Algorithmic Game Theory
Solution concepts in games

Georgios Amanatidis
amanatidis@diag.uniroma1.it

Based on slides by V. Markakis and A. Voudouris

Solution concepts

Choosing a strategy...

- Given a game, how should a player choose his strategy?
  - Recall: we assume each player knows the other players' preferences but not what the other players will choose

- The most fundamental question of game theory
  - Clearly, the answer is not always clear

- We will start with 2-player games

Prisoner's Dilemma: The Rational Outcome

D dom . . .

- Let's revisit prisoner's dilemma
- Reasoning of pl. 1:
  - If pl. 2 does not confess, then I should confess
  - If pl. 2 confesses, then I should also confess
- Similarly for pl. 2

	C	D
C	3, 3	0, 4
D	4, 0	1, 1

D dominates C for P1

- Expected outcome for rational players: they will both confess, and they will go to jail for 3 years each
  - Observation: If they had both chosen not to confess, they would go to jail only for 1 year, each of them would have a strictly better utility

Dominant strategies

- Ideally, we would like a strategy that would provide the best possible outcome, regardless of what other players choose
- Definition: A strategy s_i of pl. 1 is *dominant* if

$$u_1(s_i, t_j) \geq u_1(s', t_j)$$

for every strategy $s' \in S^1$ and every strategy $t_j \in S^2$

- Similarly for pl. 2, a strategy t_j is dominant if

$$u_2(s_i, t_j) \geq u_2(s_i, t')$$

for every strategy $t' \in S^2$ and for every strategy $s_i \in S^1$

→ row i contains the max coordinate for every single column.

Dominant strategies

Even better:

- Definition: A strategy s_i of pl. 1 is *strictly dominant* if

$$u_1 (s_i, t_j) > u_1 (s', t_j)$$

for every strategy $s' \in S^1 \setminus \{s_i\}$ and every strategy $t_j \in S^2$

- Similarly for pl. 2
- In prisoner's dilemma, strategy D (confess) is strictly dominant

Observations:

- There may be more than one dominant strategies for a player, but then they should yield the same utility under all profiles
- Every player can have at most one strictly dominant strategy
- A strictly dominant strategy is also dominant

Existence of dominant strategies

- Few games possess dominant strategies
- It may be too much to ask for
- E.g. in the Bach-or-Stravinsky game, there is no dominant strategy:
  - Strategy B is not dominant for pl. 1:
    If pl. 2 chooses S, pl. 1 should choose S
  - Strategy S is also not dominant for pl. 1:
    If pl. 2 chooses B, pl. 1 should choose B

	B	S
B	(2, 1)	(0, 0)
S	(0, 0)	(1, 2)

- In all the examples we have seen so far, only prisoner’s dilemma possesses dominant strategies

Back to choosing a strategy...

- Hence, the question of how to choose strategies still remains for the majority of games
- Model of rational choice: if a player knows or has a strong belief for the choice of the other player, then he should choose the strategy that maximizes his utility
- Suppose that someone suggests to the 2 players the strategy profile (s, t)
- When would the players be willing to follow this profile?
  - For pl. 1 to agree, it should hold that

    $u_1(s, t) \geq u_1(s', t)$ for every other strategy s' of pl. 1
  - For pl. 2 to agree, it should hold that

    $u_2(s, t) \geq u_2(s, t')$ for every other strategy t' of pl. 2

Nash Equilibria

- Definition (Nash 1950): A strategy profile (s, t) is a Nash equilibrium, if no player has a unilateral incentive to deviate, given the other player's choice
- This means that the following conditions should be satisfied:
  1. $u_1(s, t) \geq u_1(s', t)$ for every strategy $s' \in S^1$
  2. $u_2(s, t) \geq u_2(s, t')$ for every strategy $t' \in S^2$
- One of the dominant concepts in game theory from 1950s till now
- Most other concepts in noncooperative game theory are variations/extensions/generalizations of Nash equilibria

