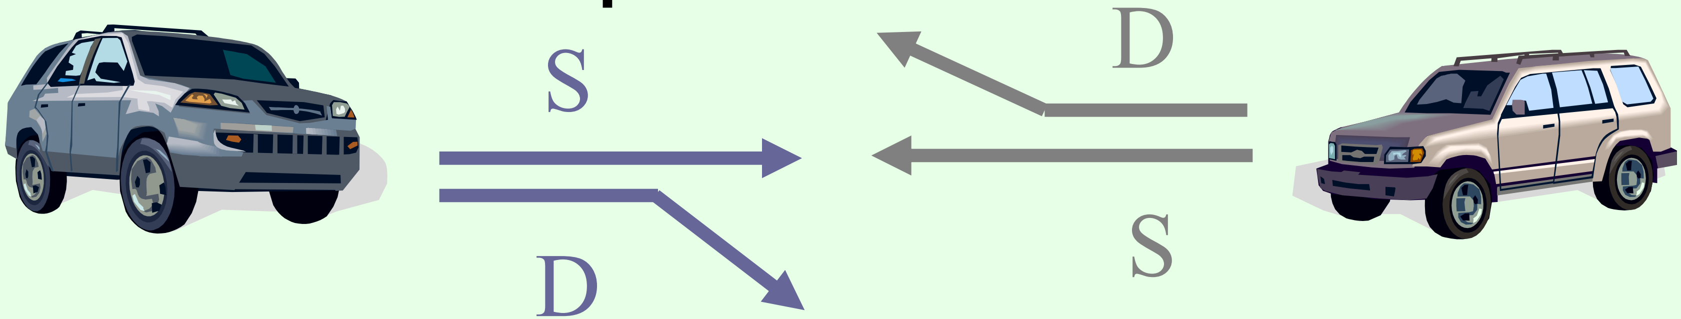

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- (D, S) and (S, D) are Nash equilibria
  - They are pure-strategy Nash equilibria: nobody randomizes
  - They are also strict Nash equilibria: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

Nash equilibria of "chicken"…

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 uses a mixed strategy?
- Recall: if a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- So we need to make player 1 indifferent between D and S
- Player 1's utility for playing D = $-p^c_S$
- Player 1's utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
  - People may die! Expected utility -1/5 for each player

Nash's Proof and PPAD

(Slides borrowed from MIT **Topics in Algorithmic Game Theory** course by Constantinos Daskalakis)

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

$\mapsto$

f: $[0,1]^2 \rightarrow [0,1]^2$, continuous
such that
fixed points $\equiv$ Nash eq.

