Game Theory and Algorithms*
Lecture 17: Network Games and Quality of Equilibria

May 10, 2011

Summary: We discuss a pair of games for *routing traffic* and *connecting networks*. They are games that take place within a network, where each player's action is to pick a path. We give examples showing how the social cost can go up unexpectedly, due to greedy users. The prices of anarchy and stability measure the social quality of Nash equilibria compared to optimal.

This week, we will see two network games:

1. network routing with a continuum of users (players drive on congested paths);
2. network formation with discrete users (players build paths & pay for their costs).

In both games we use the *potential function* method to compare equilibria to socially-optimal solutions. We will start with some small examples of the routing game; then we introduce potential functions and deal completely with the formation game; finally (next lecture) we go back to the routing game and analyze it in more detail.

1 Selfish Routing: Examples

An instance of the routing game is specified by:

- A directed graph
- A designated source node and sink node in the graph
- For each edge e, its *delay* as a function of its *congestion*, specified by a function

$$\text{delay}_e : [0, 1] \to \mathbb{R}_+.$$

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

The idea is that we have a total 1 unit of traffic flow which we want to send from the source to the sink. The delay time to traverse an edge depends on the fraction of the flow using it, and the total time of a path is the sum of the delays of the individual edges. The interesting part is now to imagine that each infinitesimal part of this flow is a greedy user, which will prefer to switch from a high-delay path to a lower-delay path, if one is available. We will study Nash equilibria in this setting.

We illustrate some details of the setup by way of an example. Suppose the graph consists of two parallel edges t and b from the source to the sink, both of whose congestion functions are the identity, $\mathrm{delay}_t(z) = \mathrm{delay}_b(z) = z$ for all $z \in [0, 1]$. We illustrate in Figure 1.

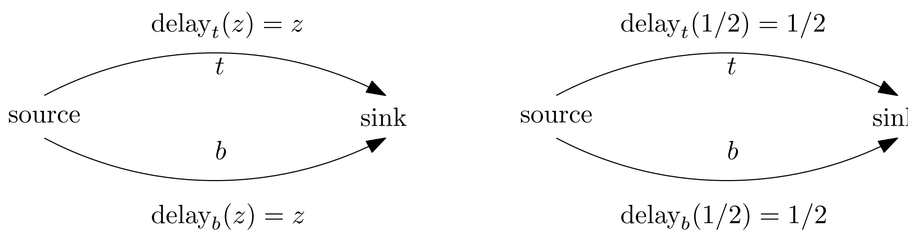

Figure 1: A simple flow routing instance, and its Nash equilibrium.

In all cases, a *flow* can be specified in the following way: for every possible source-sink path P, let f_P be the fraction of the flow using path P. Then a *flow* is defined[1] as a collection of $f_P \geq 0$ such that $\sum_P f_P = 1$. In our case, there are only two paths (each one edge long) so a flow simply assigns some fraction to use t, and the rest to use b.

Next we define

- The congestion on edge e for flow f is $f_e := \sum_{P:e\in P} f_P$;
- the delay of edge e for flow f is $\mathrm{delay}_e(f_e)$;
- the delay of path P for flow f is $\mathrm{delay}_P(f) := \sum_{e\in P} \mathrm{delay}_e(f_e)$;
- a flow is *Nash* if, for all paths P with $f_P > 0$ and all paths P',

$$\mathrm{delay}_P(f) \leq \mathrm{delay}_{P'}(f),$$

  i.e. for each path used by a positive fraction of users, no other path is faster.

In our example, which flows are Nash? Let $0 \leq x \leq 1$ be the fraction of flow assigned to the top path, so $1 - x$ is assigned to the bottom; write f^x for this flow. The delay of path t is x and the delay of path b is $1 - x$. Notice that the flow f^1 with $x = 1$, i.e. all flow through t, is not Nash since users on path t incurring delay $\mathrm{delay}_t(f^1) = 1$ would prefer to switch to the faster path b with $\mathrm{delay}_b(f^1) = 0$. In fact, the only Nash flow has $x = 1/2$: at this point no user on either path prefers to switch to the other path.

[1] Network flows also can be represented by edge values, which is a polynomial-size representation. The path representation is equivalent, and although exponential-size in the worst case, it is easier for our exposition.

1.1 Social Cost and Pigou's Example

Nash flows correspond to Nash equilibria when there are a continuum of users who all want to minimize their own delay. On the other hand, if a single designer were controlling all of the flow, they would want to minimize the *average delay/social cost*

$$\text{social-cost} := \sum_P f_P \cdot \text{delay}_P(f).$$

In the previous two-link example, the Nash flow also turned out to give a min-cost flow, i.e. it minimized the average delay. Pigou's example (1920) shows that this is not always the case. Suppose we modify the previous example by changing the top path t to have a constant delay of 1, no matter how many users take it: $\forall z, \text{delay}_t(z) = 1$. We illustrate in Figure 2.

