CS 6840: Algorithmic Game Theory	Spring 2017

Lecture 2: January 27

Lecturer: Éva Tardos — *Scribe: Albert Julius Liu*

2.1 Logistics

- Scribe notes must be submitted within 24 hours of the corresponding lecture for full credit. They are worth 4% of the course grade.
- See the first lecture's notes for additional examples of dysfunctional tournament incentives.
- Again, Roughgarden's *Twenty Lectures on Algorithmic Game Theory* is not required, but the first third of the course will closely follow the second half of the book. This lecture will follow Section 11.

2.2 Braess' paradox:

Recall from last lecture Braess' paradox:

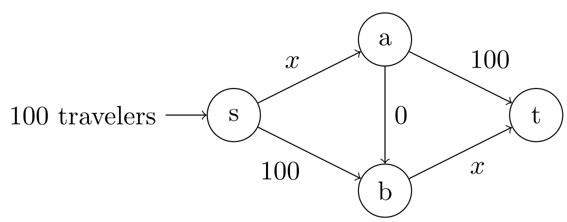

In this *atomic* problem, i.e. with discrete travelers, we had four Nash equilibria: one where all 100 travelers take the 0-cost "bridge" between a and b; two where only 99 do, and the last traveler takes either the top or bottom path; and one where 98 take the bridge and one takes each of the top and bottom path.

What if we instead consider a *non-atomic* version of the problem, where travelers are infinitely small? Let us normalize the flow of travelers to 1, and likewise the constant edge costs to 1 as well:

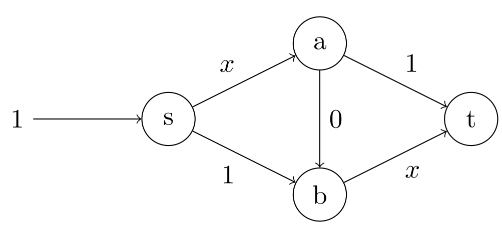

In this case there is only one Nash equilibrium, where the entire flow of travelers takes the bridge.

Proof: To see this, suppose that:

- A fraction a of the flow takes the top path.
- A fraction b takes the bottom path.
- The remaining fraction $1 - a - b$ takes the bridge path.

The costs of the three paths are then

- $2 - b$ for the top path.
- $2 - a$ for the bottom path.
- $2 - a - b$ for the bridge path.

If a and b are not both nonzero, then the flow on the top or bottom path will be incentivized to take the bridge instead, as it is strictly cheaper. If exactly one of a and b are zero, the bottom or top path (whichever one has nonzero flow) will be strictly more expensive than either of the other two paths and so will be incentivized to switch. Therefore, the unique Nash equilibrium occurs when a and b are both zero. ■

In general, this non-atomic formulation is convenient as it allows us to use the powerful tools of calculus. Even if the population being modeled is atomic in real life, this non-atomic formulation may be a good approximation if the population is large enough.

2.3 Routing games

Our Braess' paradox example is an instance of a *routing game*.

2.3.1 The game

An instance of such a game is defined by the following:

- A graph G consisting of a set of vertices V and a set of directed edges E between ordered pairs of vertices.
- For each edge, a cost $c_e(x)$ as a function of the flow on that edge. The cost functions are required to be all nonnegative (for nonnegative flow, which we will require), continuous, and monotonically nondecreasing. While it can be convenient to require that they be *strictly* monotonically increasing, this will not be necessary for this lecture, and is not true in our example above.
- A set of pairs of source and sink vertices s_i, t_i, with a desired flow r_i between each such pair of vertices. The pairs are required to be unique; that is, for any $s_i = s_j$ we require $t_i \neq t_j$, and for any $t_i = t_j$ we require $s_i \neq s_j$. Another way of thinking about it is that we collapse all duplicate source-sink pairs into a single one by summing their required flows $r_i + r_j$.

2.3.2 Solution

A legal global flow f in this game consists of an assignment of flow magnitudes f_P to each paths P such that all the required flows are satisfied, i.e. for all i,

$$r_i = \sum_{P:s_i \to t_i} f_P \tag{2.1}$$

Note that the number of possible paths $|P|$ can become exponentially large in the number of vertices $|V|$, even if only simple paths (i.e. not visiting any vertex more than once) are permitted! While this can pose a problem for computation, it is fine for our analysis here.

2.3.3 Edge flows

We will also consider flow on individual edges e, denoted by $f(e)$. This is simply the sum of the flow of all paths that go through that edge:

$$f(e) = \sum_{P \ni e} f_P \tag{2.2}$$

2.3.4 Path cost

The cost $c_P(f)$ of a given path P given a global flow f is the sum of the cost of edges on that path:

$$c_P(f) = \sum_{e \in P} c_e(f(e)) \tag{2.3}$$

2.3.5 Nash equilibrium

For this game, a Nash equilibrium is where no flow has an incentive to switch from one path to another. More formally, we are at a Nash equilibrium if, for all paths $P : s_i \to t_i$ with nonzero flow $f_P > 0$, and all possible alternative paths $Q : s_i \to t_i$ between the same source and sink nodes, we have

$$c_P(f) \le c_Q(f) \tag{2.4}$$

The converse ("if *and only if*") holds as well—that is, if this condition is not met, there is always an incentive for some flow to switch.