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

↦

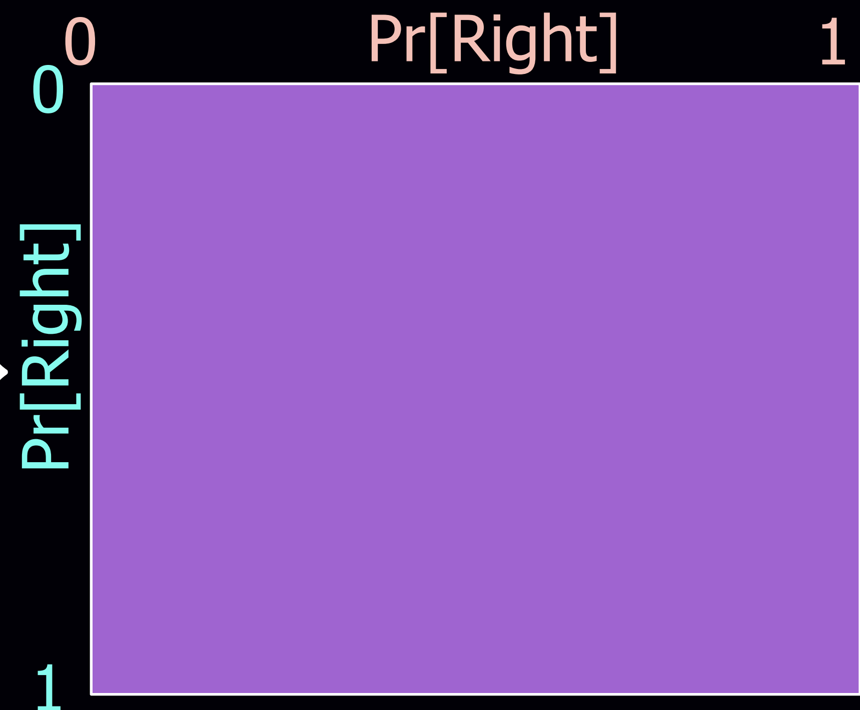

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

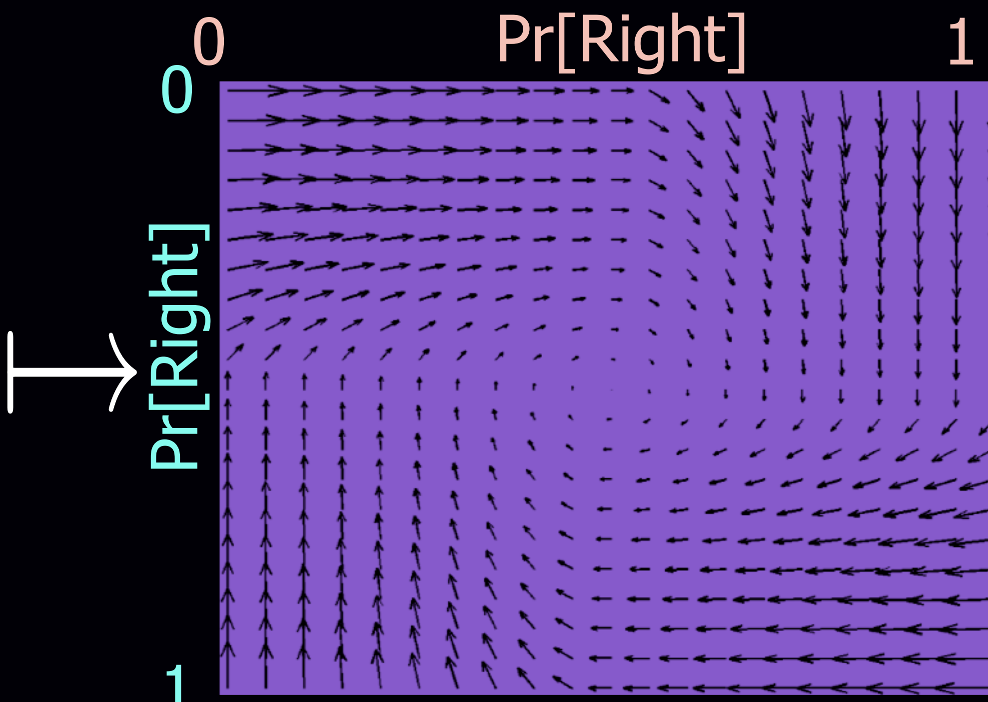

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

$\mapsto$

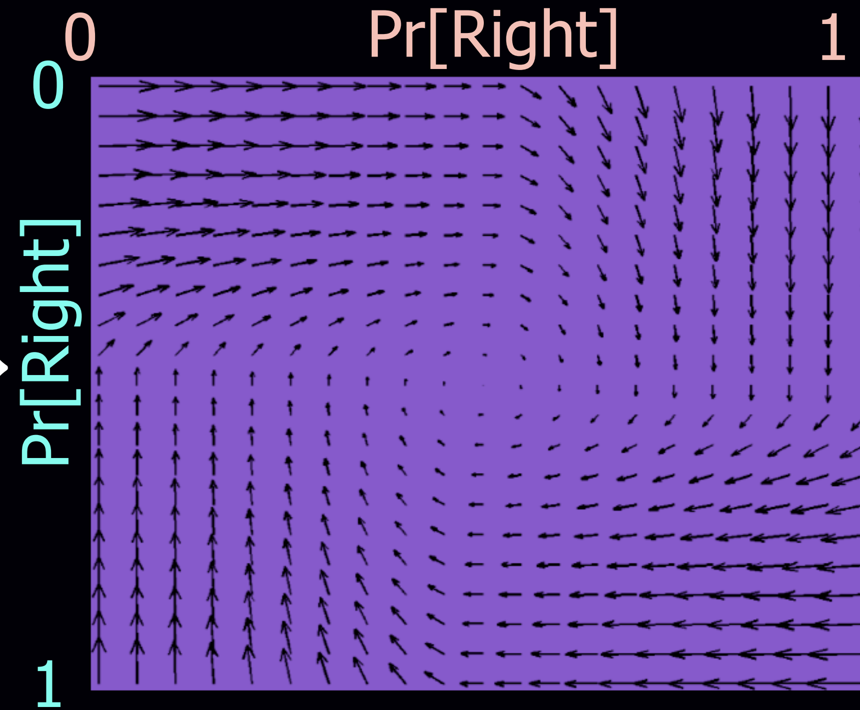

Visualizing Nash's Proof

		½	½
	Kick / Dive	Left	Right
½	Left	1 , -1	-1 , 1