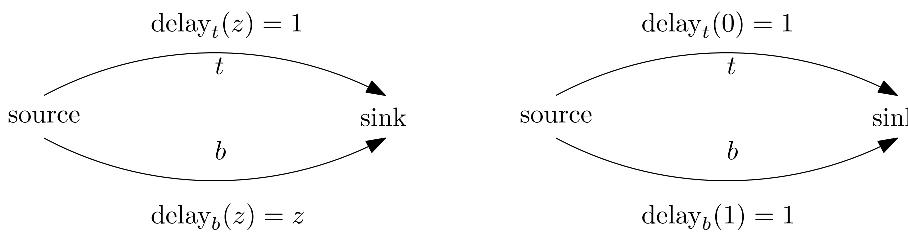

Figure 2: Pigou's example: the Nash flow (right) is not socially optimal.

In this case, the flow $f^{1/2}$ is no longer Nash: the top users have delay 1, and would prefer to switch to the bottom path with delay 1/2. In fact, if any positive fraction of users use the top path, they prefer to switch to the bottom. The unique Nash flow is f^0 which sends all of the flow along the bottom path.

On the other hand, the social cost of f^x is

$$\sum_P f_P \cdot \text{delay}_P(f^x) = x \cdot 1 + (1-x) \cdot (1-x) = 1 - x + x^2;$$

thus the social cost is 1 for the Nash flow f^0, but is minimized at 3/4 for the flow $f^{1/2}$. So *Nash flows are not socially optimal.*

Definition 1. The *price of anarchy* is the maximum social cost of any Nash flow divided by the optimal social cost:

$$PoA := \max\{\text{social-cost}(f) \mid f \text{ Nash}\} / \min\{\text{social-cost}(f) \mid \text{any } f\}.$$

In Pigou's example, the price of anarchy is $1/(3/4) = 4/3$. Can the price of anarchy be bounded? It turns out that for nondecreasing affine (= linear + constant) delay functions, 4/3 is exactly as bad as it gets! This will be an exercise next class, using *variational inequalities.*

Exercise. Modify Pigou's example by replacing the bottom function with a power function, $\text{delay}_b(z) = z^c$. What lower bound does Pigou's bound give on the price of anarchy? How does this change as $c \to \infty$?

In fact, variational inequality methods show that Pigou-like examples give the worst-case example for PoA for many natural classes of functions; see page 473 of Nisan et al.

1.1.1 Braess' Example

An even more striking example was given by Braess (1968). Consider the network at left, where the arcs have either constant unit delay or identity $\text{delay}_e(z) = z$. Then, add a "free arc" with $\text{delay}_e(z) = 0$ for all z. We illustrate in Figure 3.

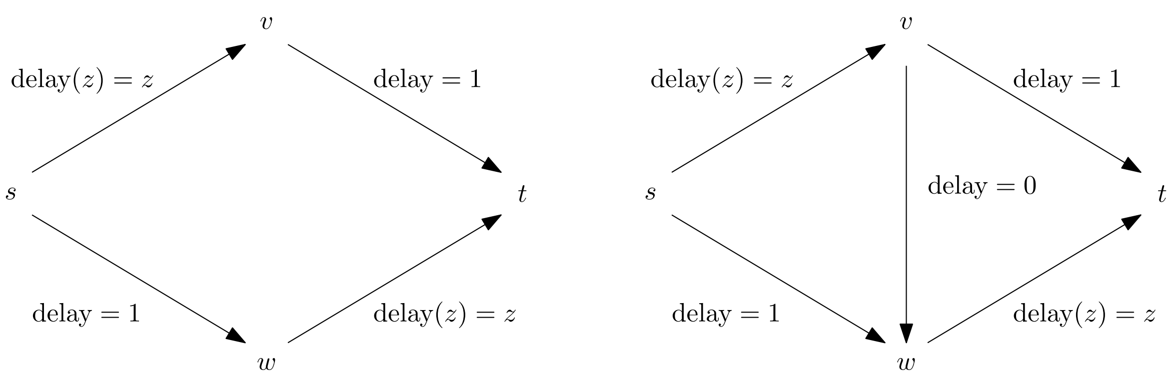

Figure 3: Braess' example: a free edge increases the social cost of the Nash equilibrium.

In the original network, a little calculation (like in our two-link example) shows that the only Nash flow puts half of the traffic on the top path, and half on the bottom path, giving delay 3/2 for every path and social cost 3/2. But in the new network, any positive traffic using a unit-delay edge has incentive to switch to the longer path $svwt$. So the only Nash flow has $f_{svwt} = 1$ and social cost 2. By adding a free arc, the delay went up! This phenomenon has been reported, for example, in New York City: when a certain road was under repair, the traffic sped up.