Pictorially:

			t		
	(,)	(,)	(x_1,)	(,)	(,)
	(,)	(,)	(x_2,)	(,)	(,)
	(,)	(,)	(x_3,)	(,)	(,)
s	(, y_1)	(, y_2)	(x, y)	(, y_4)	(, y_5)
	(,)	(,)	(x_5,)	(,)	(,)

In order for (s, t) to be a Nash equilibrium:

- x must be greater than or equal to any x_i in column t
- y must be greater than or equal to any y_j in row s

Nash Equilibria

- We should think of Nash equilibria as “stable” profiles of a game
  - At an equilibrium, each player thinks that if the other player does not change her strategy, then he also does not want to change his own strategy
- Hence, no player would regret for his choice at an equilibrium profile (s, t)
  - If the profile (s, t) is realized, pl. 1 sees that he did the best possible, against strategy t of pl. 2,
  - Similarly, pl. 2 sees that she did the best possible against strategy s of pl. 1
- Attention: If both players decide to change simultaneously, then we may have profiles where they are both better off

10 players

$S^i = \{R, B\}$

If all are blue they gain 10 €

If all are red they gain 3 €

Otherwise they gain 1 €

B 1€

- Is $(R, R, R, \ldots, R)$ a "good" profile? "NO"

3€

- Does it make sense? NO

- Is it a NE? YES

Examples of finding Nash equilibria in simple games

Example 1: Prisoner's Dilemma

In small games, we can examine all possible profiles and check if they form an equilibrium

- (C, C): both players have an incentive to deviate to another strategy
- (C, D): pl. 1 has an incentive to deviate
- (D, C): Same for pl. 2
- (D, D): Nobody has an incentive to change

	C	D
C	~~3, 3~~	~~0, 4~~
D	~~4, 0~~	1, 1

max of row

max of col

Hence: The profile (D, D) is the unique Nash equilibrium of this game

- Recall that D is a dominant strategy for both players in this game

Corollary: If s is a dominant strategy of pl. 1, and t is a dominant strategy for pl. 2, then the profile (s, t) is a Nash equilibrium

Example 1: Prisoner's Dilemma

In small games, we can examine all possible profiles and check if they form an equilibrium

- (C, C): both players have an incentive to deviate to another strategy
- (C, D): pl. 1 has an incentive to deviate
- (D, C): Same for pl. 2
- (D, D): Nobody has an incentive to change

	C	D
C	3, 3	0, 4
D	4, 0	1, 1

Hence: The profile (D, D) is the unique Nash equilibrium of this game

- Recall that D is a dominant strategy for both players in this game

Corollary: If s is a dominant strategy of pl. 1, and t is a dominant strategy for pl. 2, then the profile (s, t) is a Nash equilibrium

Example 2: Bach or Stravinsky (BoS)

	B	S
B	2, 1	~~0, 0~~
S	~~0, 0~~	1, 2

2 Nash equilibria:

- (B, B) and (S, S)
- Both derive the same total utility (3 units)
- But each player has a preference for a different equilibrium

Example 2a: Coordination games

Variation of Bach
or Stravinsky

	B	S
B	2, 2	0, 0
S	0, 0	1, 1

Again 2 Nash equilibria:

- (B, B) and (S, S)
- But now (B, B) is clearly the most preferable for both players
- Still the profile (S, S) is a valid equilibrium, no player has a unilateral incentive to deviate
  - At the profile (S, S), both players should deviate together in order to reach a better outcome

Example 3: The Hawk-Dove game

		
	2, 2	0, 4
	4, 0	-1, -1

- The most fair solution (D, D) is not an equilibrium
- 2 Nash equilibria: (D, H), (H, D)
- We have a stable situation only when one population dominates or destroys the other

Example 4: Matching Pennies

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

- In every profile, some player has an incentive to deviate
- There is no Nash equilibrium!
- Note: The same is true for Rock-Paper-Scissors

Mixed strategies in games

Existence of Nash equilibria

- We saw that not all games possess Nash equilibria
- E.g. Matching Pennies, Rock-Paper-Scissors, and many others
- What would constitute a good solution in such games?