½	Right	-1 , 1	1, -1

Penalty Shot Game

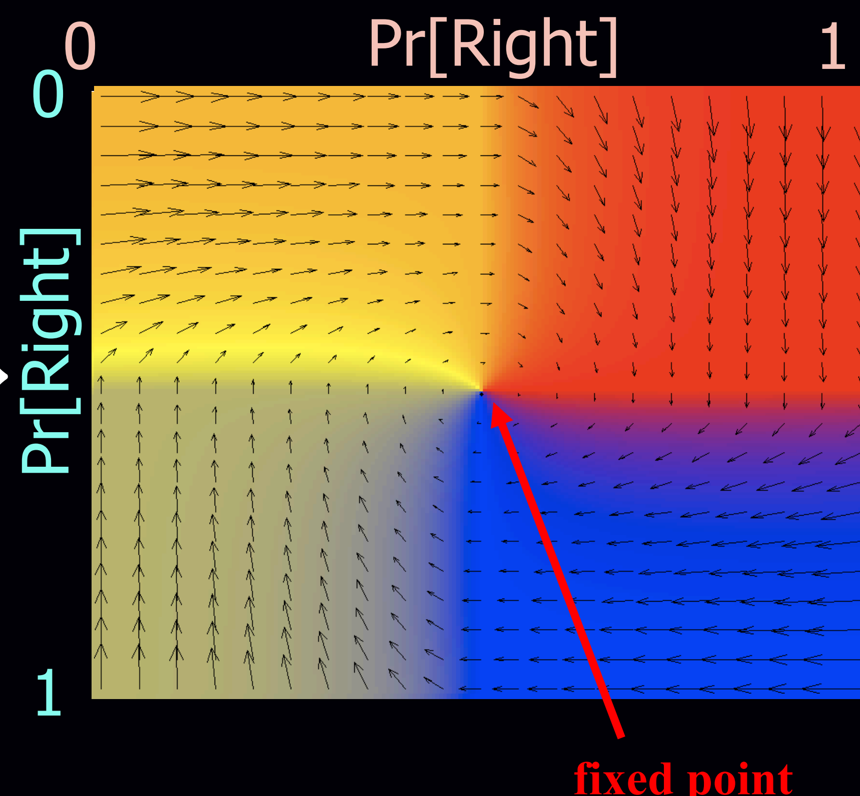

Sperner's Lemma

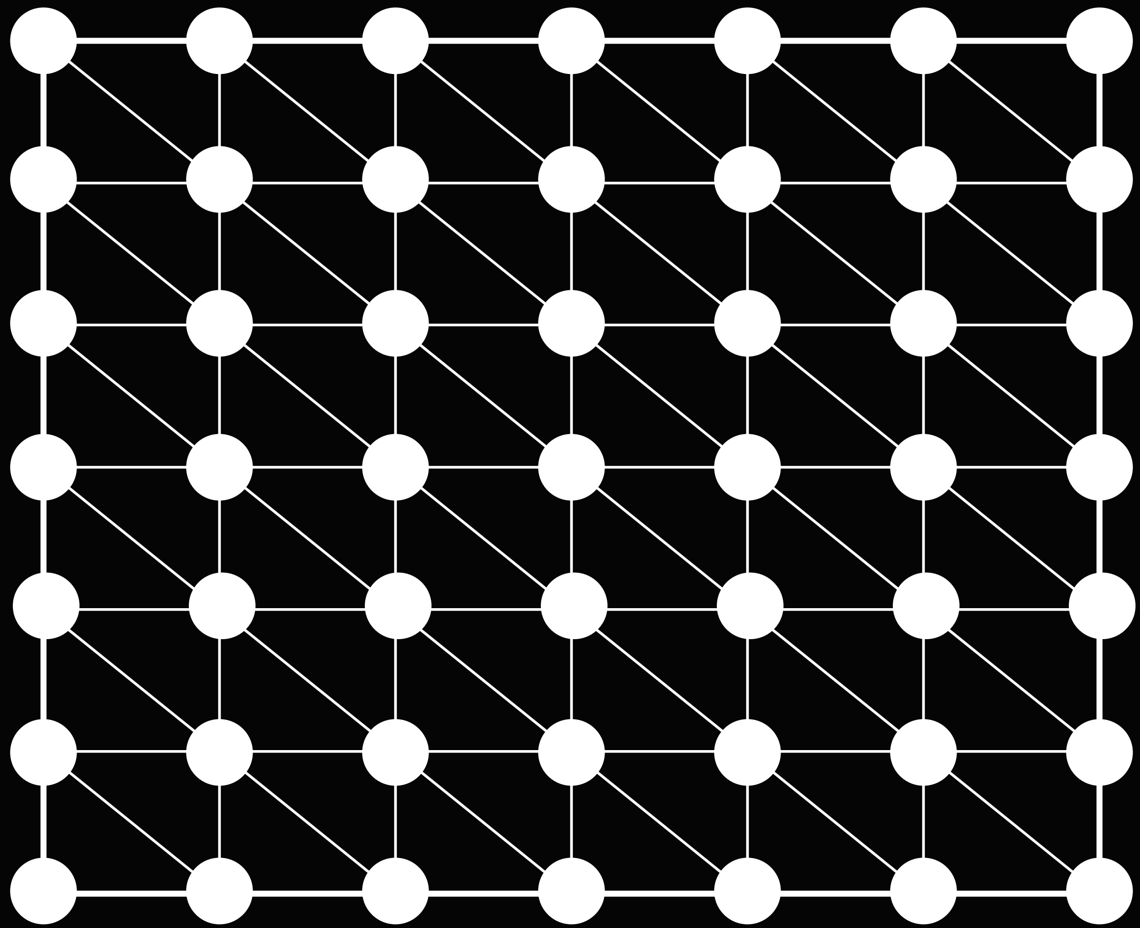

Sperner's Lemma

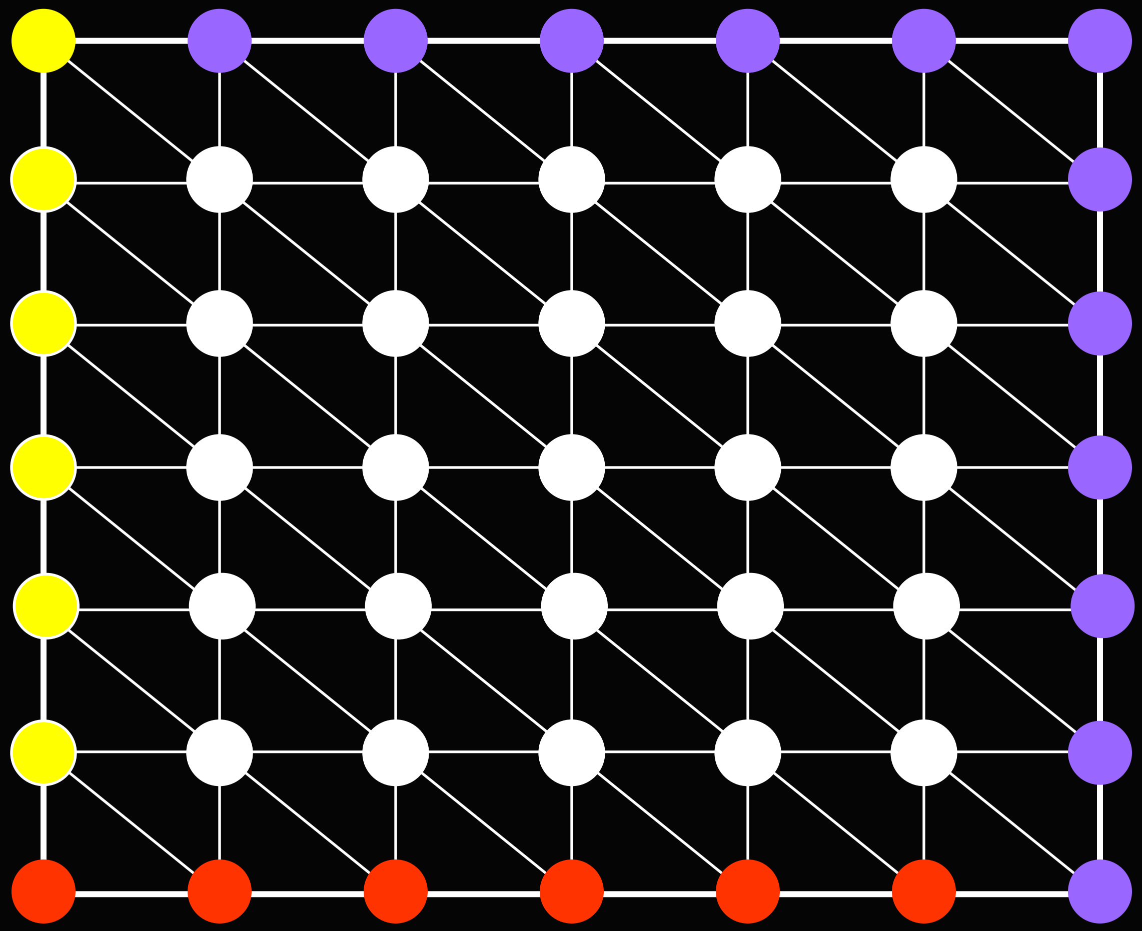

Sperner's Lemma

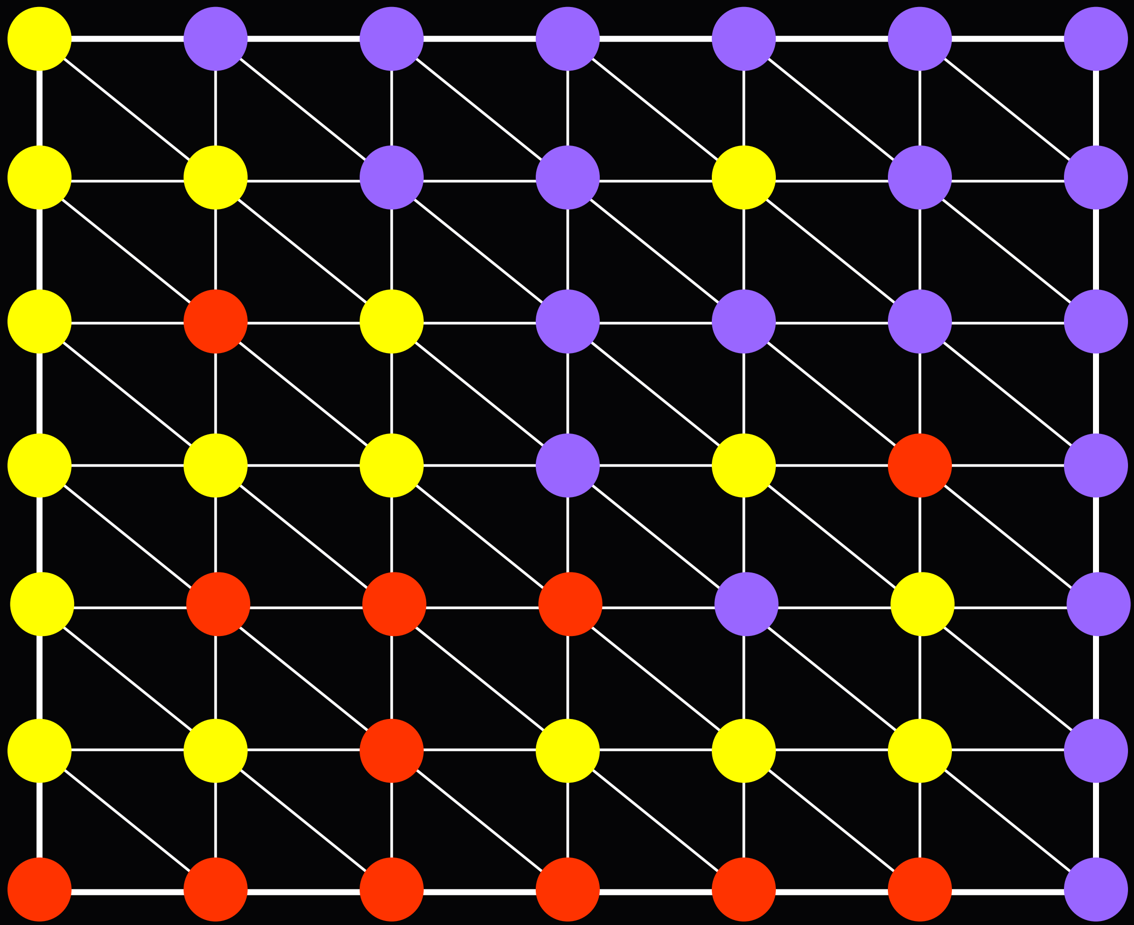

Lemma: No matter how the internal nodes are colored there exists a tri-chromatic triangle. In fact, an odd number of them.

