Networks' Challenge: Where Game Theory Meets Network Optimization

Asu Ozdaglar

International Symposium on Information Theory

July 2008

Department of Electrical Engineering & Computer Science

Massachusetts Institute of Technology, USA

Introduction

- **Central Question in Today's and Future Networks:**
  - Systematic analysis and design of network architectures and development of network control schemes
- **Traditional Network Optimization:** Single administrative domain with a single control objective and obedient users.
- **New Challenges:**
  - Large-scale with lack of access to centralized information and subject to unexpected disturbances
    - **Implication**: *Control policies have to be decentralized, scalable, and robust against dynamic changes*
  - Interconnection of heterogeneous autonomous entities, so no central party with enforcement power or accurate information about user needs
    - **Implication**: *Selfish incentives and private information of users need to be incorporated into the control paradigm*
  - Continuous upgrades and investments in new technologies
    - **Implication**: *Economic incentives of service and content providers much more paramount*

Tools and Ideas for Analysis

- These challenges necessitate the analysis of resource allocation and data processing in the presence of decentralized information and heterogeneous selfish users and administrative domains
- Instead of a central control objective, model as a multi-agent decision problem: Game theory and economic market mechanisms
- **Game Theory:** Understand incentives of selfish autonomous agents and large players such as service and content providers
  - Utility-based framework of economics (represent user preferences by utility functions)
  - Decentralized equilibrium of a multi-agent system (does not require tight closed-loop explicit controls)
  - *Mechanism Design Theory:* Inverse game theory
    - Design the system in a way that maintains decentralization, but provides appropriate incentives
- Large area of research at the intersection of Engineering, Computer Science, Economics, and Operations Research

Applications

- **Wireless Communications**
  - Power control in CDMA networks
  - Transmission scheduling in collision channels
  - Routing in multi-hop relay networks
  - Spectrum assignment in cognitive-radio networks
- **Data Networks**
  - Selfish (source) routing in overlay networks, inter-domain routing
  - Rate control using market-based mechanisms
  - Online advertising on the Internet: Sponsored search auctions
  - Network design and formation
  - Pricing and investment incentives of service providers
- **Other Networked-systems**
  - Social Networks: Information evolution, learning dynamics, herding
  - Transportation Networks, Electricity Markets

This Tutorial

- **Tools for Analysis – Part I**
  - Strategic and extensive form games
  - Solution concepts: Iterated strict dominance, Nash equilibrium, subgame perfect equilibrium
  - Existence and uniqueness results
- **Network Games – Part I**
  - Selfish routing and Price of Anarchy
  - Service provider effects:
    - Partially optimal routing
    - Pricing and capacity investments
- **Tools for Analysis – Part II**
  - Supermodular games and dynamics
  - Potential and congestion games
- **Network Games – Part II**
  - Distributed power control algorithms
  - Network design

Motivating Example

Selfish Routing for Noncooperative Users

- For simplicity, no utility from flow, just congestion effects (inelastic demand)
- Each link described by a convex latency function $l_i(x_i)$ measuring costs of delay and congestion on link i as a function of link flow x_i.

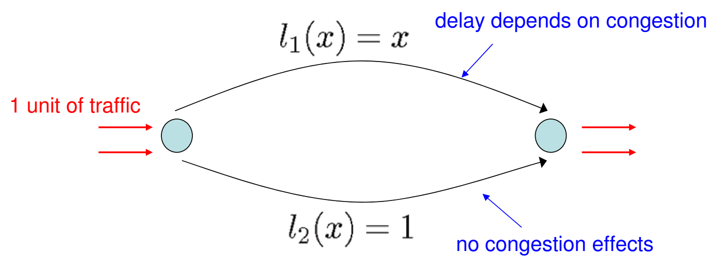

- **Traditional Network Optimization Approach:**
  - Centralized control, single metric: e.g. minimize total delay
- **Selfish Routing:**
  - Allow end users to choose routes themselves: e.g. minimize own delay
    - Applications: Transportation networks; Overlay networks
  - What is the right equilibrium notion?
    - **Nash Equilibrium:** Each user plays a "best-response" to actions of others
    - **Wardrop Equilibrium:** Nash equilibrium when "users infinitesimal"

Wardrop Equilibrium with Selfish Routing

- Consider the simple **Pigou example**:

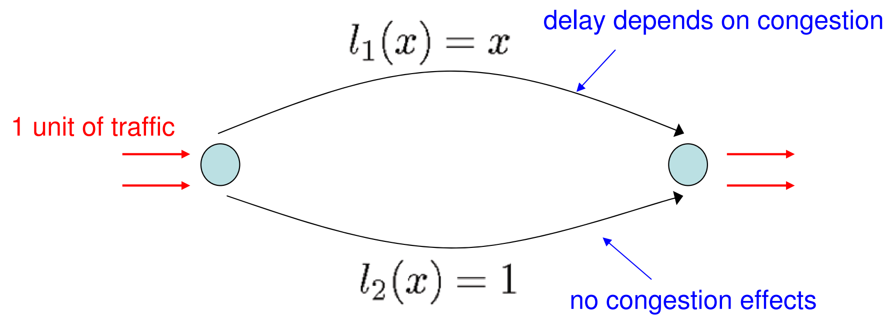

- In centralized optimum, traffic split equally between two links.
  - Cost of optimal flow: $C_{\text{system}}(x^S) = \sum_i l_i(x_i^S)x_i^S = \frac{1}{4} + \frac{1}{2} = \frac{3}{4}$
- In Wardrop equilibrium, cost equalized on paths with positive flow; all traffic goes through top link.
  - Cost of selfish routing: $C_{\text{eq}}(x^{WE}) = \sum_i l_i(x_i^{WE})x_i^{WE} = 1 + 0 = 1$
- **Efficiency metric:** Given latency functions $\{l_i\}$, we define the efficiency metric

$$\alpha = \frac{C_{\text{system}}(x^S)}{C_{\text{eq}}(x^{WE})}$$

- For the above example, we have $\alpha = \frac{3}{4}$.

Game Theory Primer–I

- A strategic (form) game is a model for a game in which all of the participants act simultaneously and without knowledge of other players' actions.

Definition (Strategic Game): A *strategic game* is a triplet $\langle \mathcal{I}, (S_i)_{i\in\mathcal{I}}, (u_i)_{i\in\mathcal{I}}\rangle$:

- $\mathcal{I}$ is a finite set of players, $\mathcal{I} = \{1, \ldots, I\}$.
- S_i is the set of available actions for player i
  - $s_i \in S_i$ is an action for player i
  - $s_{-i} = [s_j]_{j\neq i}$ is a vector of actions for all players *except* i.
  - $(s_i, s_{-i}) \in S$ is an *action profile*, or *outcome*.
  - $S = \prod_i S_i$ is the set of all action profiles
  - $S_{-i} = \prod_{j\neq i} S_j$ is the set of all action profiles for all players *except* i
- $u_i : S \to \mathbb{R}$ is the payoff (utility) function of player i

- For strategic games, we will use the terms **action** and **pure strategy** interchangeably.

Example–Finite Strategy Spaces

- When the strategy space is finite, and the number of players and actions is small, a game can be represented in matrix form.
- The cell indexed by row x and column y contains a pair, (a, b) where $a = u_1(x, y)$ and $b = u_2(x, y)$.
- **Example:** Matching Pennies.

	Heads	Tails
Heads	$-1, 1$	$1, -1$
Tails	$1, -1$	$-1, 1$

- This game represents pure conflict in the sense that one player's utility is the negative of the utility of the other player.
  - Zero-sum games: favorable structure for dynamics and computation of equilibria

Example–Infinite Strategy Spaces

- **Example:** Cournot competition.
  - Two firms producing the same good.
  - The action of a player i is a quantity, $s_i \in [0, \infty]$ (amount of good he produces).
  - The utility for each player is its total revenue minus its total cost,

$$u_i(s_1, s_2) = s_i p(s_1 + s_2) - cs_i$$

    where $p(q)$ is the price of the good (as a function of the total amount), and c is unit cost (same for both firms).

- Assume for simplicity that $c = 1$ and $p(q) = \max\{0, 2 - q\}$
- Consider the best-response correspondences for each of the firms, i.e., for each i, the mapping $B_i(s_{-i}) : S_{-i} \rightarrow S_i$ such that

$$B_i(s_{-i}) \in \text{argmax}_{s_i \in S_i} u_i(s_i, s_{-i}).$$

Example–Continued

- By using the first order optimality conditions, we have

$$
\begin{aligned}
B_i(s_{-i}) &= \text{argmax}_{s_i \geq 0}(s_i(2 - s_i - s_{-i}) - s_i) \\
&= \begin{cases} \frac{1-s_{-i}}{2} & \text{if } s_{-i} \leq 1, \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}
$$

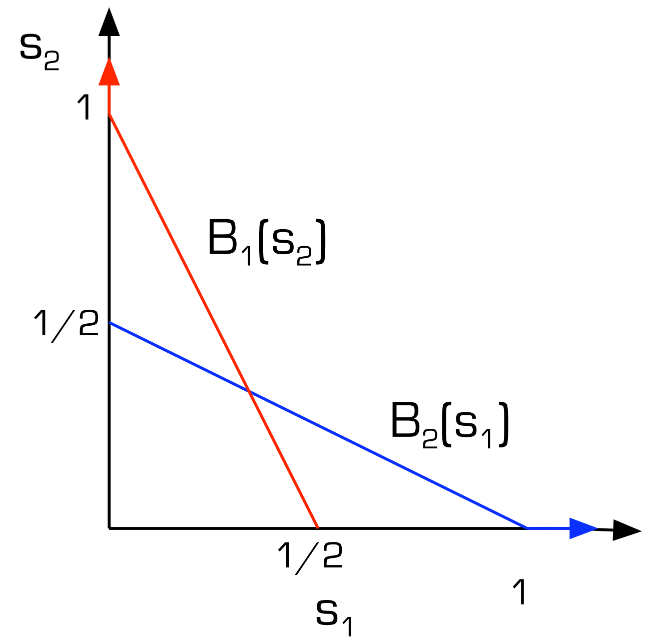

- The figure illustrates the best response functions as a function of s_1 and s_2.