Example of a game without equilibria: Matching Pennies

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

$u_2(H, H) = 1$

- In every profile, some player has an incentive to change
- Hence, no Nash equilibrium!

Q: How would we play this game in practice?

A: Maybe randomly

Matching Pennies: Randomized strategies

	½ H	½ T
½ H	1, -1	-1, 1
½ T	-1, 1	1, -1

(½, ½) (½, ½)

$$E[u_1(\tau_1, \tau_2)] = \frac{1}{4}\cdot 1 + \frac{1}{4}(-1) + \frac{1}{4}(-1) + \frac{1}{4}\cdot 1 = 0$$

- Main idea: Enlarge the strategy space so that players are allowed to play non-deterministically
- Suppose both players play
  - H with probability 1/2
  - T with probability 1/2
- Then every outcome has a probability of ¼
- For pl. 1:
  - P[win] = P[lose] = ½
  - Average utility = 0
- Similarly for pl. 2

Mixed strategies

- Definition: A mixed strategy of a player is a probability distribution on the set of his available choices
- If $S = (s_1, s_2, \ldots, s_n)$ is the set of available strategies of a player, then a mixed strategy is a vector in the form

  $\mathbf{p} = (p_1, \ldots, p_n)$, where

  $p_i \geq 0$ for $i=1, \ldots, n$, and $p_1 + \ldots + p_n = 1$
- p_j = probability for selecting the j-th strategy
- We can write it also as $p_j=p(s_j)$ = prob/ty of selecting s_j
- Matching Pennies: the uniform distribution can be written as

  $\mathbf{p} = (1/2, 1/2)$ or $p(H) = p(T) = ½$

Pure and mixed strategies

- From now on, we refer to the available choices of a player as *pure strategies* to distinguish them from mixed strategies
- For 2 players with $S^1 = \{s_1, s_2, \ldots, s_n\}$ and $S^2 = \{t_1, t_2, \ldots, t_m\}$
- Pl. 1 has n pure strategies, Pl. 2 has m pure strategies
- Every pure strategy can also be represented as a mixed strategy that gives probability 1 to only a single choice
- E.g., the pure strategy s_1 can also be written as the mixed strategy $(1, 0, 0, \ldots, 0)$
- More generally: strategy s_i can be written in vector form as the mixed strategy $e^i = (0, 0, \ldots, 1, 0, \ldots, 0)$
  - 1 at position i, 0 everywhere else
  - Some times, it is convenient in the analysis to use the vector form for a pure strategy

Utility under mixed strategies

- Suppose that each player has chosen a mixed strategy in a game
- How does a player now evaluate the outcome of a game?
- We will assume that each player cares for his expected utility
  - Justified when games are played repeatedly
  - Not justified for more risk-averse or risk-seeking players

$$E(X) = \sum_{i=1}^{k} p_i x_i$$

$x_1, \dots, x_k$ are the possible values of X and $p_1, \dots, p_k$ their prob's

Expected utility (for 2 players)

- Consider a n x m game
- Pure strategies of pl. 1: $S^1 = \{s_1, s_2, \ldots, s_n\}$
- Pure strategies of pl. 2: $S^2 = \{t_1, t_2, \ldots, t_m\}$
- Let $\mathbf{p} = (p_1, \ldots, p_n)$ be a mixed strategy of pl. 1

and $\mathbf{q} = (q_1, \ldots, q_m)$ be a mixed strategy of pl. 2

- Expected utility of pl. 1:

$E(u_1(p,q))$

$$u_1(\mathbf{p},\mathbf{q}) = \sum_{i=1}^{n}\sum_{j=1}^{m} \overbrace{p_i \cdot q_j}^{\text{probability of cell } (i,j)} \cdot \underbrace{u_1(s_i,t_j)}_{P_1\text{'s utility in cell } (i,j)} = \sum_{i=1}^{n}\sum_{j=1}^{m} p(s_i)\cdot q(t_j)\cdot u_1(s_i,t_j)$$

- Similarly for pl. 2 (replace u_1 by u_2)