Sperner's Lemma

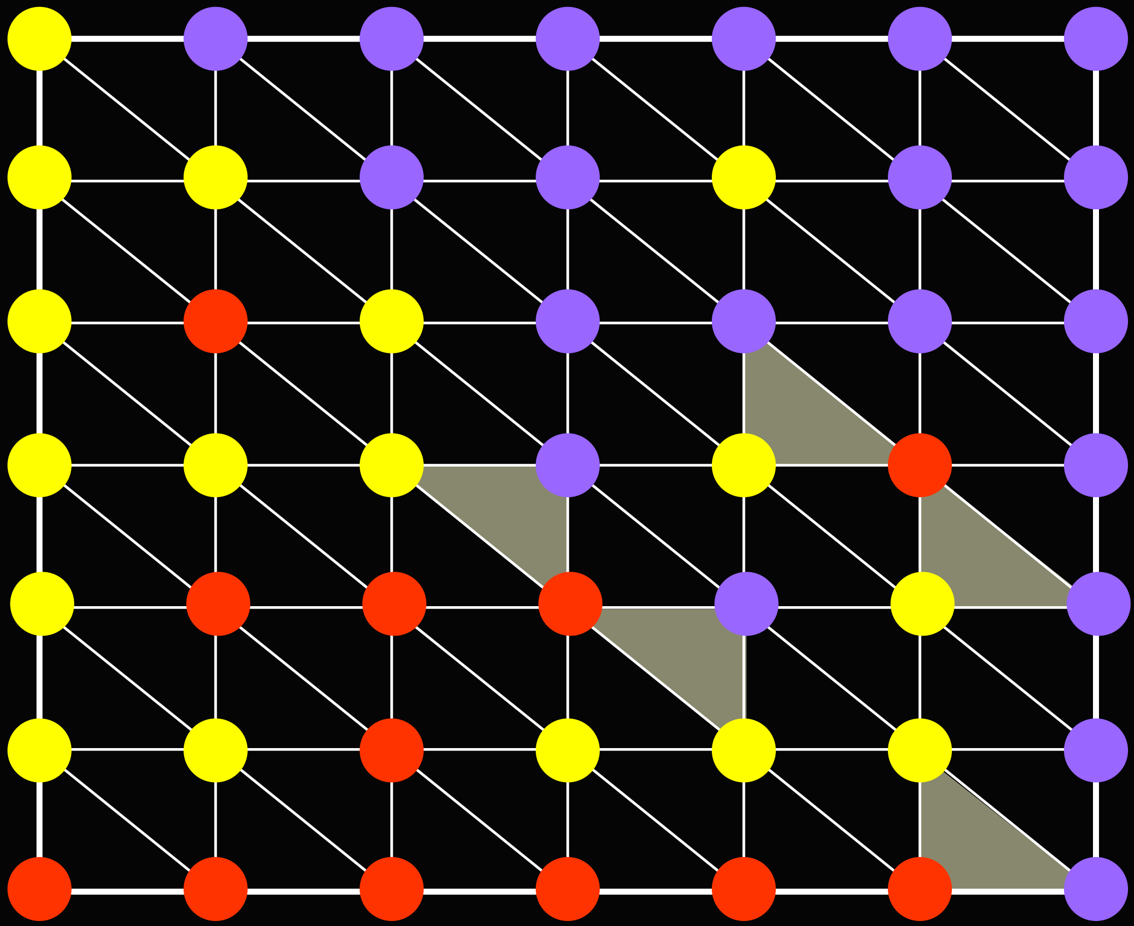

Lemma: No matter how the internal nodes are colored there exists a tri-chromatic triangle. In fact, an odd number of them.

Sperner's Lemma

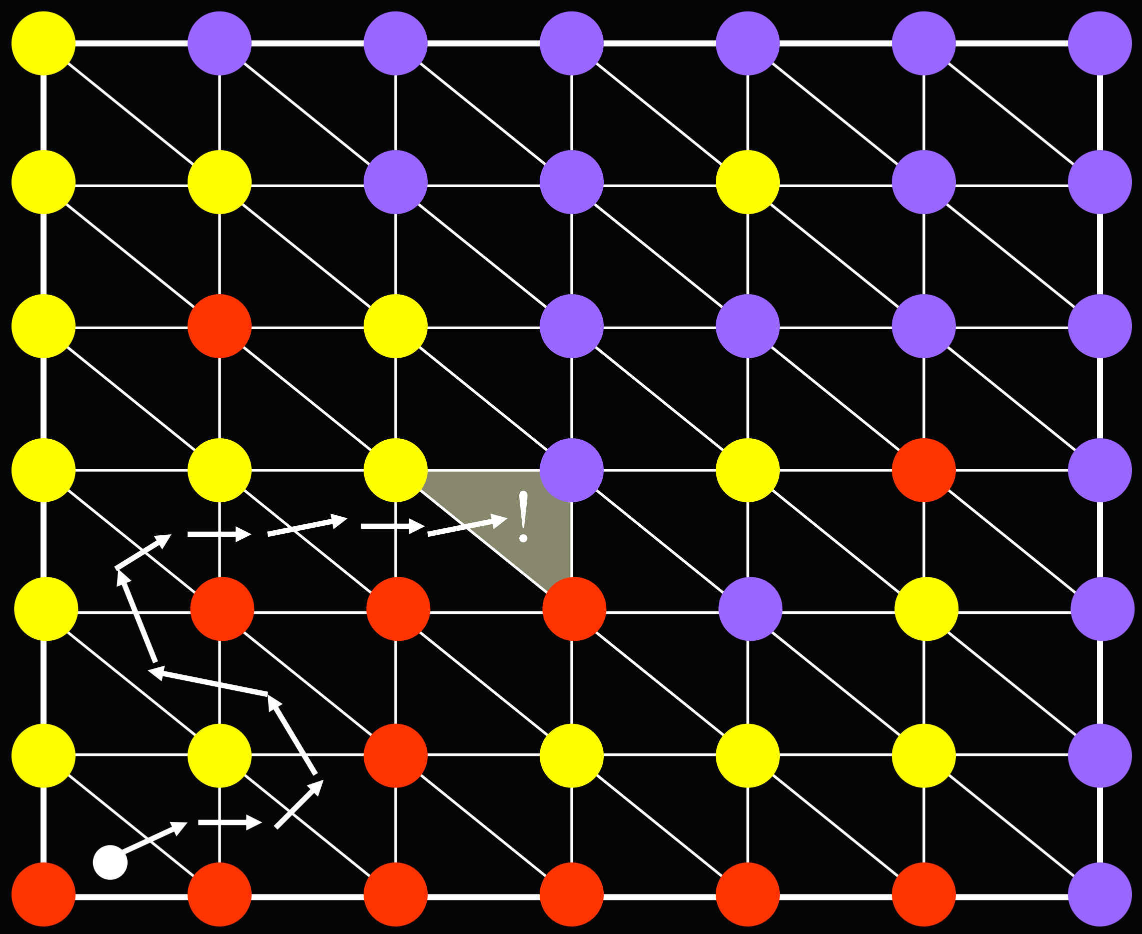

Lemma: No matter how the internal nodes are colored there exists a tri-chromatic triangle. In fact, an odd number of them.

Sperner's Lemma

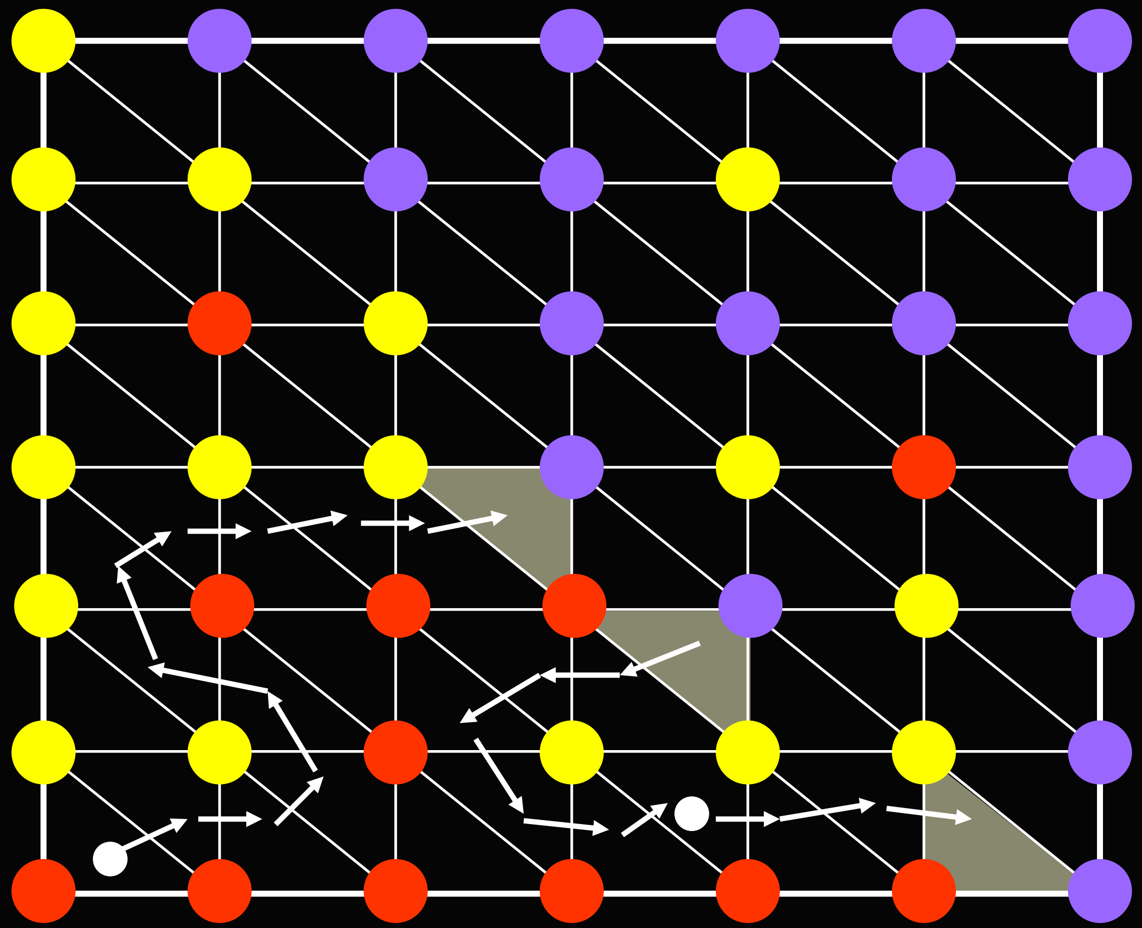

Lemma: No matter how the internal nodes are colored there exists a tri-chromatic triangle. In fact, an odd number of them.