- Assuming that players are rational and fully knowledgable about the structure of the game and each other's rationality, what should the outcome of the game be?

Dominant Strategies

- **Example:** Prisoner's Dilemma.
  - Two people arrested for a crime, placed in separate rooms, and the authorities are trying to extract a confession.

	Cooperate	Don't Cooperate
Cooperate	$2, 2$	$5, 1$
Don't Cooperate	$1, 5$	$4, 4$

- What will the outcome of this game be?
  - Regardless of what the other player does, playing "DC" is better for each player.
  - The action "DC" **strictly dominates** the action "C"
- Prisoner's dilemma paradigmatic example of a self-interested, rational behavior not leading to jointly (socially) optimal result.

Prisoner's Dilemma and ISP Routing Game

- Consider two Internet service providers that need to send traffic to each other
- Assume that the unit cost along a link (edge) is 1

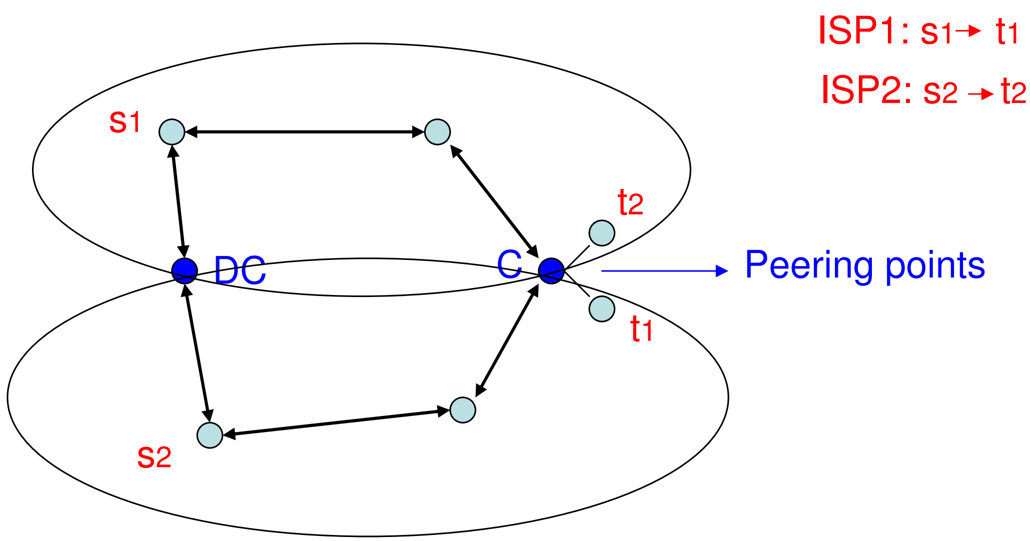

- This situation can be modeled by the "Prisoner's Dilemma" payoff matrix

	C	DC
C	2, 2	5, 1
DC	1, 5	4, 4

Dominated Strategies

Definition (Strictly Dominated Strategy): A strategy $s_i \in S_i$ is *strictly dominated* for player i if there exists some $s_i' \in S_i$ such that

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \qquad \text{for all } s_{-i} \in S_{-i}.$$

Definition (Weakly Dominated Strategy): A strategy $s_i \in S_i$ is *weakly dominated* for player i if there exists some $s_i' \in S_i$ such that

$$u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) \qquad \text{for all } s_{-i} \in S_{-i},$$

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \qquad \text{for some } s_{-i} \in S_{-i}.$$

- No player would play a strictly dominated strategy
- Common knowledge of payoffs and rationality results in **iterative elimination of strictly dominated strategies**

Example: Iterated Elimination of Strictly Dominated Strategies.

	Left	Middle	Right
Up	$4,3$	$5,1$	$6,2$
Middle	$2,1$	$8,4$	$3,6$
Down	$3,0$	$9,6$	$2,8$

Revisiting Cournot Competition

- Apply iterated strict dominance to Cournot model to predict the outcome

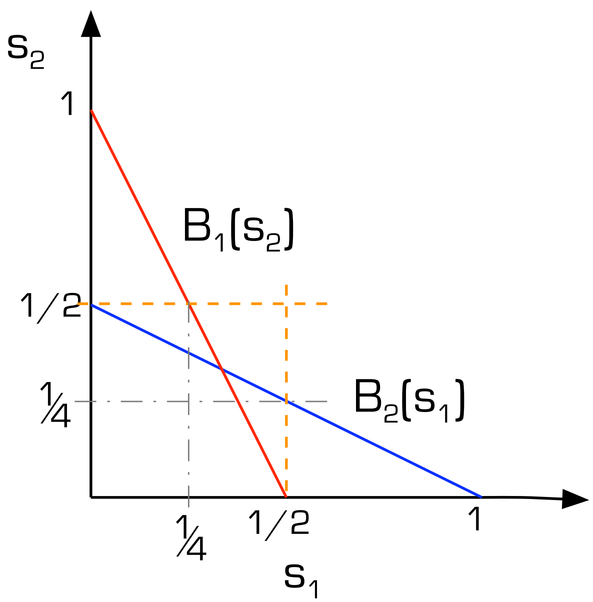

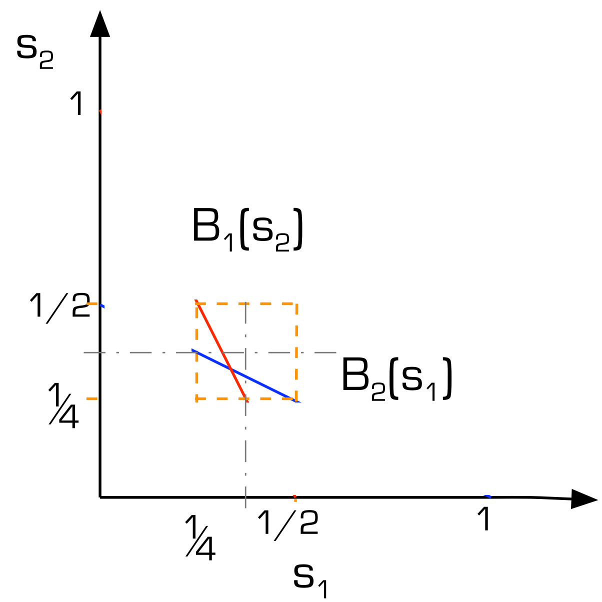

- One round of elimination yields $S_1^1 = [0, 1/2]$, $S_2^1 = [0, 1/2]$
- Second round of elimination yields $S_1^1 = [1/4, 1/2]$, $S_2^1 = [1/4, 1/2]$
- It can be shown that the endpoints of the intervals converge to the intersection
- Most games not solvable by iterated strict dominance, need a stronger **equilibrium notion**

Pure Strategy Nash Equilibrium

Definition (Nash equilibrium): A (pure strategy) Nash Equilibrium of a strategic game $\langle \mathcal{I}, (S_i)_{i\in\mathcal{I}}, (u_i)_{i\in\mathcal{I}}\rangle$ is a strategy profile $s^* \in S$ such that for all $i \in \mathcal{I}$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \qquad \text{for all } s_i \in S_i.$$

- No player can profitably deviate given the strategies of the other players
- Why should one expect Nash equilibrium to arise?
  - Introspection
  - Self-enforcing
  - Learning or evolution
- Recall the best-response correspondence $B_i(s_{-i})$ of player i,

$$B_i(s_{-i}) \in \arg\max_{s_i\in S_i} u_i(s_i, s_{-i}).$$

- An action profile s^* is a Nash equilibrium if and only if

$$s_i^* \in B_i(s_{-i}^*) \qquad \text{for all } i \in \mathcal{I}.$$

- **Question:** When iterated strict dominance yields a unique strategy profile, is this a Nash equilibrium?

Examples

Example: Battle of the Sexes (players wish to coordinate but have conflicting interests)

	BALLET	SOCCER
BALLET	$2, 1$	$0, 0$
SOCCER	$0, 0$	$1, 2$

- Two Nash equilibria, (Ballet, Ballet) and (Soccer, Soccer).

Example: Matching Pennies.

	HEADS	TAILS
HEADS	$1, -1$	$-1, 1$
TAILS	$-1, 1$	$1, -1$

Matching Pennies

- No pure Nash equilibrium
- There exists a "stochastic steady state", in which each player chooses each of her actions with $1/2$ probability $\Rightarrow$ **Mixed strategies**

Mixed Strategies and Mixed Strategy Nash Equilibrium

- Let Σ_i denote the set of probability measures over the pure strategy (action) set S_i.
- We use $\sigma_i \in \Sigma_i$ to denote the **mixed strategy of player i**, and $\sigma \in \Sigma = \prod_{i \in \mathcal{I}} \Sigma_i$ to denote a **mixed strategy profile**.
- Note that this implicitly assumes that **players randomize independently**.
- We similarly define $\sigma_{-i} \in \Sigma_{-i} = \prod_{j \neq i} \Sigma_j$.
- Following Von Neumann-Morgenstern expected utility theory, we extend the payoff functions u_i from S to Σ by

$$u_i(\sigma) = \int_S u_i(s) d\sigma(s).$$

Definition (Mixed Nash Equilibrium): A mixed strategy profile σ^* is a (mixed strategy) Nash Equilibrium if for each player i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

- Note that it is sufficient to check pure strategy deviations, i.e., σ^* is a mixed Nash equilibrium if and only if for all i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*) \qquad \text{for all } s_i \in S_i.$$

Characterization of Mixed Strategy Nash Equilibria

Lemma: Let $G = \langle \mathcal{I}, (S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$ be a finite strategic game. Then, $\sigma^* \in \Sigma$ is a Nash equilibrium if and only if for each player $i \in \mathcal{I}$, every pure strategy in the support of σ_i^* is a best response to σ_{-i}^*.