Example

		B	S
		1/2	1/2
4/5	B	2, 1 (4/10)	0, 0 (4/10)
1/5	S	0, 0 (1/10)	1, 2 (1/10)

- Let **p** = (4/5, 1/5), **q** = (1/2, 1/2)
- u_1(**p**, **q**) = 4/5 x 1/2 x 2 + 1/5 x 1/2 x 1 = 0.9
- u_2(**p**, **q**) = 4/5 x 1/2 x 1 + 1/5 x 1/2 x 2 = 0.6
- When can we have an equilibrium with mixed strategies?

u_i is really $E(u_i)$

Nash equilibria with mixed strategies

- Definition: A profile of mixed strategies (**p**, **q**) is a Nash equilibrium if
  - $u_1(\mathbf{p}, \mathbf{q}) \geq u_1(\mathbf{p}', \mathbf{q})$ for any other mixed strategy **p**' of pl. 1
  - $u_2(\mathbf{p}, \mathbf{q}) \geq u_2(\mathbf{p}, \mathbf{q}')$ for any other mixed strategy **q**' of pl. 2
- Again, we just demand that no player has a unilateral incentive to deviate to another strategy
- How do we verify that a profile is a Nash equilibrium?
  - There is an infinite number of mixed strategies!
  - Infeasible to check all these deviations

Nash equilibria with mixed strategies

- $\boxed{\text{Corollary:}}$ It suffices to check only deviations to pure strategies
  - Because each mixed strategy is a convex combination of pure strategies

- Equivalent definition: A profile of mixed strategies (**p**, **q**) is a Nash equilibrium if
  - $u_1(\mathbf{p}, \mathbf{q}) \geq u_1(e^i, \mathbf{q})$ for every pure strategy e^i of pl. 1
  - $u_2(\mathbf{p}, \mathbf{q}) \geq u_2(\mathbf{p}, e^j)$ for every pure strategy e^j of pl. 2

- Hence, we only need to check n+m inequalities as in the case of pure equilibria

$\left((1/3, 2/3), (4/5, 1/5) \right)$?

Mixed equilibria

- **Mixed equilibrium:** A profile of *mixed* strategies such that each player maximizes its expected utility, given the strategies of the other players

Theorem [Nash, 1951]

Every finite strategic game of n players has at least one mixed equilibrium

- Every pure equilibrium is also a mixed equilibrium
  - Every pure strategy can be seen as a probability distribution over all strategies that assigns probability 1 to this one pure strategy

Matching Pennies: mixed equilibria

		odd		
		heads	**tails**	
even	**heads**	**1, -1**	**-1, 1**	x
	tails	**-1, 1**	**1, -1**	$1-x$
		y	$1-y$	

- Even player selects heads with probability x and tails with $1 - x$
- Odd player selects heads with probability y and tails with $1 - y$

- $p(\text{heads, heads}) = xy$
- $p(\text{heads, tails}) = x(1 - y)$
- $p(\text{tails, heads}) = (1 - x)y$
- $p(\text{tails, tails}) = (1 - x)(1 - y)$

Matching Pennies: mixed equilibria

		odd		
		heads	**tails**	
even	**heads**	**1, -1**	**-1, 1**	x
	tails	**-1, 1**	**1, -1**	$1-x$
		y	$1-y$	

- $\mathbb{E}_p[u_\text{e}]$
  $= xy \cdot 1 + x(1-y) \cdot (-1) + (1-x)y \cdot (-1) + (1-x)(1-y) \cdot 1$
  $= 4xy - 2x - 2y + 1$
  $= \boldsymbol{x(4y-2) - 2y + 1}$

- $\mathbb{E}_p[u_\text{o}]$
  $= xy \cdot (-1) + x(1-y) \cdot 1 + (1-x)y \cdot 1 + (1-x)(1-y) \cdot (-1)$
  $= \boldsymbol{y(2-4x) + 2x - 1}$