Sperner's Lemma

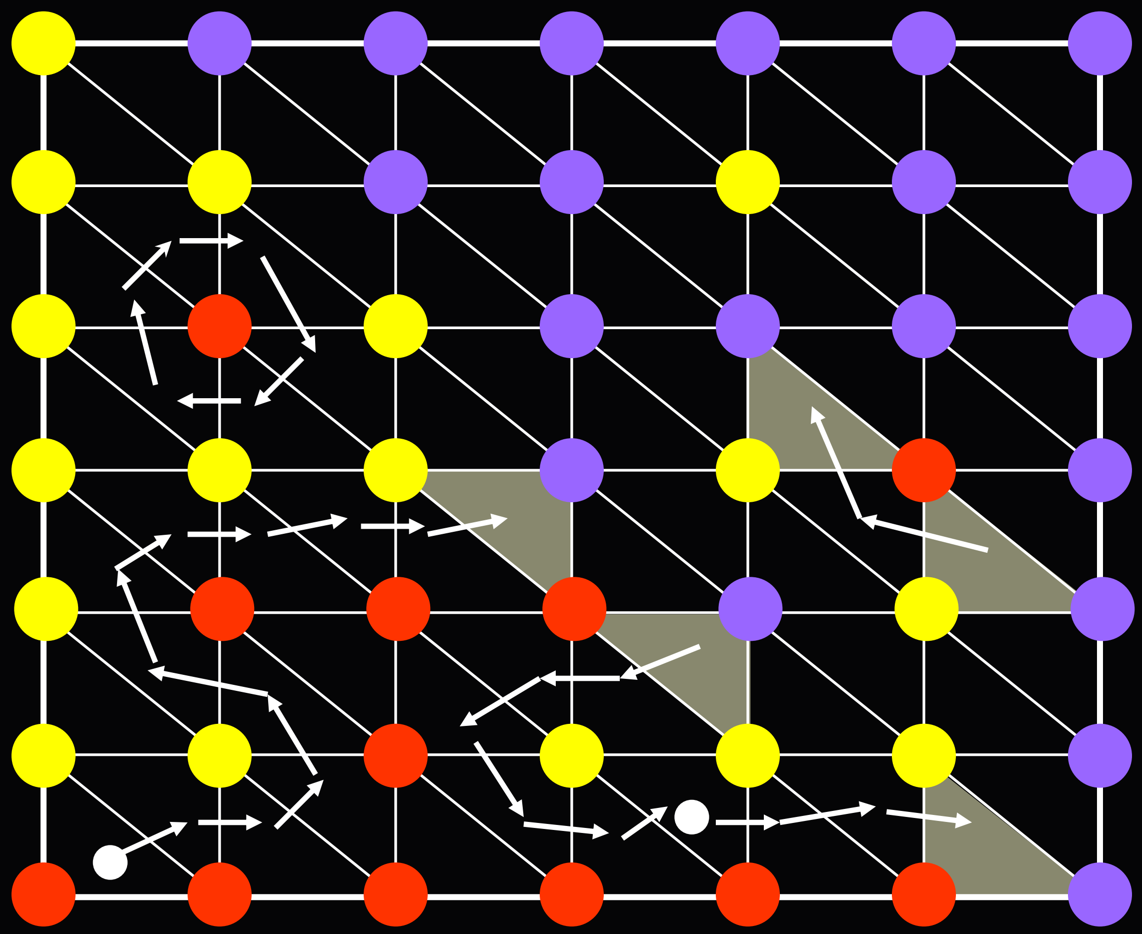

Lemma: No matter how the internal nodes are colored there exists a tri-chromatic triangle. In fact, an odd number of them.

Sperner's Lemma

Space of Triangles

*Transition Rule: If ∃ **red** - **yellow** door cross it with **yellow** on your left hand*