- It follows that every action in the support of any player's equilibrium mixed strategy yields the same payoff.
- The characterization result extends to **infinite games**: $\sigma^* \in \Sigma$ is a Nash equilibrium if and only if for each player $i \in \mathcal{I}$,
  - (i) no action in S_i yields, given σ_{-i}^*, a payoff that exceeds his equilibrium payoff,
  - (ii) the set of actions that yields, given σ_{-i}^*, a payoff less than his equilibrium payoff has σ_i^*-measure zero.
- **Example:** Recall Battle of the Sexes Game.

	Ballet	Soccer
Ballet	$2, 1$	$0, 0$
Soccer	$0, 0$	$1, 2$

This game has two pure Nash equilibria and a mixed Nash equilibrium $\left((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})\right)$.

Existence of Nash Equilibria –I

Theorem [Nash 50]: Every finite game has a mixed strategy Nash equilibrium.

Proof Outline:

- σ^* mixed Nash equilibrium if and only if $\sigma_i^* \in B_i(\sigma_{-i}^*)$ for all $i \in \mathcal{I}$, where

$$B_i(\sigma_{-i}^*) \in \arg \max_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_{-i}^*).$$

- This can be written compactly as $\sigma^* \in B(\sigma^*)$, where $B(\sigma) = [B_i(\sigma_{-i})]_{i \in \mathcal{I}}$, i.e., σ^* is a **fixed point of the best-response correspondence**.

- Use Kakutani's fixed point theorem to establish the existence of a fixed point.

Linearity of expectation in probabilities play a key role; extends to (quasi)-concave payoffs in infinite games

Theorem [Debreu, Glicksberg, Fan 52]: Assume that the S_i are nonempty compact convex subsets of an Euclidean space. Assume that the payoff functions $u_i(s_i, s_{-i})$ are quasi-concave in s_i and continuous in s, then there exists a pure strategy Nash equilibrium.

Existence of Nash Equilibria –II

- Can we relax quasi-concavity?
- **Example:** Consider the game where two players pick a location $s_1, s_2 \in \mathbb{R}^2$ on the circle. The payoffs are $u_1(s_1, s_2) = -u_2(s_1, s_2) = d(s_1, s_2)$, where $d(s_1, s_2)$ denotes the Euclidean distance between $s_1, s_2 \in \mathbb{R}^2$.
  - No pure Nash equilibrium.
  - The profile where both mix uniformly on the circle is a mixed Nash equilibrium.

Theorem [Glicksberg 52]: Every continuous game has a mixed strategy Nash equilibrium.

- Existence results for discontinuous games! [Dasgupta and Maskin 86]
- Particularly relevant for price competition models.

Uniqueness of Pure Nash Equilibrium in Infinite Games

- Concavity of payoffs $u_i(s_i, s_{-i})$ in s_i not sufficient to establish uniqueness
- Assume that $S_i \subset \mathbb{R}^{m_i}$. We use the notation:

$$\nabla_i u(x) = \left[\frac{\partial u(x)}{\partial x_i^1}, \ldots, \frac{\partial u(x)}{\partial x_i^{m_i}}\right]^T, \qquad \nabla u(x) = \left[\nabla_1 u_1(x), \ldots, \nabla_I u_I(x)\right]^T.$$

Definition: We say that the payoff functions $(u_1, \ldots, u_I)$ are *diagonally strictly concave for* $x \in S$, if for every $x^*, \bar{x} \in S$, we have

$$(\bar{x} - x^*)^T \nabla u(x^*) + (x^* - \bar{x})^T \nabla u(\bar{x}) > 0.$$

- Let $U(x)$ denote the Jacobian of $\nabla u(x)$, i.e., for $m_i = 1$, $[U(x)]_{ij} = \frac{\partial^2 u_i(x)}{\partial x_j x_i}$
- A sufficient condition for diagonal strict concavity is that the symmetric matrix $(U(x) + (U^T(x)))$ is negative definite for all $x \in S$.

Theorem [Rosen 65]: Assume that the payoff functions $(u_1, \ldots, u_I)$ are diagonally strictly concave for $x \in S$. Then the game has a unique pure strategy Nash equilibrium.

Extensive Form Games

- Extensive-form games model multi-agent sequential decision making.
- Our focus is on multi-stage games with observed actions

- Extensive form represented by tree diagrams
- Additional component of the model, histories (i.e., sequences of action profiles)
- Let H^k denote the set of all possible stage-k histories
- Strategies are maps from all possible histories into actions: $s_i^k : H^k \rightarrow S_i$

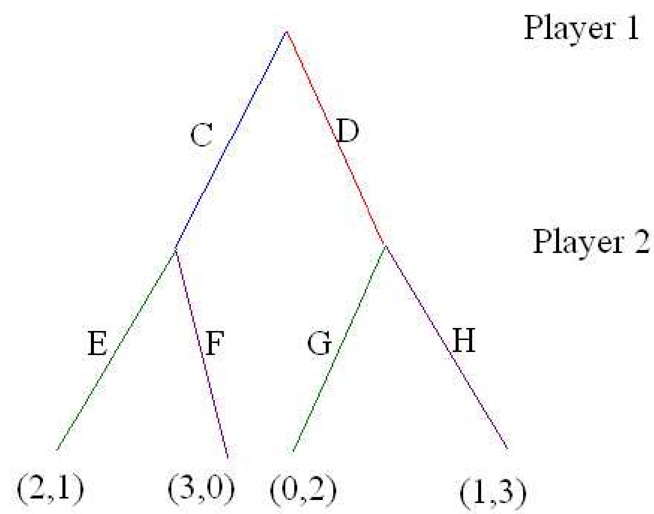

Example:

- Player 1's strategies: $s_1 : H^0 = \emptyset \rightarrow S_1$; two possible strategies: C,D
- Player 2's strategies: $s^2 : H^1 = \{C, D\} \rightarrow S_2$; four possible strategies: EG,EH,FG, FH

Subgame Perfect Equilibrium

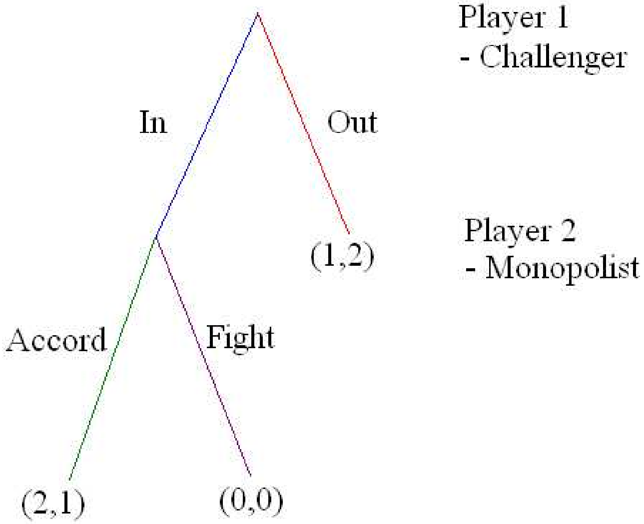

- Equivalent strategic form representation

	Accommodate	Fight
In	2,1	0,0
Out	1,2	1,2

- Two pure Nash equilibria: (In,A) and (Out,F)
- The equilibrium (Out,F) is sustained by a **non-credible threat** of the monopolist

- Equilibrium notion for extensive form games: **Subgame Perfect (Nash) Equilibrium**
  - Requires each player's strategy to be "optimal" not only at the start of the game, but also after every history
  - For finite horizon games, found by backward induction
  - For infinite horizon games, characterization in terms of **one-stage deviation principle**

Revisit Routing Models

- Directed network $N = (V, E)$
- Origin-destination pairs (s_j, t_j), $j = 1, \ldots, k$ with rate r_j
- $\mathcal{P}_j$ denotes the set of paths between s_j and t_j; $\mathcal{P} = \cup_j \mathcal{P}_j$
- x_p denotes the flow on path $p \in \mathcal{P}$ (can be non-integral)
- Each link $i \in E$ has a **latency function $l_i(x_i)$**, which captures congestion effects ($x_i = \sum_{\{p \in \mathcal{P} | i \in p\}} x_p$)
  - Assume $l_i(x_i)$ nonnegative, differentiable, and nondecreasing