Exercise (Short but important)**.** Using the fact that the price of anarchy is at most 4/3, show that the factor 4/3 is the worst that can arise in Braess' paradox, for nondecreasing affine delay functions.

2 Network Formation and Potential Functions

Now we switch gears to another game, where there are a discrete set of players building a network. We have:

- A directed graph,

- n players each with a source s_i and sink t_i (possibly overlapping),
- construction costs c_e for each directed edge e.

From this we construct a strategic game:

- The action set of each player is the set of all s_i-t_i paths; P_i is player i's choice;
- all edges on $\bigcup_i P_i$ are built;
- the cost of each edge e is split amongst all players using it, so

  $$u_i(P) := -\mathrm{cost}_i(P) = -\sum_{e \in P_i} c_e/\mathrm{usage}(e, P)$$

  where $\mathrm{usage}(e, P) = |\{j \mid e \in P_j\}|$.

2.1 Price of Anarchy

This time, an action profile P is Nash if it forms a Nash equilibrium of the strategic game, i.e. if no player can save money by switching unilaterally to a different path. The *social cost* of a strategy profile is the negative sum of all utilities, which equals the sum of all edges constructed:

$$\text{social-cost}(P) = \sum_i \mathrm{cost}_i = \sum_{e \in \bigcup_i P_i} c_e.$$

Claim 2. *The price of anarchy is at most n, the number of players.*

Proof. Consider a P with optimal social cost. We claim that in any Nash equilibrium, $\mathrm{cost}_i \leq \text{social-cost}(P)$; indeed, i can switch to their path P_i in P and even if they don't share edges, they pay only for $P_i \subseteq \bigcup_i P_i$ and so their personal cost is at most social-cost(P).

So in any Nash equilibrium P',

$$\text{social-cost}(P') = \sum_{i=1}^{n} \mathrm{cost}_i(P') \leq \sum_{i=1}^{n} \text{social-cost}(P),$$

proving the price of anarchy is at most n. □

In fact, this bound is tight. Consider the graph shown on the left of Figure 4. The social optimum sends everyone on the top path, and has social cost 1. But if everyone uses the bottom path, they each pay $n/n = 1$ and cannot improve by deviating to the top path. The bottom path has social cost n.

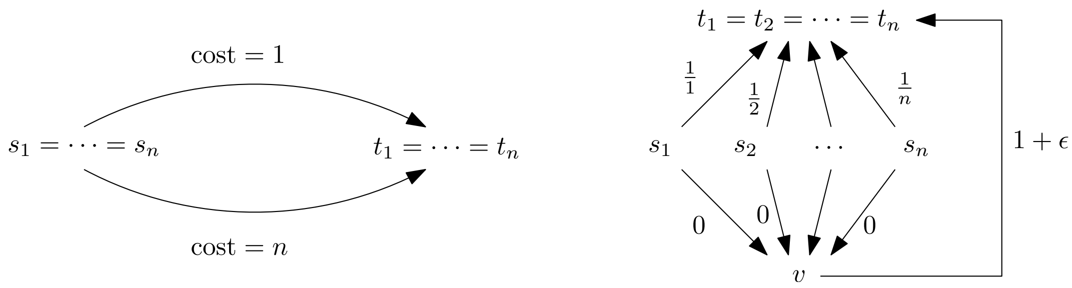

Figure 4: Left: the price of anarchy is at least n since some Nash equilibrium has social cost n. Right: the price of stability is at least $H_n/(1+\epsilon)$ since *all* Nash equilibria have social cost H_n. Edge costs are shown.

2.2 Stability and Potentials

The network formation game has a much better cost if we compare to the *best* equilibrium.

Definition 3. The *price of stability* is the *minimum* social cost of any Nash flow divided by the optimal social cost:

$$PoS := \min\{\text{social-cost}(f) \mid f \text{ Nash}\} / \min\{\text{social-cost}(f) \mid \text{any } f\}.$$

So you can imagine that the social designer only needs to put the users in a stable state (Nash equilibrium) where they will stay; whereas in anarchy the users will arrive at some Nash equilibrium the designer cannot control. Today we will show:

Theorem 4. *In the network formation game, the price of stability is* $H_n = 1 + \frac{1}{2} + \cdots + \frac{1}{n} \sim \ln n$.

The lower bound is shown by a simple example; the upper bound introduces the very important *potential function method.*

Example 5 (Lower bound)**.** Consider the network shown on the right of Figure 4, where all users have separate sources, and one common sink, as well as a vertex v providing an alternate path to the sink. There is only one Nash equilibrium, since user n is forced to use the direct route in a Nash equilibrium, which forces user $n-1$ to use the direct route, etc. Since the all-direct route has social cost H_n, and the all-indirect route has social cost $1+\epsilon$, the best bound on the price of stability cannot be better than H_n.