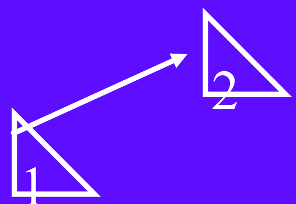

Sperner's Lemma

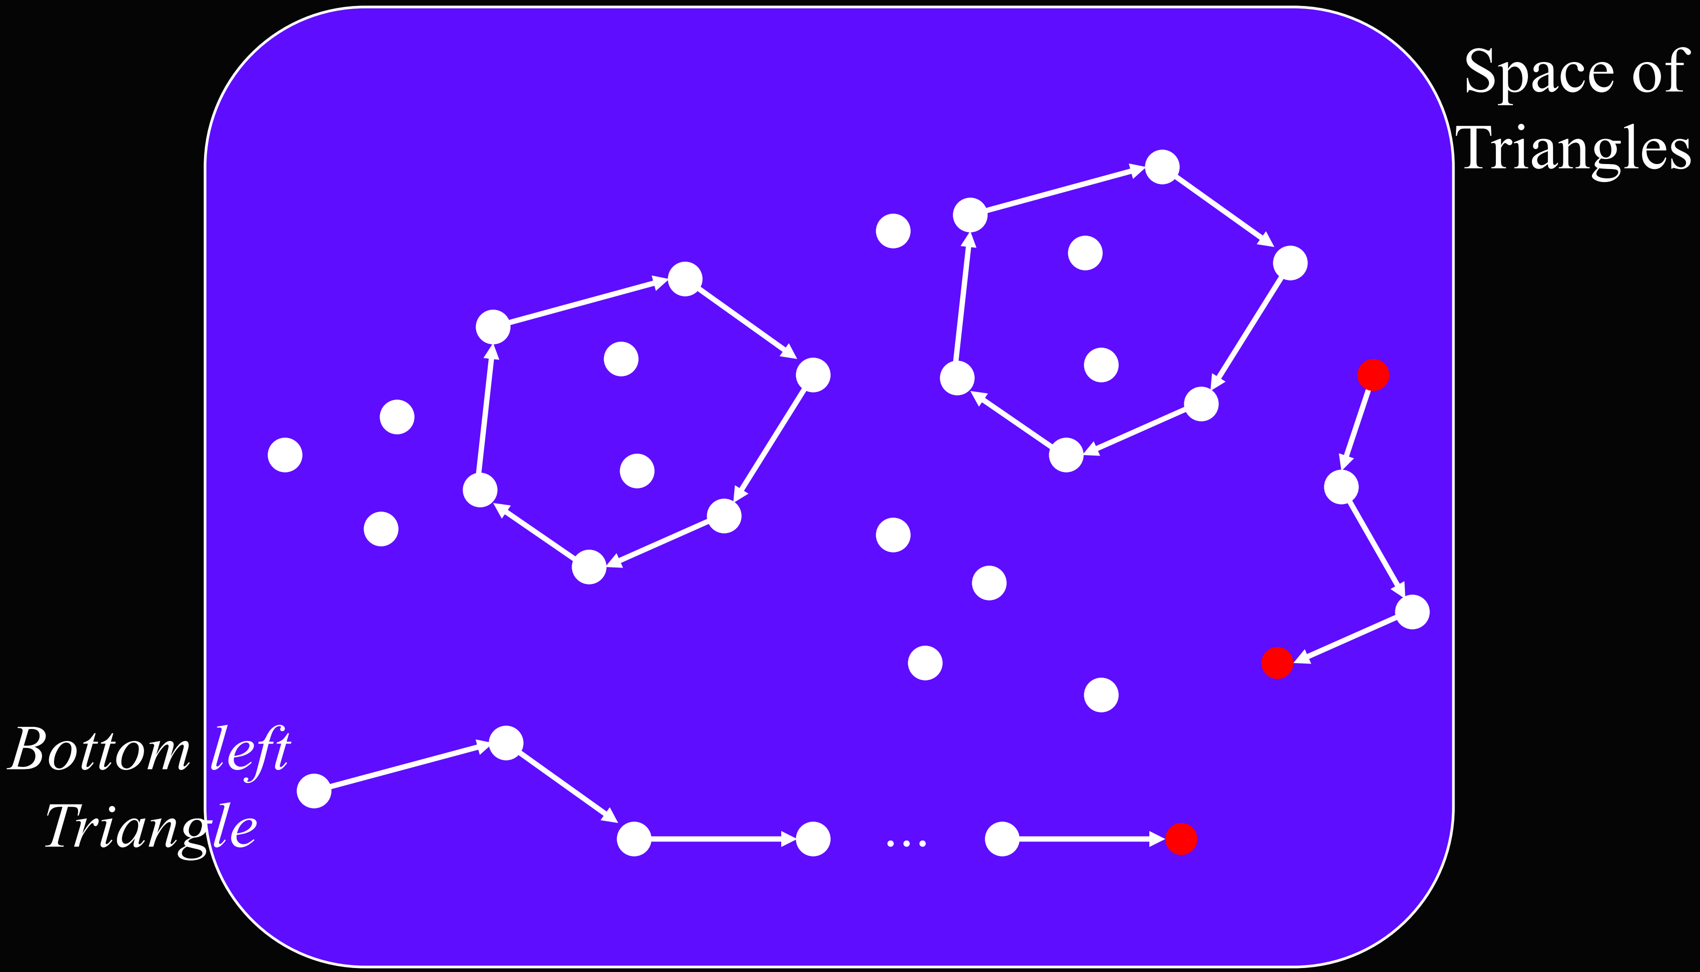

The PPAD Class [Papadimitriou'94]

The class of all problems with guaranteed solution by dint of the following graph-theoretic lemma

> *A directed graph with an unbalanced node (node with indegree ≠ outdegree) must have another.*

Such problems are defined by a directed graph *G*, and an unbalanced node *u* of *G; they require finding another unbalanced node.*
e.g. finding a Sperner triangle is in PPAD

But wait a second…given an unbalanced node in a directed graph, why is it not trivial to find another?

Solving SPERNER

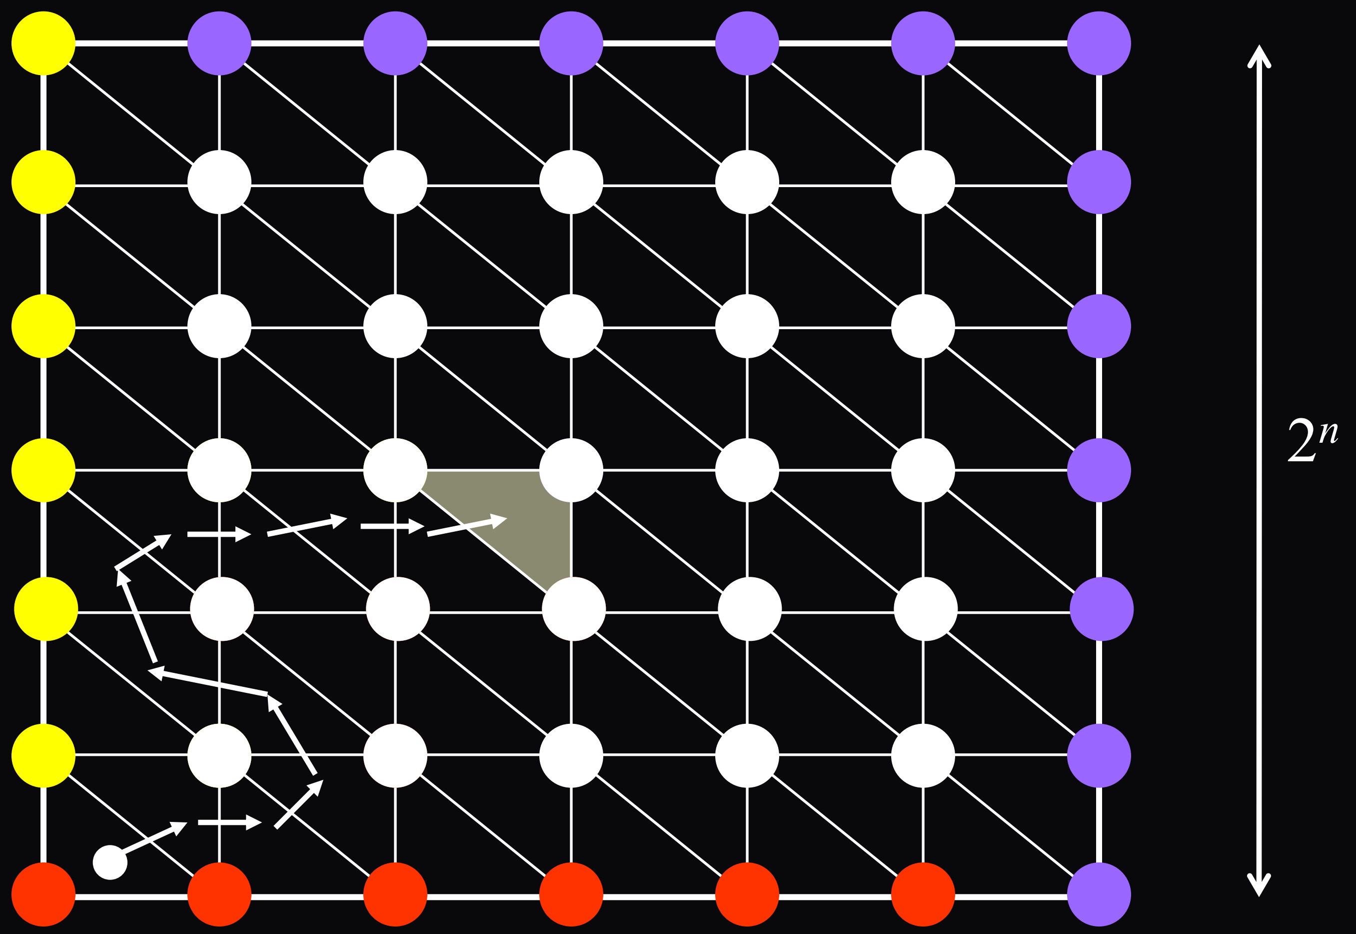

However, the walk may wonder in the box for a long time, before locating the tri-chromatic triangle. Worst-case: 2^{2n}.

The PPAD Class

The class of all problems with guaranteed solution by dint of the following graph-theoretic lemma

> *A directed graph with an unbalanced node (node with indegree $\neq$ outdegree) must have another.*

Nash, Fixed Point, Sperner $\in$ PPAD

Where is PPAD located w.r.t. NP?