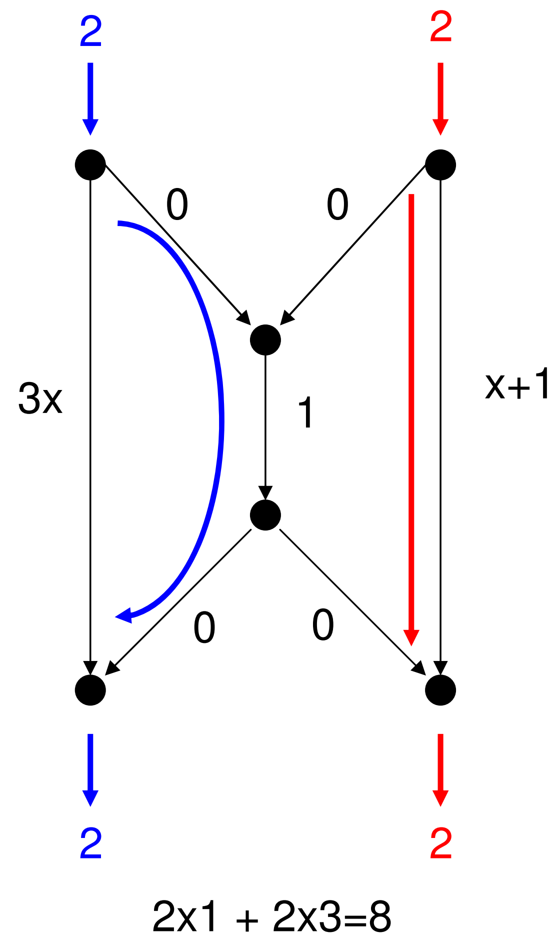

- We call the tuple $R = (V, E, (s_j, t_j, r_j)_{j=1,\ldots,k}, (l_i)_{i \in E})$ a **routing instance**
- The total latency cost of a flow x is: $C(x) = \sum_{i \in E} x_i l_i(x_i)$

Socially Optimal Routing

Given a routing instance $R = (V, E, (s_j, t_j, r_j), (l_i))$:

- We define the **social optimum** x^S, as the optimal solution of the multicommodity min-cost flow problem

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in E} x_i l_i(x_i) \\ \text{subject to} \quad & \sum_{\{p \in \mathcal{P} | i \in p\}} x_p = x_i, \quad i \in E, \\ & \sum_{p \in \mathcal{P}_j} x_p = r_j, \ j = 1, \ldots, k, \quad x_p \geq 0, \quad p \in \mathcal{P}. \end{aligned}$$

- We refer to a feasible solution of this problem as a **feasible flow**.

Wardrop (User) Equilibrium

- When traffic routes "selfishly," all nonzero flow paths must have equal latency.
  - **Nonatomic users** $\Rightarrow$ Aggregate flow of many "small" users.

Definition: A feasible flow is a **Wardrop equilibrium** x^{WE} if

$$\sum_{i \in p_1} l_i(x_i) \leq \sum_{i \in p_2} l_i(x_i), \qquad \text{for all } p_1, p_2 \in \mathcal{P} \text{ with } x_{p_1} > 0.$$

- A feasible flow is a Wardrop equilibrium x^{WE} iff it is an optimal solution of

$$\begin{aligned} \text{minimize} \quad & \sum_{i \in E} \int_0^{x_i} l_i(z)\, dz \\ \text{subject to} \quad & \sum_{\{p \in \mathcal{P} \mid i \in p\}} x_p = x_i, \quad i \in E, \\ & \sum_{p \in \mathcal{P}_j} x_p = r_j, \; j = 1, \ldots, k, \quad x_p \geq 0, \; p \in \mathcal{P}. \end{aligned}$$

- Existence and "essential" uniqueness of a Wardrop equilibrium follows from the previous optimization formulation [Beckmann, McGuire, Winsten 56]
- A feasible flow x^{WE} is a Wardrop equilibrium iff [Smith 79]

$$\sum_{i \in E} l_i(x_i^{WE})(x_i^{WE} - x_i) \leq 0, \qquad \text{for all feasible } x.$$

Recall Pigou Example

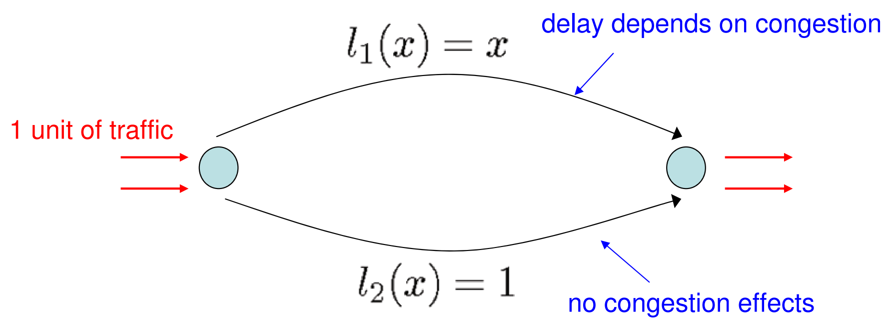

- In social optimum, traffic split equally between two links.
  - Cost of optimal flow: $C(x^S) = \sum_i l_i(x_i^S)x_i^S = \frac{1}{4} + \frac{1}{2} = \frac{3}{4}$
- In Wardrop equilibrium, cost equalized on paths with positive flow; all traffic goes through top link.
  - Cost of selfish routing: $C(x^{WE}) = \sum_i l_i(x_i^{WE})x_i^{WE} = 1 + 0 = 1$
- **Efficiency metric:** Given the routing instance R, we define the efficiency metric

$$\alpha(R) = \frac{C(x^S(R))}{C(x^{WE}(R))}$$

- For the above example, we have $\alpha(R) = \frac{3}{4}$.

Selfish Routing and Price of Anarchy

- Let $\mathcal{R}'$ denote the set of all routing instances.
- Worst case efficiency over all instances: $\inf_{R \in \mathcal{R}'} \frac{C(x^S(R))}{C(x^{WE}(R))}$
  - **Price of Anarchy**: Measure of lack of centralized coordination [Koutsoupias, Papadimitriou 99]

Theorem : [Roughgarden, Tardos 02] Let $\mathcal{R}^{aff}$ ($\mathcal{R}^{conv}$) denote routing instances with affine (convex) latency functions.

(a) Let $R \in \mathcal{R}^{aff}$. Then,

$$\frac{C(x^S(R))}{C(x^{WE}(R))} \geq \frac{3}{4}.$$

Furthermore, the bound above is tight.

(b)

$$\inf_{R \in \mathcal{R}^{conv}} \frac{C(x^S(R))}{C(x^{WE}(R))} = 0.$$

- Bounds for capacitated networks and polynomial latency functions [Correa, Schulz, Stier-Moses 03, 05]
- Genericity analysis [Friedman 04], [Qiu et al. 03]
  - Likely outcomes rather than worst cases

Further Paradoxes of Decentralized Equilibrium: Braess' Paradox

- **Idea:** Addition of an intuitively helpful link negatively impacts network users

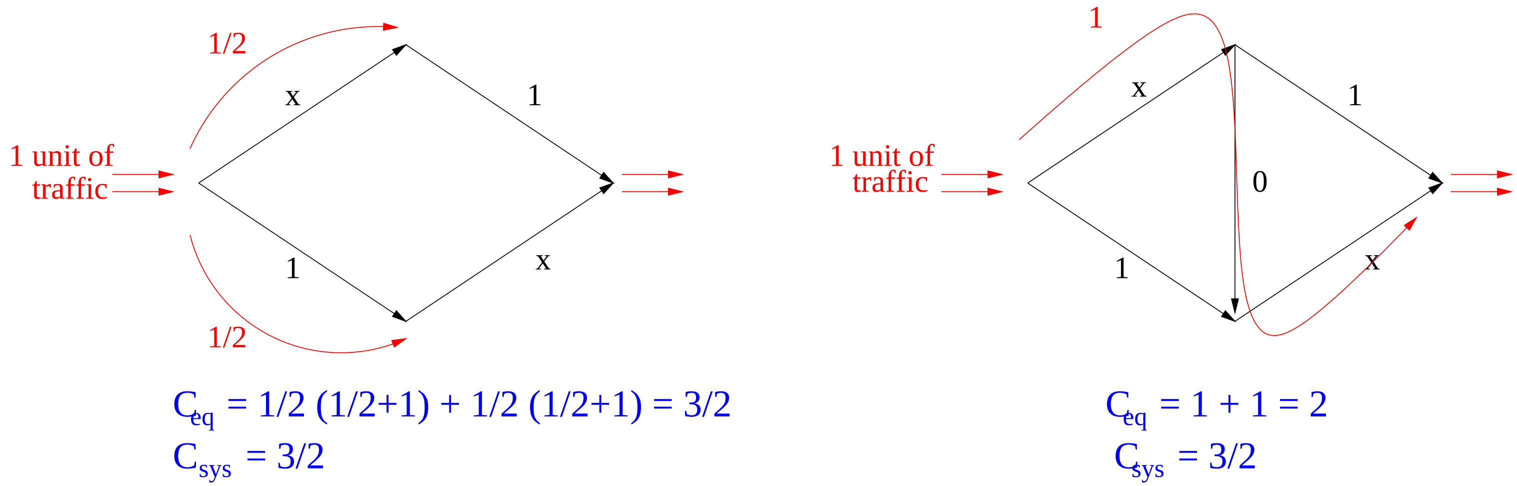

- Introduced in transportation networks [Braess 68], [Dafermos, Nagurney 84]
  - Studied in the context of communication networks, distributed computing, queueing networks [Altman et al. 03]
- Motivated research in methods of upgrading networks without degrading network performance
  - Leads to limited methods under various assumptions.

Selfish Routing

- Is this the right framework for thinking about network routing?
- **No, for 2 reasons:**
  - It ignores providers' role in routing traffic
  - It ignores providers' pricing and profit incentives

New Routing Paradigm for Noncooperative Users and Providers

- Most large-scale networks, such as Internet, consist of interconnected administrative domains that control traffic within their networks.
- Obvious conflicting interests as a result:
  - Users care about end-to-end performance.
  - Individual network providers optimize their own objectives.
- The study of routing patterns and performance requires an analysis of **Partially Optimal Routing (POR)**: [Acemoglu, Johari, Ozdaglar 06]
  - End-to-end route selection selfish
    - Transmission follows minimum latency route for each source.
  - Network providers route traffic within their own network to achieve minimum intradomain latency.