Matching Pennies: mixed equilibria

- $\mathbb{E}_p[u_{\mathrm{e}}] = x(4y - 2) - 2y + 1$
- $\mathbb{E}_p[u_{\mathrm{o}}] = y(2 - 4x) + 2x - 1$

- The expected utility of each player is a **linear function** in terms of her corresponding probability
- To analyze how a player is going to act, we need to see whether the slope of the linear function is negative or positive

- **Negative:** the function is decreasing and the player aims to set a small value for the probability
- **Positive:** the function is increasing and the players aims to set a high value for the probability

Matching Pennies: mixed equilibria

- $\mathbb{E}_p[u_{\mathsf{e}}] = x(4y-2) - 2y + 1$
- $\mathbb{E}_p[u_{\mathsf{o}}] = y(2-4x) + 2x - 1$

- **Even player:** the slope is $4y-2$ and it depends on y, the probability with which the odd player selects heads
- $\boldsymbol{y < 1/2}$
  - ⇨ the slope $4y-2$ is **negative**
  - ⇨ the function $\mathbb{E}_p[u_{\mathsf{e}}]$ is **decreasing in $\boldsymbol{x}$**
  - ⇨ even player sets $\boldsymbol{x = 0}$ to maximize $\mathbb{E}_p[u_{\mathsf{e}}]$
  - ⇨ the slope $2-4x = 2$ of the odd player is **positive**
  - ⇨ the function $\mathbb{E}_p[u_{\mathsf{o}}]$ is **increasing in $\boldsymbol{y}$**
  - ⇨ odd player sets $\boldsymbol{y = 1}$ to maximize $\mathbb{E}_p[u_{\mathsf{o}}]$
  - ⇨ **contradiction**

Matching Pennies: mixed equilibria

- $\mathbb{E}_p[u_{\text{e}}] = x(4y - 2) - 2y + 1$
- $\mathbb{E}_p[u_{\text{o}}] = y(2 - 4x) + 2x - 1$

- **Even player:** the slope is $4y - 2$ and it depends on y, the probability with which the odd player selects heads
- $\boldsymbol{y > 1/2}$
  - ⇨ the slope $4y - 2$ is **positive**
  - ⇨ the function $\mathbb{E}_p[u_{\text{e}}]$ is **increasing in $\boldsymbol{x}$**
  - ⇨ even player sets $\boldsymbol{x = 1}$ to maximize $\mathbb{E}_p[u_{\text{e}}]$
  - ⇨ the slope $2 - 4x = -2$ of the odd player is **negative**
  - ⇨ the function $\mathbb{E}_p[u_{\text{o}}]$ is **decreasing in $\boldsymbol{y}$**
  - ⇨ odd player sets $\boldsymbol{y = 0}$ to maximize $\mathbb{E}_p[u_{\text{o}}]$
  - ⇨ **contradiction**

Matching Pennies: mixed equilibria

- $\mathbb{E}_p[u_{\mathrm{e}}] = x(4y - 2) - 2y + 1$
- $\mathbb{E}_p[u_{\mathrm{o}}] = y(2 - 4x) + 2x - 1$

- It must be $\boldsymbol{y = 1/2}$
- Following the same reasoning for the odd player, we can see that it must also be $\boldsymbol{x = 1/2}$
- For these values of x and y both slopes are equal to 0 and the linear functions are maximized
- The pair $(x, y) = (1/2, 1/2)$ corresponds to a mixed equilibrium, which is actually unique for this game

Multi-player games

Games with more than 2 players

- All the definitions we have seen can be generalized for multi-player games
  - Dominant strategies, Nash equilibria
- But: we can no longer have a representation with 2-dimensional arrays
- For n-player games we would need n-dimensional arrays (unless there is a more concise representation)

Definitions for n-player games

Definition: A game in normal form consists of

- A set of players $N = \{1, 2,\dots, n\}$
- For every player i, a set of available pure strategies S^i
- For every player i, a utility function
  $u_i: S^1 \times \dots \times S^n \to R$