(Believed) Location of PPAD

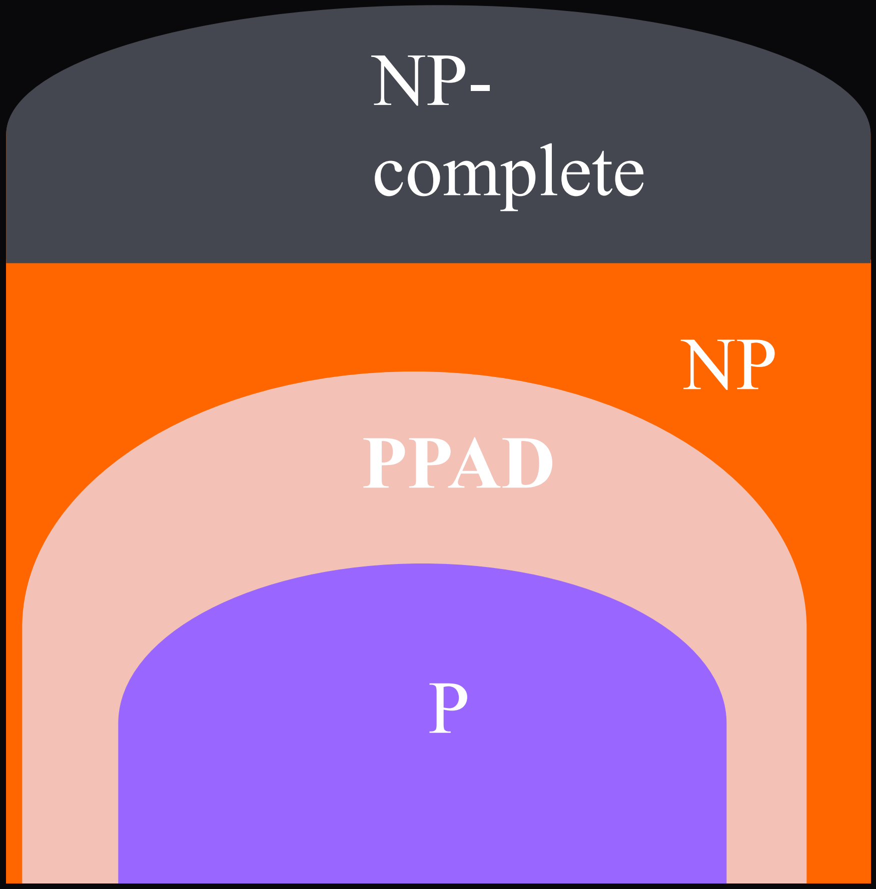

Finding Nash Equilibrium

Games and Computation

- [Nash 50] Every finite game has an equilibrium point
  - Finding it requires solving hard problems

Games and Computation

If one can find an (approximate) equilibrium

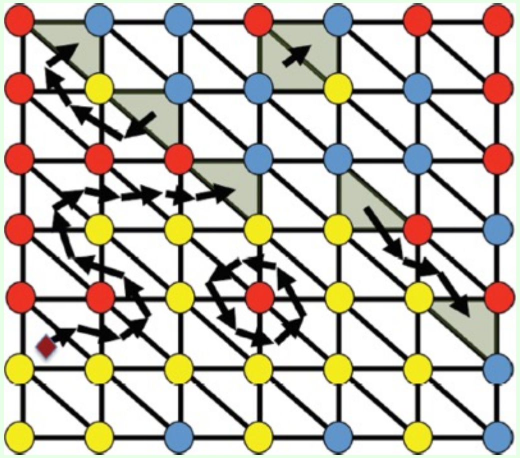

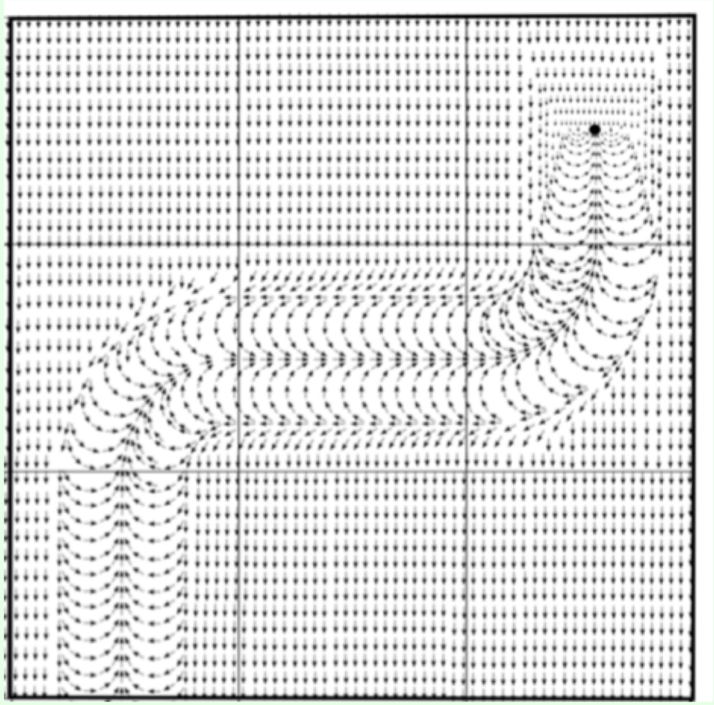

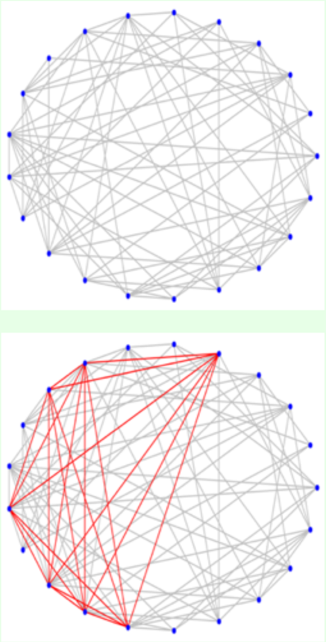

How hard is it to compute *one* (any) Nash equilibrium?

- Complexity was open for a long time
  - [Papadimitriou STOC01]: "together with factoring […] the most important concrete open question on the boundary of P today"
- Recent sequence of papers shows that computing one (any) Nash equilibrium is PPAD-complete (even in 2-player games) [Daskalakis, Goldberg, Papadimitriou 2006; Chen, Deng 2006]
- All known algorithms require exponential time (in the worst case)

The presentation game

		Presenter	
		Put effort into presentation (E)	*Do not put effort into presentation (NE)*
Audience	*Pay attention (A)*	4, 4	-16, -14
	Do not pay attention (NA)	0, -2	0, 0

- Pure-strategy Nash equilibria: (A, E), (NA, NE)
- Mixed-strategy Nash equilibrium:
  ((1/10 A, 9/10 NA), (4/5 E, 1/5 NE))
  – Utility 0 for audience, -14/10 for presenter
  – Can see that some equilibria are strictly better for both players than other equilibria, i.e., some equilibria Pareto-dominate other equilibria

The "equilibrium selection problem"

- You are about to play a game that you have never played before with a person that you have never met
- According to which equilibrium should you play?
- Possible answers:
  - Equilibrium that maximizes the sum of utilities (social welfare)
  - Or, at least not a Pareto-dominated equilibrium
  - So-called focal equilibria
    - "Meet in Paris" game - you and a friend were supposed to meet in Paris at noon on Sunday, but you forgot to discuss where and you cannot communicate. All you care about is meeting your friend. Where will you go?
  - Equilibrium that is the convergence point of some learning process
  - An equilibrium that is easy to compute
  - …
- Equilibrium selection is a difficult problem

What if we want to compute a Nash equilibrium with a specific property?

- For example:
  - An equilibrium that maximizes the expected social welfare (= the sum of the agents' utilities)
  - An equilibrium that maximizes the expected utility of the worst-off player
  - An equilibrium that is not Pareto-dominated
  - An equilibrium that maximizes the expected utility of a given player
  - An equilibrium in which a given pure strategy is played with positive probability
  - An equilibrium in which a given pure strategy is played with zero probability
  - …
- All of these are NP-hard (and the optimization questions are inapproximable assuming P ≠ NP), even in 2-player games [Gilboa, Zemel 89; Conitzer & Sandholm IJCAI-03/GEB-08]

Some properties of Nash equilibria

- If you can eliminate a strategy using strict dominance or even iterated strict dominance, it will not occur in any (i.e., it will be played with probability 0 in every) Nash equilibrium
  - Weakly dominated strategies may still be played in some Nash equilibrium
- In 2-player zero-sum games, a profile is a Nash equilibrium if and only if both players play minimax strategies
  - Hence, in such games, if (σ_1, σ_2) and (σ_1', σ_2') are Nash equilibria, then so are (σ_1, σ_2') and (σ_1', σ_2)
    - No equilibrium selection problem here!