Partially Optimal Routing

- Consider a subnetwork inside of N, denoted $N_0 = (V_0, E_0)$.
- Assume first that N_0 has a **unique entry and exit point**, denoted by $s_0 \in V_0$ and $t_0 \in V_0$. $\mathcal{P}_0$ denotes paths from s_0 to t_0.
- We call $R_0 = (V_0, E_0, s_0, t_0)$ a **subnetwork** of N : $R_0 \subset R$.
- Given an incoming amount of flow X_0, the network operator chooses the routing by:

$$
\begin{aligned}
L(X_0) = \text{minimize} \quad & \sum_{i \in E_0} x_i l_i(x_i) \\
\text{subject to} \quad & \sum_{\{p \in P_0 | i \in p\}} x_p = x_i, \quad i \in E_0, \\
& \sum_{p \in \mathcal{P}_0} x_p = X_0, \qquad x_p \geq 0, \quad p \in \mathcal{P}_0.
\end{aligned}
$$

- Define $l_0(X_0) = L(X_0)/X_0$ as the **effective latency** of POR in the subnetwork R_0.

POR Flows

- Given a routing instance $R = (V, E, (s_j, t_j, r_j), (l_i))$, and a subnetwork $R_0 = (V_0, E_0, s_0, t_0)$ defined as above, we define a new routing instance $R' = (V', E', (s_j, t_j, r_j), (l'_i))$ as follows:

$$V' = (V \setminus V_0) \bigcup \{s_0, t_0\};$$

$$E' = (E \setminus E_0) \bigcup \{(s_0, t_0)\};$$

- $(l'_i) = \{l_i\}_{i \in E \setminus E_0} \bigcup \{l_0\}$.
- We refer to R' as the equivalent POR instance for R with respect to R_0.
- The overall network flow in R with partially optimal routing in R_0, $x^{POR}(R, R_0)$, is defined as:

$$x^{POR}(R, R_0) = x^{WE}(R').$$

Price of Anarchy for Partially Optimal Routing

- Let $\mathcal{R}^{aff}(\mathcal{R}^{conv})$ denote routing instances with affine (convex) latency functions.

Proposition: Let $\mathcal{R}'$ denote set of all routing instances.

$$\inf_{\substack{R\in\mathcal{R}' \\ R_0\subset R}} \frac{C(x^S(R))}{C(x^{POR}(R,R_0))} \leq \inf_{R\in\mathcal{R}'} \frac{C(x^S(R))}{C(x^{WE}(R))}.$$

$$\inf_{\substack{R\in\mathcal{R}^{aff} \\ R_0\subset R}} \frac{C(x^S(R))}{C(x^{POR}(R,R_0))} \geq \inf_{R\in\mathcal{R}^{conc}} \frac{C(x^S(R))}{C(x^{WE}(R))}.$$

Theorem:

(a)

$$\inf_{\substack{R\in\mathcal{R}^{conv} \\ R_0\subset R}} \frac{C(x^S(R))}{C(x^{POR}(R,R_0))} = 0.$$

(b) Consider a routing instance R where l_i is an affine latency function for all $i \in E$; and a subnetwork R_0 of R.

$$\frac{C(x^S(R))}{C(x^{POR}(R,R_0))} \geq \frac{3}{4}.$$

Furthermore, the bound above is tight.

Price of Anarchy for Partially Optimal Routing

Proof of part (b): The proof relies on the following two results:

Lemma: Assume that the latency functions l_i of all the links in the subnetwork are nonnegative affine functions. Then, the effective latency of POR, $l_0(X_0)$, is a nonnegative concave function of X_0.

Proposition: Let $R \in \mathcal{R}^{conc}$ be a routing instance where all latency functions are concave.

$$\frac{C(x^S(R))}{C(x^{WE}(R))} \geq \frac{3}{4}.$$

Furthermore, this bound is tight.

Price of Anarchy for Partially Optimal Routing

Proof of Proposition: From the variational inequality representation of WE, for all feasible x, we have

$$C(x^{WE}) = \sum_{j \in E} x_j^{WE} l_j(x_j^{WE}) \leq \sum_{j \in E} x_j l_j(x_j^{WE})$$
$$= \sum_{j \in E} x_j l_j(x_j) + \sum_{j \in E} x_j (l_j(x_j^{WE}) - l_j(x_j)).$$

For all feasible x, we have

$$x_j(l_j(x_j^{WE}) - l_j(x_j)) \leq \frac{1}{4} x_j^{WE} l_j(x_j^{WE}).$$

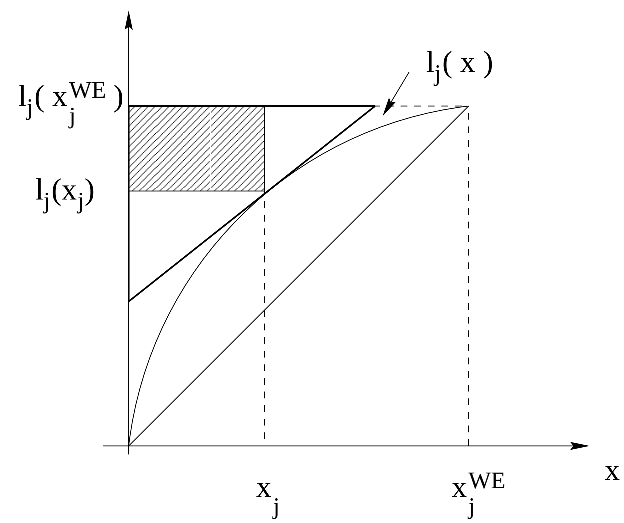

- For subnetworks with multiple entry-exit points, even for linear latencies, efficiency loss of POR can be arbitrarily high.
- Need for regulation and pricing!

Congestion and Provider Price Competition

- If a network planner can charge appropriate prices (taxes), system optimal solution can be decentralized even with selfish routing.
- Where do **prices** come from?
  - In newly-emerging large-scale networks, for-profit entities charge prices
  - Efficiency implications of profit-maximizing prices

Model:

- I parallel links.
- Interested in routing d units of traffic $\Rightarrow$ inelastic traffic

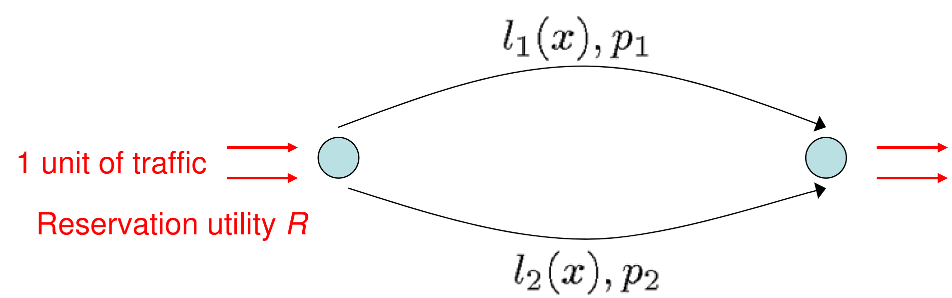

- Users have a reservation utility R and do not send their flow if the effective cost exceeds the reservation utility.
- Each link owned by a different service provider: charges a price p_i per unit bandwidth on link i (extends to arbitrary market structure).

Wardrop Equilibrium

- Assume $l_i(x_i)$: convex, continuously differentiable, nondecreasing.

Wardrop's principle: Flows routed along paths with minimum *"effective cost"*.

Definition: Given $p \geq 0$, x^* is a *Wardrop Equilibrium* (WE) if

$$\begin{aligned} l_i(x_i^*) + p_i &= \min_j \{l_j(x_j^*) + p_j\}, \qquad \text{for all } i \text{ with } x_i^* > 0, \\ l_i(x_i^*) + p_i &\leq R, \qquad \text{for all } i \text{ with } x_i^* > 0, \end{aligned}$$

and $\sum_{i \in \mathcal{I}} x_i^* \leq d$, with $\sum_{i \in \mathcal{I}} x_i^* = d$ if $\min_j \{l_j(x_j) + p_j\} < R$.

We denote the set of WE at a given p by $W(p)$.

- For any $p \geq 0$, the set $W(p)$ is nonempty.
- If l_i strictly increasing, $W(p)$ is a singleton and a continuous function of p.

Social Problem and Optimum

Definition: A flow vector x^S is a *social optimum* if it is an optimal solution of the *social problem*

$$\text{maximize}_{\substack{x \geq 0 \\ \sum_{i \in \mathcal{I}} x_i \leq d}} \sum_{i \in \mathcal{I}} (R - l_i(x_i)) x_i,$$

- It follows from the Karush-Kuhn-Tucker optimality conditions that $x^S \in \mathbb{R}_+^I$ is a social optimum iff

$$\begin{aligned} l_i(x_i^S) + x_i^S l_i'(x_i^S) &= \min_{j \in \mathcal{I}} \{ l_j(x_j^S) + x_j^S l_j'(x_j^S) \}, \qquad \forall\, i \text{ with } x_i^S > 0, \\ l_i(x_i^S) + x_i^S l_i'(x_i^S) &\leq R, \qquad \forall\, i \text{ with } x_i^S > 0, \end{aligned}$$

$\sum_{i \in \mathcal{I}} x_i^S \leq d$, with $\sum_{i \in \mathcal{I}} x_i^S = d$ if $\min_j \{ l_j(x_j^S) + x_j^S l_j'(x_j^S) \} < R$.

- $(l_i)'(x_i^S) x_i^S$**: Marginal congestion cost, Pigovian tax.**

Oligopoly Equilibrium

- Given the prices of other providers $p_{-i} = [p^j]_{j \neq i}$, SP i sets p_i to maximize his profit

$$\Pi_i(p_i, p_{-i}, x) = p_i x_i,$$

where $x \in W(p_i, p_{-i})$.