- Let $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$ be a profile of mixed strategies for the players
- Each $\mathbf{p}_i$ is a probability distribution on S^i
- Expected utility of pl. i under $\mathbf{p}$ =

$$u_i(\mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{(s_1,\dots,s_n)\in S^1\times\cdots\times S^n} \mathbf{p}_1(s_1)\dots\mathbf{p}_n(s_n)\underline{u_1(s_1,\dots,s_n)}$$

$(s_1, s_2, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n) \to (s_1, \dots, s_{i-1}, s_i', s_{i+1}, \dots, s_n)$

Notation

(s_i, s_{-i}) (s_i', s_{-i})

- Given a vector $s = (s_1, \dots, s_n)$, we denote by s_{-i} the vector where we have removed the i-th coordinate:

$$s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$$

- E.g., if $s = (\cancel{3}, 5, 7, 8)$, then
  - $s_{-3} = (3, 5, 8)$
  - $s_{-1} = (5, 7, 8)$
- We can write a strategy profile s as $s = (s_i, s_{-i})$

Definitions for n-player games

- A strategy $\mathbf{p}_i$ of pl. i is *dominant* if

$$u_i(\mathbf{p}_i, \mathbf{p}_{-i}) \geq u_i(e^j, \mathbf{p}_{-i})$$

  for every pure strategy e^j of pl. i, and every profile $\mathbf{p}_{-i}$ of the other players
- Replace $\geq$ with $>$ for strictly dominant

- A profile $\mathbf{p} = (\mathbf{p}_1, \ldots, \mathbf{p}_n)$ is a Nash equilibrium if for every player i and every pure strategy e^j of pl. i, we have

$$u_i(\mathbf{p}) \geq u_i(e^j, \mathbf{p}_{-i})$$

  – As in 2-player games, it suffices to check only deviations to pure strategies

$u_i(p_i, p_{-i})$

Nash equilibria in multi-player games

At a first glance:

- Even finding pure Nash equilibria looks already more difficult than in the 2-player case
- We can try with brute force all possible profiles
- Suppose we have n players, and each of them has m strategies: $|S^i| = m$
- There are m^n pure strategy profiles!
- However, in some cases, we can exploit symmetry or other properties to reduce our search space

Example: Congestion games

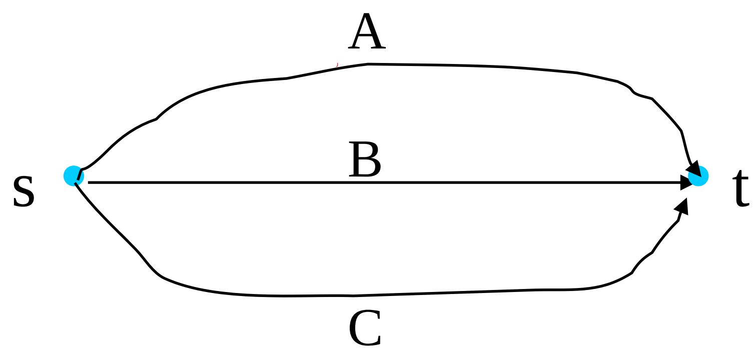

A simple example of a congestion game:

- A set of network users wants to move from s to t
- 3 possible routes, A, B, C
- Time delay in a route: depends on the number of users who have chosen this route
- $d_A(x) = 5x$, $d_B(x) = 7.5x$, $d_C(x) = 10x$,