Search-based approaches (for 2 players)

- Suppose we know the support X_i of each player i's mixed strategy in equilibrium
  - That is, which pure strategies receive positive probability
- Then, we have a linear feasibility problem:
  - for both i, for any $s_i \in S_i - X_i$, $\mathbf{p_i(s_i)} = 0$
  - for both i, for any $s_i \in X_i$, $\Sigma \mathbf{p_{-i}(s_{-i})} u_i(s_i, s_{-i}) = \mathbf{u_i}$
  - for both i, for any $s_i \in S_i - X_i$, $\Sigma \mathbf{p_{-i}(s_{-i})} u_i(s_i, s_{-i}) \leq \mathbf{u_i}$
- Thus, we can search over possible supports
  - This is the basic idea underlying methods in [Dickhaut & Kaplan 91; Porter, Nudelman, Shoham AAAI04/GEB08]
- Dominated strategies can be eliminated

Solving for a Nash equilibrium using MIP (2 players)

[Sandholm, Gilpin, Conitzer AAAI05]

- maximize *whatever you like (e.g., social welfare)*
- subject to
  - for both i, for any s_i, $\Sigma_{s_{-i}} \mathbf{p}_{\mathbf{s}_{-i}} u_i(s_i, s_{-i}) = \mathbf{u}_{\mathbf{s}_i}$
  - for both i, for any s_i, $\mathbf{u}_i \geq \mathbf{u}_{\mathbf{s}_i}$
  - for both i, for any s_i, $\mathbf{p}_{\mathbf{s}_i} \leq \mathbf{b}_{\mathbf{s}_i}$
  - for both i, for any s_i, $\mathbf{u}_i - \mathbf{u}_{\mathbf{s}_i} \leq M(1- \mathbf{b}_{\mathbf{s}_i})$
  - for both i, $\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} = 1$
- $\mathbf{b}_{\mathbf{s}_i}$ is a binary variable indicating whether s_i is in the support, M is a large number

Lemke-Howson algorithm (1-slide sketch!)

	GREEN	ORANGE
RED	1, 0	0, 1
BLUE	0, 2	1, 0

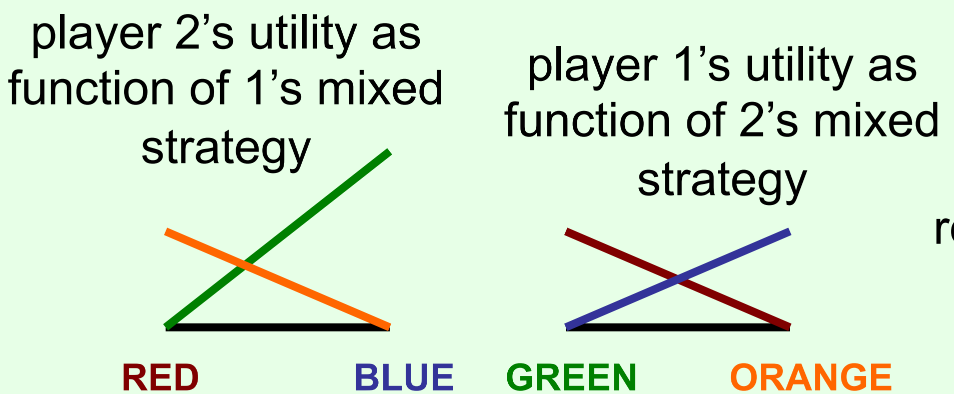

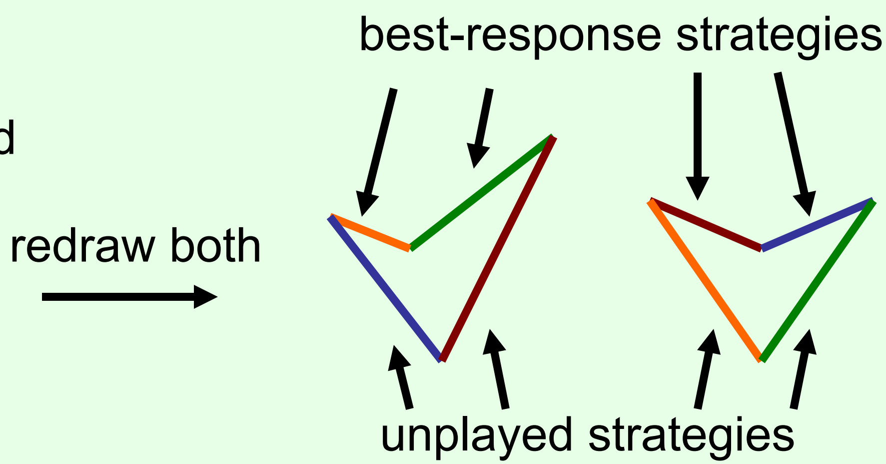

- Strategy profile = pair of points
- Profile is an equilibrium iff every pure strategy is either a best response or unplayed
- I.e. equilibrium = pair of points that includes all the colors
  - … except, pair of bottom points doesn't count (the "artificial equilibrium")
- Walk in some direction from the artificial equilibrium; at each step, throw out the color used twice

Correlated equilibrium [Aumann 74]

- Suppose there is a trustworthy mediator who has offered to help out the players in the game
- The mediator chooses a profile of pure strategies, perhaps randomly, then tells each player what her strategy is in the profile (but not what the other players' strategies are)
- A correlated equilibrium is a distribution over pure-strategy profiles so that every player wants to follow the recommendation of the mediator (if she assumes that the others do so as well)
- Every Nash equilibrium is also a correlated equilibrium
  - Corresponds to mediator choosing players' recommendations independently
- … but not vice versa
- *(Note: there are more general definitions of correlated equilibrium, but it can be shown that they do not allow you to do anything more than this definition.)*

A correlated equilibrium for "chicken"

	D	S
D	0, 0 20%	-1, 1 40%
S	1, -1 40%	-5, -5 0%

- Why is this a correlated equilibrium?
- Suppose the mediator tells the row player to Dodge
- From Row's perspective, the conditional probability that Column was told to Dodge is 20% / (20% + 40%) = 1/3
- So the expected utility of Dodging is (2/3)*(-1) = -2/3
- But the expected utility of Straight is (1/3)*1 + (2/3)*(-5) = -3
- So Row wants to follow the recommendation
- If Row is told to go Straight, he knows that Column was told to Dodge, so again Row wants to follow the recommendation
- Similar for Column

A nonzero-sum variant of rock-paper-scissors (Shapley's game [Shapley 64])

		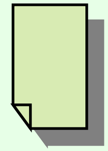	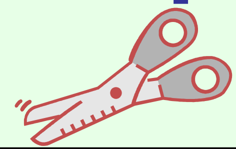
	0, 0 0	0, 1 1/6	1, 0 1/6
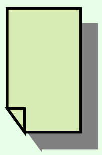	1, 0 1/6	0, 0 0	0, 1 1/6
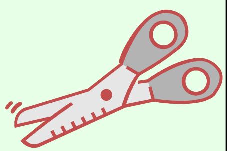	0, 1 1/6	1, 0 1/6	0, 0 0

- If both choose the same pure strategy, both lose
- These probabilities give a correlated equilibrium:
- E.g. suppose Row is told to play Rock
- Row knows Column is playing either paper or scissors (50-50)
  - Playing Rock will give ½; playing Paper will give 0; playing Scissors will give ½
- So Rock is optimal (not uniquely)

Solving for a correlated equilibrium using linear programming (n players!)

- Variables are now $\mathbf{p_s}$ where s is a profile of pure strategies
- maximize *whatever you like (e.g., social welfare)*
- subject to
  - for any i, s_i, s_i', $\Sigma_{s_{-i}} \mathbf{p_{(s_i, s_{-i})}} u_i(s_i, s_{-i}) \geq \Sigma_{s_{-i}} \mathbf{p_{(s_i, s_{-i})}} u_i(s_i', s_{-i})$
  - $\Sigma_s \mathbf{p_s} = 1$