- We refer to the game among SPs as the *price competition game*.

Definition: A vector $(p^{OE}, x^{OE}) \geq 0$ is a (pure strategy) Oligopoly Equilibrium (OE) if $x^{OE} \in W(p_i^{OE}, p_{-i}^{OE})$ and for all $i \in \mathcal{I}$,

$$\Pi_i(p_i^{OE}, p_{-i}^{OE}, x^{OE}) \geq \Pi_i(p_i, p_{-i}^{OE}, x), \qquad \forall\ p_i \geq 0,\ \forall\ x \in W(p_i, p_{-i}^{OE}). \qquad (1)$$

We refer to p^{OE} as the *OE price*.

- Equivalent to the subgame perfect equilibrium notion.

Example

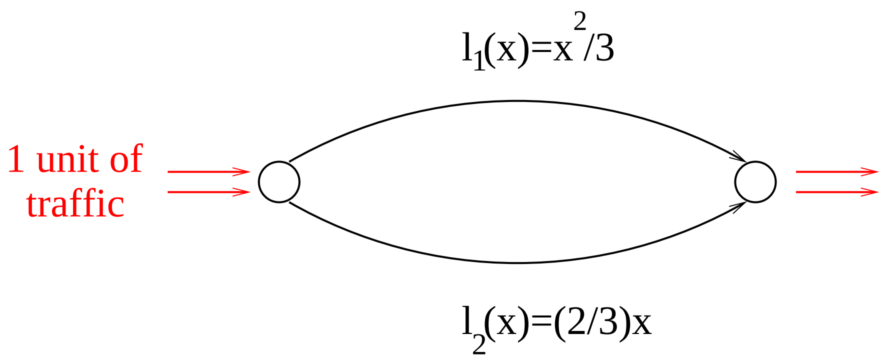

- Social Optimum: $x_1^S = 2/3,\qquad x_2^S = 1/3$
- WE: $x_1^{WE} = 0.73 > x_1^S,\qquad x_2^{WE} = 0.27$
- Single Provider: $x_1^{ME} = 2/3,\qquad x_2^{ME} = 1/3$
- Multiple Providers: $x_1^{OE} = 0.58,\qquad x_2^{OE} = 0.42$
  - The monopolist internalizes the congestion externalities.
  - Increasing competition decreases efficiency!
  - There is an additional source of "differential power" in the oligopoly case that distorts the flow pattern.

Existence and Price Characterization

Proposition: Assume that the latency functions are linear. Then the price competition game has a (pure strategy) OE.

- Existence of a mixed strategy equilibrium can be established for arbitrary convex latency functions.
- Oligopoly Prices: Let (p^{OE}, x^{OE}) be an OE. Then,

$$p_i^{OE} = (l_i)'(x_i^{OE})x_i^{OE} + \frac{\sum_{j \in \mathcal{I}_s} x_j^{OE}}{\sum_{j \notin \mathcal{I}_s} \frac{1}{l_j'(x_j^{OE})}}$$

- In particular, for two links, the OE prices are given by

$$p_i^{OE} = x_i^{OE}(l_1'(x_1^{OE}) + l_2'(x_2^{OE})).$$

  - Increase in price over the marginal congestion cost.

Efficiency Bound for Parallel Links

- **Recall our efficiency metric**: Given a set of latency functions $\{l_i\}$ and an equilibrium flow x^{OE}, we define the efficiency metric as

$$\alpha(\{l_i\}, x^{OE}) = \frac{R\sum_{i=1}^{I} x_i^{OE} - \sum_{i=1}^{I} l_i(x_i^{OE})x_i^{OE}}{R\sum_{i=1}^{I} x_i^{S} - \sum_{i=1}^{I} l_i(x_i^{S})x_i^{S}}.$$

Theorem [Acemoglu, Ozdaglar 05]: Consider a parallel link network with inelastic traffic. Then

$$\alpha(\{l_i\}, x^{OE}) \geq \frac{5}{6}, \qquad \forall\ \{l_i\}_{i\in\mathcal{I}},\ x^{OE},$$

and the bound is tight irrespective of the number of links and market structure.

Proof Idea:

- Lower bound the infinite dimensional optimization problem by a finite dimensional problem.
- Use the special structure of parallel links to analytically solve the optimization problem.

Contrasts (superficially) with the intuition that with large number of oligopolists equilibrium close to competitive.

Extensions

- **General network topologies** [Acemoglu, Ozdaglar 06],[Chawla, Roughgarden 08]
  - With serial provider competition, efficiency can be worse.
  - Bounds under additional assumptions on network and demand structure
  - Regulation and cooperation may be necessary
- **Elastic traffic:** Routing and flow control [Hayrapetyan et al. 06], [Ozdaglar 06], [Musacchio and Wu 07]
- Models for investment and capacity upgrade decisions [Acemoglu, Bimpikis, Ozdaglar 07], [Weintraub, Johari, Van Roy 06]
- **Atomic players:** Users that control large portion of traffic (models coalitions) [Cominetti, Correa, Stier-Moses 06, 07], [Bimpikis, Ozdaglar 07]
- **Two-sided markets:** interactions of content providers, users, and service providers [Musacchio, Schwartz, Walrand 07]
- Are networks leading to worst case performance likely? (Genericity analysis)

Game Theory Primer–II

- Convexity often fails in many game-theoretic situations, including wireless network games
- Are there any other structures that we can exploit in games for:
  - analysis of equilibria
  - design of distributed dynamics that lead to equilibria
- **Games with special structure**
  - Supermodular Games
  - Potential Games

Supermodular Games

- Supermodular games are those characterized by **strategic complementarities**
- Informally, this means that the **marginal utility of increasing a player's strategy raises with increases in the other players' strategies**.
  - Implication $\Rightarrow$ best response of a player is a nondecreasing function of other players' strategies
- Why interesting?
  - They arise in many models.
  - Existence of a pure strategy equilibrium without requiring the quasi-concavity of the payoff functions.
  - Many solution concepts yield the same predictions.
  - The equilibrium set has a smallest and a largest element.
  - They have nice sensitivity (or comparative statics) properties and behave well under a variety of distributed dynamic rules.

Monotonicity of Optimal Solutions

- The machinery needed to study supermodular games is lattice theory and monotonicity results in lattice programming
  - Methods used are **non-topological and they exploit order properties**
- We first study the monotonicity properties of optimal solutions of parametric optimization problems:

$$x(t) \in \arg\max_{x \in X} f(x,t),$$

  where $f : X \times T \to \mathbb{R}$, $X \subset \mathbb{R}$, and T is some partially ordered set.
  - We will focus on $T \subset \mathbb{R}^K$ with the usual **vector order**, i.e., for some $x, y \in T$, $x \geq y$ if and only if $x_i \geq y_i$ for all $i = 1, \ldots, k$.
  - Theory extends to general lattices
- We are interested in conditions under which we can establish that $x(t)$ is a nondecreasing function of t.

Increasing Differences

- Key property: Increasing differences

Definition: Let $X \subseteq \mathbb{R}$ and T be some partially ordered set. A function $f : X \times T \to \mathbb{R}$ has **increasing differences** in (x, t) if for all $x' \geq x$ and $t' \geq t$, we have

$$f(x', t') - f(x, t') \geq f(x', t) - f(x, t).$$

- incremental gain to choosing a higher x (i.e., x' rather than x) is greater when t is higher, i.e., $f(x', t) - f(x, t)$ is nondecreasing in t.

Lemma: Let $X \subseteq \mathbb{R}$ and $T \subset \mathbb{R}^k$ for some k, a partially ordered set with the usual vector order. Let $f : X \times T \to \mathbb{R}$ be a twice continuously differentiable function. Then, the following statements are equivalent:

(a) The function f has increasing differences in (x, t).

(b) For all $t' \geq t$ and all $x \in X$, we have

$$\frac{\partial f(x, t')}{\partial x} \geq \frac{\partial f(x, t)}{\partial x}.$$

(c) For all $x \in X$, $t \in T$, and all $i = 1, \ldots, k$, we have

$$\frac{\partial^2 f(x, t)}{\partial x \partial t_i} \geq 0.$$

Examples–I

Example: Network effects (positive externalities).

- A set $\mathcal{I}$ of users can use one of two technologies X and Y (e.g., Blu-ray and HD DVD)
- $B_i(J, k)$ denotes payoff to i when a subset J of users use technology k and $i \in J$
- There exists a network effect or positive externality if

  $$B_i(J, k) \leq B_i(J', k), \qquad \text{when } J \subset J',$$

  i.e., player i better off if more users use the same technology as him.
- Leads naturally to a strategic form game with actions $S_i = \{X, Y\}$
- Define the order $Y \succeq X$, which induces a lattice structure
- Given $s \in S$, let $X(s) = \{i \in \mathcal{I} \mid s_i = X\}$, $Y(s) = \{i \in \mathcal{I} \mid s_i = Y\}$.
- Define the payoffs as

  $$u_i(s_i, s_{-i}) = \begin{cases} B_i(X(s), X) & \text{if } s_i = X, \\ B_i(Y(s), Y) & \text{if } s_i = Y \end{cases}$$

- Show that the payoff functions of this game feature increasing differences.

Examples –II

Example: Cournot duopoly model.

- Two firms choose the quantity they produce $q_i \in [0, \infty)$.
- Let $P(Q)$ with $Q = q_i + q_j$ denote the inverse demand (price) function. Payoff function of each firm is $u_i(q_i, q_j) = q_i P(q_i + q_j) - cq_i$.
- Assume $P'(Q) + q_i P''(Q) \leq 0$ (firm i's marginal revenue decreasing in q_j).
- Show that the payoff functions of the transformed game defined by $s_1 = q_1$, $s_2 = -q_2$ has increasing differences in (s_1, s_2).