Proof: Suppose for a global flow f there are some P and Q with $f_P > 0$ such that

$$c_P(f) > c_Q(f) \tag{2.5}$$

Equivalently, for some $\delta > 0$,

$$c_P(f) - c_Q(f) = \delta \tag{2.6}$$

This difference is the incentive to switch from P to Q. The question is whether we can maintain an incentive for a nonzero amount of flow switching.

Let ϵ be an amount of flow that we propose to switch from P to Q.

c_P and c_Q are both sums of edge costs, which are continuous; since sums and differences of continuous functions are also continuous, this means that $c_P - c_Q$ varies continuously with ϵ. By the definition of continuity, for some sufficiently small ϵ, $c_P - c_Q$ remains positive, that is, the incentive to switch, remains larger than zero. ■

2.4 Price of anarchy

How bad is it if everyone acts selfishly, compared to if a central coordinator tries to minimize global cost?

2.4.1 Global cost

To answer this question, we must first define global cost. We could define this as the total of all path costs, weighted by the amount of flow taking that path:

$$c(f) = \sum_P f_P c_P(f) \tag{2.7}$$

Or we could define this as the total of all *edge* costs, weighted by the amount of flow taking that edge:

$$c(f) = \sum_e f(e) c_e(f(e)) \tag{2.8}$$

In fact, these two definitions are equivalent.

Proof: Starting with the path-based definition of cost (Equation 2.7), we substitute in the definition of path cost (Equation 2.3):

$$\sum_P f_P \sum_{e \in P} c_e(f(e)) \tag{2.9}$$

We can move f_P inside the inner summation:

$$\sum_P \sum_{e \in P} f_P c_e(f(e)) \tag{2.10}$$

Now we are summing over all P, e pairs such that $e \in P$. It is therefore equivalent to switch the nesting of the summations:

$$\sum_e \sum_{P \ni e} f_P c_e(f(e)) \tag{2.11}$$

Since $c_e(f(e))$ does not depend on P, we can move it outside the inner summation:

$$\sum_e c_e(f(e)) \sum_{P \ni e} f_P \tag{2.12}$$

Finally, using the definition of edge flow (Equation 2.2) we can replace the inner summation

$$\sum_e c_e(f(e)) f(e) \tag{2.13}$$

matching the edge-based global cost definition (Equation 2.8) as desired. ■

2.4.2 Defining price of anarchy

Now we can define the price of anarchy. Let f^* be the global flow that minimizes global cost, and f be the Nash equilibrium (the one with the greatest global cost if it is not unique). The price of anarchy is defined as

$$\text{price of anarchy} = \frac{c(f)}{c(f^*)} \tag{2.14}$$

2.4.3 Example: Braess' paradox

In our Braess' paradox scenario: the socially optimal solution is for half of the flow to take the top path and half the bottom. Every unit of flow experiences a cost of 1.5, so the global cost is 1.5. Meanwhile in the Nash equilibrium the entire flow takes the bridge for a cost of 2, so this example has a price of anarchy of 4/3.

2.4.4 Another example

Here is another example with just two vertices and two edges.

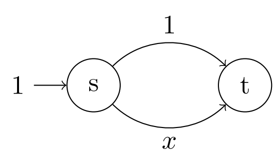

There is only one Nash equilibrium, namely the entire flow taking the bottom edge—otherwise, the bottom edge costs less than the top edge, and the top flow has an incentive to switch to the bottom edge. The global cost for the Nash equilibrium is 1.

To find the optimal solution, we can use calculus. Suppose x flow takes the bottom path, and the rest takes the top path. The global cost is then

$$c(x) = \underbrace{1 - x}_{\text{cost from top edge}} + \underbrace{x^2}_{\text{cost from bottom edge}} \tag{2.15}$$

Remember that the optimum of a differentiable function will occur at the bounds (not in this case, since both bounds have a global cost of 1), or at a local extremum, where the derivative is zero:

$$\frac{dc}{dx} = -1 + 2x \tag{2.16}$$

The minimum thus occurs at $x = \frac{1}{2}$, just like our Braess' paradox example. Plugging x back into the global cost (Equation 2.15), we have a total cost of $\frac{3}{4}$, and a price of anarchy of $\frac{4}{3}$—again just like our Braess' paradox example.

However, unlike our Braess' paradox example, in this case the social solution is in some sense "unfair": the flow on the top pays a unit cost of 1, while the flow on the bottom pays a unit cost of $\frac{1}{2}$.

A Appendix: Undirected graphs

Note that our definition was for a directed graph. It is possible to reduce a game on an undirected graph (where the cost is a function of the total flow in both directions) to a game on a directed graph via the use of a gadget:

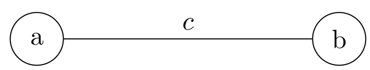

becomes

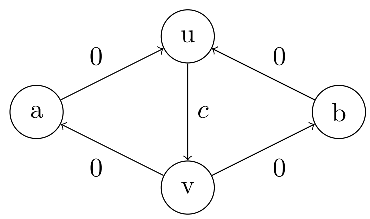

B Appendix: Multigraphs

In our graph example above (Section 2.4.4), we allow multiple edges between the same pair of nodes (with different costs). Again, a game on such a graph can be reduced to a game on a conventional directed graph via the use of a gadget:

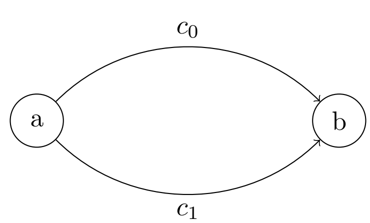

becomes