Example: Congestion games

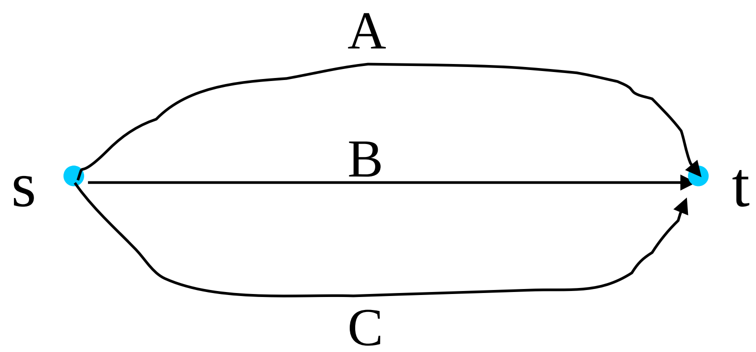

- Suppose we have n = 5 players
- For each player i, S^i = {A, B, C}
- Number of possible pure strategy profiles: 3^5 = 243
- Utility function of a player: should increase when delay decreases (e.g., we can define it as u = – delay)
- At profile s = (A, C, A, B, A)
  - $u_1(s) = -15$, $u_2(s) = -10$, $u_3(s) = -15$, $u_4(s) = -7.5$, $u_5(s) = -15$

Example: Congestion games

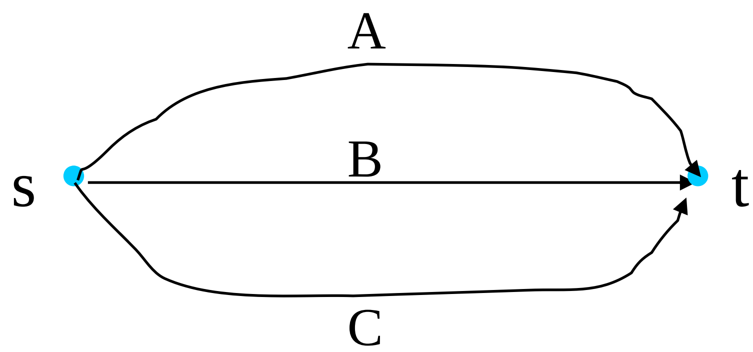

- There is no need to examine all 243 possible profiles to find a pure equilibrium
- Exploiting symmetry:
  - In every route, the delay does not depend on who chose the route but only how many did so
- We can also exploit further properties
  - E.g. There can be no equilibrium where one of the routes is not used by some player

Homework: Find the pure Nash equilibria of this game (if there are any)

Existence of Nash equilibria

Nash equilibria: Recap

Recall the problematic issues we have identified for pure Nash equilibria:

1. Non-existence: there exist games that do not possess an equilibrium with pure strategies
2. Non-uniqueness: there are games that have many Nash equilibria
3. Welfare guarantees: The equilibria of a game do not necessarily have the same utility for the players

Have we made any progress by considering equilibria with mixed strategies?

Existence of Nash equilibria

- Theorem [Nash 1951]: Every finite game possesses at least one equilibrium when we allow mixed strategies
  - Finite game: finite number of players, and finite number of pure strategies per player
- Corollary: if a game does not possess an equilibrium with pure strategies, then it definitely has one with mixed strategies
- One of the most important results in game theory
- Nash's theorem resolves the issue of non-existence
  - By allowing a richer strategy space, existence is guaranteed, no matter how big or complex the game might be

Examples

- In Prisoner's dilemma or Bach-or-Stravinsky, there exist equilibria with pure strategies
  - For such games, Nash's theorem does not add any more information. However, in addition to pure equilibria, we may also have some mixed equilibria

- Matching-Pennies: For this game, Nash's theorem guarantees that there exists an equilibrium with mixed strategies
  - In fact, it is the profile we saw: ((1/2, 1/2), (1/2, 1/2))

- Rock-Paper-Scissors?
  - Again the uniform distribution: ((1/3, 1/3, 1/3), (1/3, 1/3, 1/3))

Nash equilibria: Computation

- Nash's theorem only guarantees the existence of Nash equilibria
  - Proof reduces to using Brouwer's fixed point theorem
- Brouwer's theorem: Let $f{:}D \rightarrow D$, be a continuous function, and suppose D is convex and compact. Then there exists x such that $f(x) = x$
  - Many other versions of fixed point theorems also available

Nash equilibria: Computation

- So far, we are not aware of efficient algorithms for finding fixed points [Hirsch, Papadimitriou, Vavasis ’91]
  - There exist exponential time algorithms for finding approximate fixed points
- Can we design polynomial time algorithms for 2-player games?
  - After all, it seems to be only a special case of the general problem of finding fixed points
- For games with more players?