Monotonicity of Optimal Solutions

Theorem [Topkis 79]: Let $X \subset \mathbb{R}$ be a compact set and T be some partially ordered set. Assume that the function $f : X \times T \to \mathbb{R}$ is upper semicontinuous in x for all $t \in T$ and has increasing differences in (x, t). Define $x(t) = \arg\max_{x \in X} f(x, t)$. Then, we have:

1. For all $t \in T$, $x(t)$ is nonempty and has a greatest and least element, denoted by $\bar{x}(t)$ and $\underline{x}(t)$ respectively.
2. For all $t' \geq t$, we have $\bar{x}(t') \geq \bar{x}(t)$ and $\underline{x}(t') \geq \underline{x}(t)$.

- If f has increasing differences, the set of optimal solutions $x(t)$ is non-decreasing in the sense that the largest and the smallest selections are non-decreasing.

Supermodular Games

Definition: The strategic game $\langle \mathcal{I}, (S_i), (u_i) \rangle$ is a supermodular game if for all i:

1. S_i is a compact subset of $\mathbb{R}$ (or more generally S_i is a complete lattice in $\mathbb{R}^{m_i}$),
2. u_i is upper semicontinuous in s_i, continuous in s_{-i},
3. u_i has increasing differences in (s_i, s_{-i}) [or more generally u_i is supermodular in (s_i, s_{-i}), which is an extension of the property of increasing differences to games with multi-dimensional strategy spaces].

- Apply Topkis' Theorem to best response correspondences

Corollary: Assume $\langle \mathcal{I}, (S_i), (u_i) \rangle$ is a supermodular game. Let

$$B_i(s_{-i}) = \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}).$$

Then:

1. $B_i(s_{-i})$ has a greatest and least element, denoted by $\bar{B}_i(s_{-i})$ and $\underline{B}_i(s_{-i})$.
2. If $s'_{-i} \geq s_{-i}$, then $\bar{B}_i(s'_{-i}) \geq \bar{B}_i(s_{-i})$ and $\underline{B}_i(s'_{-i}) \geq \underline{B}_i(s_{-i})$.

Existence of a Pure Nash Equilibrium

- Follows from Tarski's fixed point theorem

Theorem [Tarski 55]: Let S be a compact sublattice of $\mathbb{R}^k$ and $f : S \to S$ be an increasing function (i.e., $f(x) \le f(y)$ if $x \le y$). Then, the set of fixed points of f, denoted by E, is nonempty.

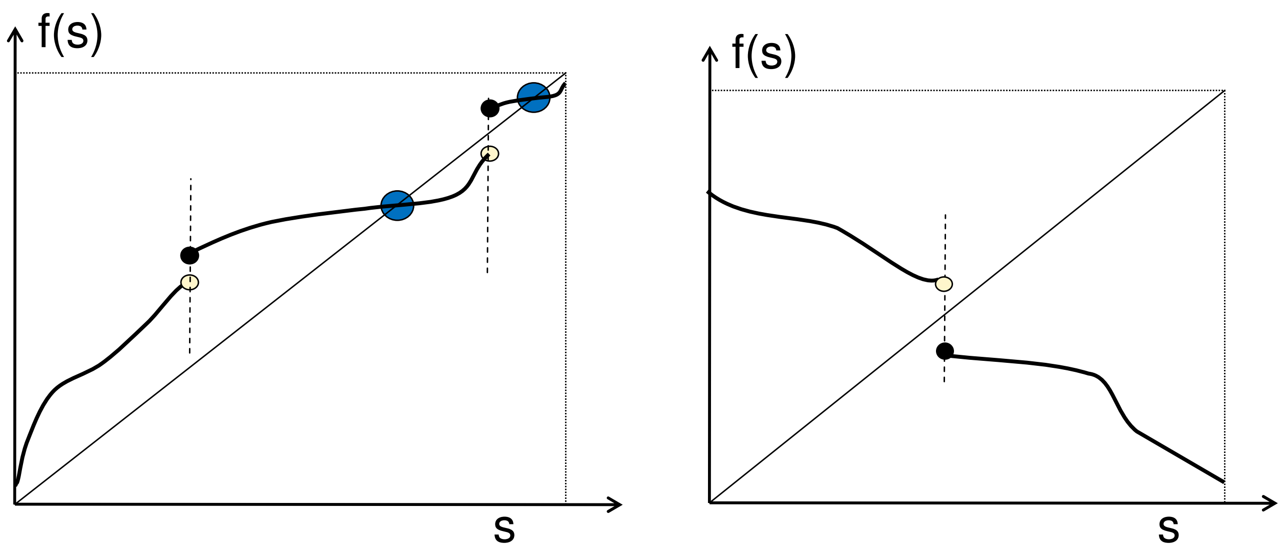

- Apply Tarski's fixed point theorem to best response correspondences

Main Result

- A different approach to understand the structure of Nash equilibria.

Theorem [Milgrom, Roberts 90]: Let $\langle \mathcal{I}, (S_i), (u_i) \rangle$ be a supermodular game. Then the set of strategies that survive iterated strict dominance (i.e., iterated elimination of strictly dominated strategies) has greatest and least elements $\bar{s}$ and $\underline{s}$, which are both pure strategy Nash Equilibria.

Proof idea: Start from the largest or smallest strategy profile and iterate the best-response mapping.

Corollary: Supermodular games have the following properties:

1. Pure strategy NE exist.

2. The largest and smallest strategies are compatible with iterated strict dominance (ISD), rationalizability, correlated equilibrium, and Nash equilibrium are the same.

3. If a supermodular game has a unique NE, it is dominance solvable (and lots of learning and adjustment rules converge to it, e.g., best-response dynamics).

Potential Games

Example: Cournot competition.

- n firms choose quantity $q_i \in (0, \infty)$
- The payoff function for player i given by $u_i(q_i, q_{-i}) = q_i(P(Q) - c)$.
- We define the function $\Phi(q_1, \cdots, q_n) = q_1 \cdots q_n(P(Q) - c)$
- Note that for all i and all q_{-i},

$$u_i(q_i, q_{-i}) - u_i(q_i', q_{-i}) > 0 \quad \text{iff} \quad \Phi(q_i, q_{-i}) - \Phi(q_i', q_{-i}) > 0, \text{ for all } q_i, q_i' \in (0, \infty).$$

- Φ is an **ordinal potential function** for this game.

Example: Cournot competition.

- $P(Q) = a - bQ$ and arbitrary costs $c_i(q_i)$
- We define the function
  $\Phi^*(q_1, \cdots, q_n) = a \sum_{i=1}^n q_i - b \sum_{i=1}^n q_i^2 - b \sum_{1 \leq i < l \leq n}^n q_i q_l - \sum_{i=1}^n c_i(q_i)$.
- It can be shown that for all i and all q_{-i},

$$u_i(q_i, q_{-i}) - u_i(q_i', q_{-i}) = \Phi^*(q_i, q_{-i}) - \Phi^*(q_i, q_{-i}'), \text{ for all } q_i, q_i' \in (0, \infty).$$

- Φ is an **(exact) potential function** for this game.

Potential Functions

Definition [Monderer and Shapley 96]:

(i) A function $\Phi : S \to \mathbb{R}$ is called an **ordinal potential function** for the game G if for all i and all $s_{-i} \in S_{-i}$,

$$u_i(x, s_{-i}) - u_i(z, s_{-i}) > 0 \quad \text{iff} \quad \Phi(x, s_{-i}) - \Phi(z, s_{-i}) > 0, \quad \text{for all } x, z \in S_i.$$

(ii) A function $\Phi : S \to \mathbb{R}$ is called a **potential function** for the game G if for all i and all $s_{-i} \in S_{-i}$,

$$u_i(x, s_{-i}) - u_i(z, s_{-i}) = \Phi(x, s_{-i}) - \Phi(z, s_{-i}), \quad \text{for all } x, z \in S_i.$$

G is called an ordinal (exact) potential game if it admits an ordinal (exact) potential.

Remarks:

- A global maximum of an ordinal potential function is a pure Nash equilibrium (there may be other pure NE, which are local maxima)
  - Every finite ordinal potential game has a pure Nash equilibrium.
- Many learning dynamics (such as 1-sided better reply dynamics, fictitious play, spatial adaptive play) "converge" to a pure Nash equilibrium [Monderer and Shapley 96], [Young 98], [Marden, Arslan, Shamma 06, 07]

Congestion Games

- Congestion games arise when users need to share resources in order to complete certain tasks
  - For example, drivers share roads, each seeking a minimal cost path.
  - The cost of each road segment adversely affected by the number of other drivers using it.
- **Congestion Model:** $C = \langle N, M, (S_i)_{i\in N}, (c^j)_{j\in M}\rangle$ where
  - $N = \{1, 2, \cdots, n\}$ is the set of players,
  - $M = \{1, 2, \cdots, m\}$ is the set of resources,
  - S_i consists of sets of resources (e.g., paths) that player i can take.
  - $c^j(k)$ is the cost to each user who uses resource j if k users are using it.
- Define congestion game $\langle N, (S_i), (u_i)\rangle$ with utilities $u_i(s_i, s_{-i}) = \sum_{j\in s_i} c^j(k_j)$, where k_j is the number of users of resource j under strategies s.

Theorem [Rosenthal 73]: Every congestion game is a potential game.