2.3 Potential Games

A strategic game $(A_1, A_2, \ldots, A_n, u_1, u_2, \ldots, u_n)$ is called a *potential game* when there exists a function $\Phi : A_1 \times A_2 \times \cdots \times A_n \to \mathbb{R}$ such that for all i, all a, and all a'_i,

$$u(a'_i, a_{-i}) - u(a_i, a_{-i}) = -(\Phi(a'_i, a_{-i}) - \Phi(a_i, a_{-i})),$$

or in words *every unilateral improvement in utility (decrease in cost) corresponds to an equal decrease in the potential function.*

Not every strategic game is a potential game. Also, a potential game has more than one potential, but **Exercise**: show that all potential functions Φ for a given potential game differ only by a constant.

Why are potential games important? Here is the first reason.

Theorem 6. *Every potential game with a finite number of actions has a pure Nash equilibrium.*

Proof. Let a^* be an action profile which minimizes the potential function. Any deviation by any player cannot decrease the potential function, so it cannot increase their utility. $\square$

The theorem also generalizes to infinite action sets that are *compact*, provided the utilities are continuous. Moreover, the proof also shows that

- *best response dynamics* must eventually terminate in potential games, since the potential function cannot be decreased an infinite number of times;
- Nash equilibria are the same as coordinate-wise *local minima* of the potential function.

We will see several other applications of potential functions this week. Both the network routing and formation games are potential games; we show the latter next.

Proposition 7. *The network formation game is a potential game with*

$$\Phi(P) = \sum_e c_e \cdot H_{\text{usage}(e,P)} = \sum_e c_e(1 + \frac{1}{2} + \cdots + \frac{1}{\text{usage}(e,P)}).$$

Proof. We need to show that if a single player deviates, their decrease in cost equals the decrease in the potential function. Let P be some action profile and (P_{-i}, P'_i) be the result of player i deviating.

What is the effect of the deviation on player i's utility? Their cost shares for edges in $P_i \cap P'_i$, or in $E \backslash (P_i \cup P'_i)$, do not change. For each edge e in $P_i \backslash P'_i$ they save $c_e/\text{usage}(e,P)$; for each edge e in $P'_i \backslash P_i$ they are lose $c_e/(\text{usage}(e,P)+1)$.

Likewise, look at the definition of Φ and how the contribution by each edge e changes from P to (P_{-i}, P'_i): it goes down by $c_e/\text{usage}(e,P)$ for edges in $P_i \backslash P'_i$, up by $c_e/(\text{usage}(e,P)+1)$ for edges in $P'_i \backslash P_i$, and is unchanged for other edges. $\square$

Now we combine one easy general fact with one easy specific fact.

Proposition 8. *In a game with potential Φ, if for some constants C, C' and all action profiles a we have*

$$C \cdot \text{social-cost}(a) \le \Phi(a) \le C' \cdot \text{social-cost}(a),$$

then the price of stability is at most C'/C.

Proposition 9. *In the network formation game,* social-cost$(P) \leq \Phi(P) \leq H_n \cdot$social-cost$(P)$*.*

Together, this gives the tight upper bound Theorem 4 on the price of stability for formation games.

Proof of Proposition 8. Essentially, the proposition holds because minimizing the potential gives a Nash equilibrium, and minimizing the potential is "almost" the same as minimizing the social cost.

In detail, take any minimizer a of the potential function, which is a Nash equilibrium. For any other a',

$$\text{social-cost}(a) \leq \Phi(a)/C \leq \Phi(a')/C \leq \text{social-cost}(a') \cdot C'/C.$$

Sine this holds when a' is socially optimal, the price of stability is at most C'/C. □

Proof of Proposition 9. Recall

$$\text{social-cost}(P) = \sum_{e:\text{usage}(e,P)>0} c_e \qquad \text{and} \qquad \Phi(P) = \sum_e c_e \cdot H_{\text{usage}(e,P)}.$$

Unused edges contribute 0 to both sums; used edges contribute c_e to the social cost, and between c_e and $c_e \cdot H_n$ to the potential function. □

This potential function method will also give us a bound on the price of stability for network flow games; and it will turn out that for a large class of delay functions, there is a unique Nash equilibrium, so the price of anarchy equals the price of stability. Afterwards, we will show a few more facts about the network routing game.

Remark: Finding the optimal social cost for this network formation game amounts to the *directed Steiner forest* problem, which is quite hard to approximate ("label-cover-hard"). Moreover, be wary that additional ingredients (e.g. fast convergence of best response dynamics) would need to be added to the previous results to get a polynomial-time algorithm for finding a Nash equilibrium with $O(\alpha \ln n)$-approximately optimal social cost. (I do not know if such a result is actually known.)