Proof idea: Verify that the following is a potential function for the congestion game:

$$\Phi(s) = \sum_{j\in\cup s_i} \Big(\sum_{k=1}^{k_j} c^j(k)\Big)$$

Network Games–II

- In the presence of heterogeneity in QoS requests, the resource allocation problem becomes nonstandard
  - Traditional network optimization techniques **information intensive, rely on tight closed-loop controls, and non-robust against dynamic changes**
- Recent literature used game-theoretic models for resource allocation among heterogeneous users in wireline and wireless networks
  - User terminals: players competing for network resources
  - Compatible with self-interested nature of users
  - Leads to distributed control algorithms
- **Utility-maximization framework** of market economics, to provide different access privileges to users with different QoS requirements in a distributed manner [Kelly 97], [Kelly, Maulloo, Tan 98], [Low and Lapsley 99], [Srikant 04]
  - Each user (or equivalently application) represented by a utility function that is a measure of his preferences over transmission rates.
- **For wireless network games:** Negative externality due to interference effects

Wireless Games

- Most focus on infrastructure networks, where users transmit to a common concentration point (base station in a cellular network or access point)
- **Actions:** Transmit power, transmission rate, modulation scheme, multi-user receiver, carrier allocation strategy etc.
- **Utilities:** Received signal-to-interference-noise ratio (SINR) measure of quality of signal reception for the wireless user:

$$\gamma_i = \frac{p_i h_i}{\sigma^2 + \sum_{j \neq i} p_j h_j},$$

where σ^2 is the noise variance (assuming an additive white Gaussian noise channel), and h_i is the channel gain from mobile i to the base station.

Examples of Utility Functions

- **Spectral Efficiency** [Alpcan, Basar, Srikant, Altman 02], [Gunturi, Paganini 03]

$$u_i = \xi_i \log(1 + \gamma_i) - c_i p_i,$$

where ξ_i is a user dependent constant and c_i is the price per unit power.

- **Energy Efficiency** [Goodman, Mandayam 00]

$$u_i = \frac{Throughput}{power} = \frac{R_i f(\gamma_i)}{p_i} \quad bits/joule,$$

where R_i is the transmission rate for user i and $f(\cdot)$ is an **efficiency function** that represents packet success rate (assuming packet retransmission if one or more bit errors)

  - $f(\gamma)$ depends on details of transmission: modulation, coding, packet size
  - **Examples**: $f(\gamma) = (1 - 2Q(\sqrt{2\gamma}))^M$ (BPSK modulation), (where M is the packet size, and $Q(\cdot$ is the complementary cumulative distribution function of a standard normal random variable), $f(\gamma) = (1 - e^{-\gamma/2})$ (FSK modulation)
  - In most practical cases, $f(\gamma)$ is strictly increasing and has a sigmoidal shape.

Wireless Power Control Game

- Power control in cellular CDMA wireless networks
- It has been recognized that in the presence of interference, the strategic interactions between the users is that of **strategic complementarities** [Saraydar, Mandayam, Goodman 02], [Altman and Altman 03]

Model:

- Let $L = \{1, 2, ..., n\}$ denote the set of users (nodes) and

$$\mathcal{P} = \prod_{i \in L} [P_i^{\text{min}}, P_i^{\text{max}}] \subset \mathbb{R}^n$$

  denote the set of power vectors $p = [p_1, \ldots, p_n]$.

- Each user is endowed with a utility function $f_i(\gamma_i)$ as a function of its SINR γ_i.
- The payoff function of each user represents a tradeoff between the payoff obtained by the received SINR and the power expenditure, and takes the form

$$u_i(p_i, p_{-i}) = f_i(\gamma_i) - cp_i.$$

Increasing Differences

- Assume that each utility function satisfies the following assumption regarding its **coefficient of relative risk aversion**:

$$\frac{-\gamma_i f_i''(\gamma_i)}{f_i'(\gamma_i)} \geq 1, \qquad \text{for all } \gamma_i \geq 0.$$

  - Satisfied by α-fair functions $f(\gamma) = \frac{\gamma^{1-\alpha}}{1-\alpha}$, $\alpha > 1$ [Mo, Walrand 00], and the efficiency functions introduced earlier

- Show that for all $i = 1 \ldots, n$, the function $u_i(p_i, p_{-i})$ has increasing differences in (p_i, p_{-i}).

Implications:

- Power control game has a pure Nash equilibrium.
- The Nash equilibrium set has a largest and a smallest element, and there are **distributed algorithms** that will converge to any of these equilibria.
- These algorithms involve each user updating their power level locally (based on total received power at the base station).

Extensions–I

- **Distributed multi-user power control in digital subscriber lines** [Yu, Ginis, Cioffi 02], [Luo, Pang 06]
  - Model as a Gaussian parallel interference channel
  - Each user chooses its power spectral density to maximize rate subject to a power budget
  - Existence of a Nash equilibrium (due to concavity)
  - Best-response dynamics leads to a **distributed iterative water filling algorithm**, where each user optimizes its power spectrum treating other users' interference as noise
  - Convergence analysis by [Luo, Pang 06] based on Linear Complementarity Problem Formulation
- **Power control in CDMA-based networks (no fading)**
  - Single and multi-cell [Saraydar, Mandayam, and Goodman 02], [Alpcan, Basar, Srikant and Altman 02]
  - Joint power control-receiver design [Meshkati, Poor, Schwartz, Mandayam 05]
  - Adaptive modulation and delay constraints [Meshkati, Goldsmith, Poor, Schwartz 07]

Extensions–II

- **Distributed control in collision channels**
  - Aloha-like framework [Altman, El-Azouzi, Jimenez 04], [Inaltekin, Wicker 06]
  - Single packet perspective [MacKenzie, Wicker 03], [Fiat, Mansour, Nadav 07]
  - Rate-based equilibria [Jin, Kesidis 02], [Menache, Shimkin 07]
- **Power control and transmission scheduling in wireless fading channels ("Water-Filling" Games)**
  - CDMA-like networks [Altman, Avrachenko, Miller, Prabhu 07], [Lai and El-Gamal 08]
  - Collision Channels [Menache, Shimkin 08], [Cho, Hwang, Tobagi 08]
- **Jamming Games** [Altman, Avrachenkov, Garnaev 07], [Gohary, Huang, Luo, Pang 08]

Extensions–III

Game Theory for Nonconvex Distributed Optimization:

- Distributed Power Control for Wireless Adhoc Networks [Huang,Berry,Honig 05]
  - Two models: Single channel spread spectrum, Multi-channel orthogonal frequency division multiplexing
  - Asynchronous distributed algorithm for optimizing total network performance
  - Convergence analysis in the presence of nonconvexities using **supermodular game theory**
- Distributed Cooperative Control–“Constrained Consensus” [Marden, Arslan, Shamma 07]
  - Distributed algorithms to reach consensus in the “values of multiple agents” (e.g. averaging and rendezvous problems)
  - Nonconvex constraints in agent values
  - Design a game (i.e., utility functions of players) such that
    - The resulting game is a **potential game** and the Nash equilibrium “coincides” with the social optimum
    - Use learning dynamics for potential games to design distributed algorithms with favorable convergence properties

Network Design

- Sharing the cost of a designed network among participants [Anshelevich et al. 05]

Model:

- Directed graph $N = (V, E)$ with edge cost $c_e \geq 0$, k players
- Each player i has a set of nodes T_i he wants to connect
- A strategy of player i set of edges $S_i \subset E$ such that S_i connects to all nodes in T_i

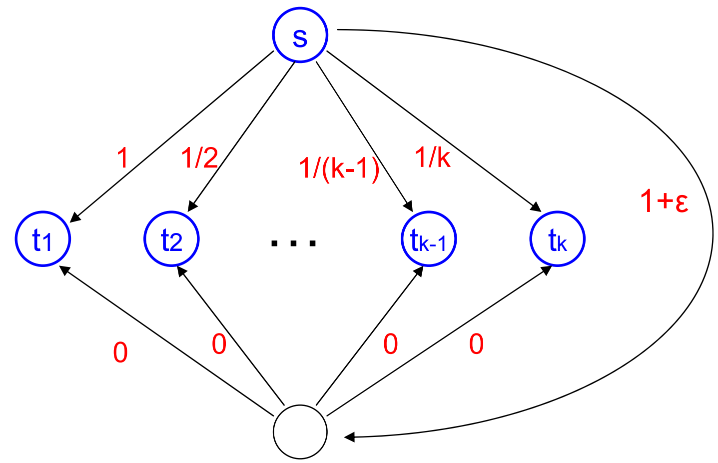

Optimum cost: 1+ε

Unique NE cost: $\sum_{i=1}^{k} 1/i = H(k)$

- **Cost sharing mechanism:** All players using an edge split the cost equally
- Given a vector of player's strategies $S = (S_1, \ldots, S_k)$, the cost to agent i is $C_i(S) = \sum_{e \in S_i} (c_e / x_e)$, where x_e is the number of agents whose strategy contains edge e

Price of Stability for Network Design Game

- The price of anarchy can be arbitrarily bad.
  - Consider k players with common source s and destination t, and two parallel edges of cost 1 and k.
- We consider the worst performance of the **best Nash equilibrium** relative to the system optimum.
  - Price of Stability

Theorem: The network design game has a pure Nash equilibrium and the price of stability is at most $H(k) = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{k}$.

Proof idea: The game is a **congestion game**, implying existence of a pure Nash equilibrium from [Rosenthal 73]. Use the potential function to establish the bound.

Extensions:

- Congestion effects
- More general cost-sharing mechanisms and their performance

Concluding Remarks

- New emerging control paradigm for large-scale networked-systems based on game theory and economic market mechanisms
- Many applications of decentralized network control
  - Sensor networks, mobile ad hoc networks
  - Large-scale data networks, Internet
  - Transportation networks
  - Power networks
- Future Challenges
  - Models for understanding when local competition yields efficient outcomes
  - Dynamics of agent interactions over large-scale networks
  - Distributed algorithm design in the presence of incentives and network effects