CS 684 Algorithmic Game Theory **Scribe: Zoya Svitkina**
Instructor: Eva Tardos **May 3, 2004**

1 How bad is Braess' paradox

Let us review what Braess' paradox is, using the same example as we had before. The setting is a routing game on a graph. The delay $d_e(x)$ on edge e is a function of the amount of traffic x sent on e. We assume that $d_e(x)$ is monotone and continuous. In this example, the traffic rate from s to t is 1.

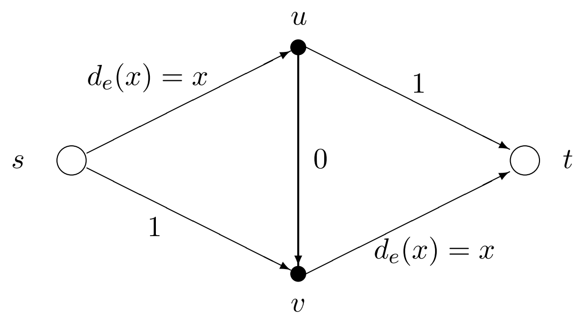

The cost of a flow in Nash equilibrium is evaluated as the delay on a path from s to t (recall that at Nash equilibrium, all $s-t$ paths have the same delay). In the example, if the edge uv is not present, then the Nash flow is 1/2 on each of the two paths, and the resulting delay is 1.5. However, if the edge uv with zero delay is added to the graph, then the equilibrium flow is 1 on the path $s-u-v-t$, and the delay is 2. Braess' paradox is the fact that the seemingly beneficial action of adding a new edge in fact only increases the delay.

Before, we had a function representing social cost, $\sum_p f_p d_p(f)$, where f_p is flow on a path p, and we studied the "price of anarchy", which is $\frac{\textit{social cost of Nash}}{\textit{opt social cost}}$. Now we will focus on Braess' paradox and study the ratio $\frac{\textit{social cost of Nash}}{\textit{social cost of Nash on best subgraph}}$.

The first observation is that the ratio for Braess' paradox is at most the price of anarchy. This is because routing the flow on a subgraph is one of the things that *opt* can do, so the optimal delay is not more than the best delay on a subgraph. For example, when delays are linear, $d_e(x) = a_e x + b_e$, since the price of anarchy is at most 4/3, so is the Braess' paradox ratio.

Let us try to construct an example with high Braess' paradox ratio. The bad example that we had for the price of anarchy was a two-edge two-node graph with delays x^N and 1 on the two edges, and a traffic rate of 1. In Nash equilibrium, all the flow goes on the first edge, with a delay of 1, whereas the optimum is to send ϵ flow on the edge with delay 1, and the rest on the other edge, for an average delay of $\epsilon \cdot 1 + (1-\epsilon)(1-\epsilon)^N \approx \epsilon$. However, there is no Braess' paradox on this graph (the ratio is 1).

As a second attempt to construct an example with high ratio, let us change the delay function on the variable-delay edges in the picture above from x to x^N. As a result, when the middle edge is absent, the delay at Nash is close to 1, whereas when it is present, the delay is 2.

Here are some facts which are known about Braess' paradox:

1. For a single source-sink pair, the worst Braess' paradox ratio is $n/2$, where $n = |V|$ (Roughgarden 2001).
2. It's NP-hard to find a good subgraph, even if the class of examples is guaranteed to only contain networks with Braess' paradox ratio which is either $n/2$ or 1 (Roughgarden 2001).
3. If the graph contains 2 $s-t$ pairs, then the Braess' paradox can be as bad as $\sim 2^{O(n)}$.

Now we will prove an upper bound of n for the first fact, and will show an example illustrating the third fact. We will consider a single source - single sink network and look at the $s-t$ delay on it.

Theorem 1 *For flow with a single source-sink pair, the Braess' paradox ratio is at most n.*

Let us say that f is a Nash flow that has delay D, and f^* is Nash flow on the best subgraph, with a delay of D^*. For some edge e, if $f^*(e) > 0$ and $f(e) \leq f^*(e)$, then we will call such an edge *good*, which means that it is not deleted and is part of the best subgraph.

Claim 2 *If e is good, then $d_e(f) \leq D^*$.*
Proof.

$$d_e(f) \leq d_e(f^*) \leq D^*$$

where the first inequality follows from monotonicity of d_e, and the second one follows because $f^*(e) > 0$, which means that e is part of some flow path in f^*, and therefore its delay is at most that of the whole path. ■

Note: If there is an $s-t$ path of good edges, then $D \leq (n-1)D^*$.

Looking at the example in the figure, we see that in this case only the edges sv and ut are good. So not every graph contains such a path of good edges. However, we will use something similar in order to establish the theorem.

Claim 3 *Nash flow is acyclic (i.e., edges with $f(e) > 0$ do not form cycles.*
Proof sketch: Nash flow is optimal for another objective function ($c_e(x) = \int_0^x d_e(y)dy$), and decreasing flow around a cycle decreases this cost function. ■

Suppose that we sort the nodes $s, v_1, v_2, ..., t$ in the order of delay on the path from s to v, so that $d(v_i) \leq d(v_{i+1})$. Notice that if $f(ij) > 0$, then $d(i) \leq d(j)$. As a result, since the graph of edges that carry flow f is acyclic, all edges that carry flow will go from earlier nodes to later nodes in this ordering. Our example would become:

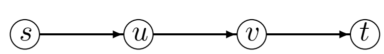

(the edges shown are the ones carrying flow f).

Claim 4 $d(v_{i+1}) \leq d(v_i) + D^*$ *for all i.*
Notice that this implies the claimed bound of $D \leq (n-1)D^*$.
Proof. Let $S = \{s, v_1, v_2, ..., v_i\}$ be some prefix of the ordering. We will prove that there exists a good edge uv from S to $\bar{S}$. This suffices to prove the claim because $d(v_{i+1}) - d(v_i) \leq d(v) - d(u) \leq D^*$

by Claim 2. If r is the total $s-t$ flow, then, using the fact that edges with $f(e) > 0$ all go from left to right, we know that the flow f out of S is exactly r. The flow f^* out of S is at least as much, even if there are paths that reenter S. So there must be some edge for which $f^*(uv) \geq f(uv)$, which proves the claim. ■

We have shown that, although there may not be a path of good edges from s to t, there is a good edge crossing every cut in the ordered list of nodes. In the example, these are edges sv and ut.

2 Bad example for two commodities

The following figure shows that if the graph contains two source-sink pairs, then the ratio of Braess' can be exponential in n.

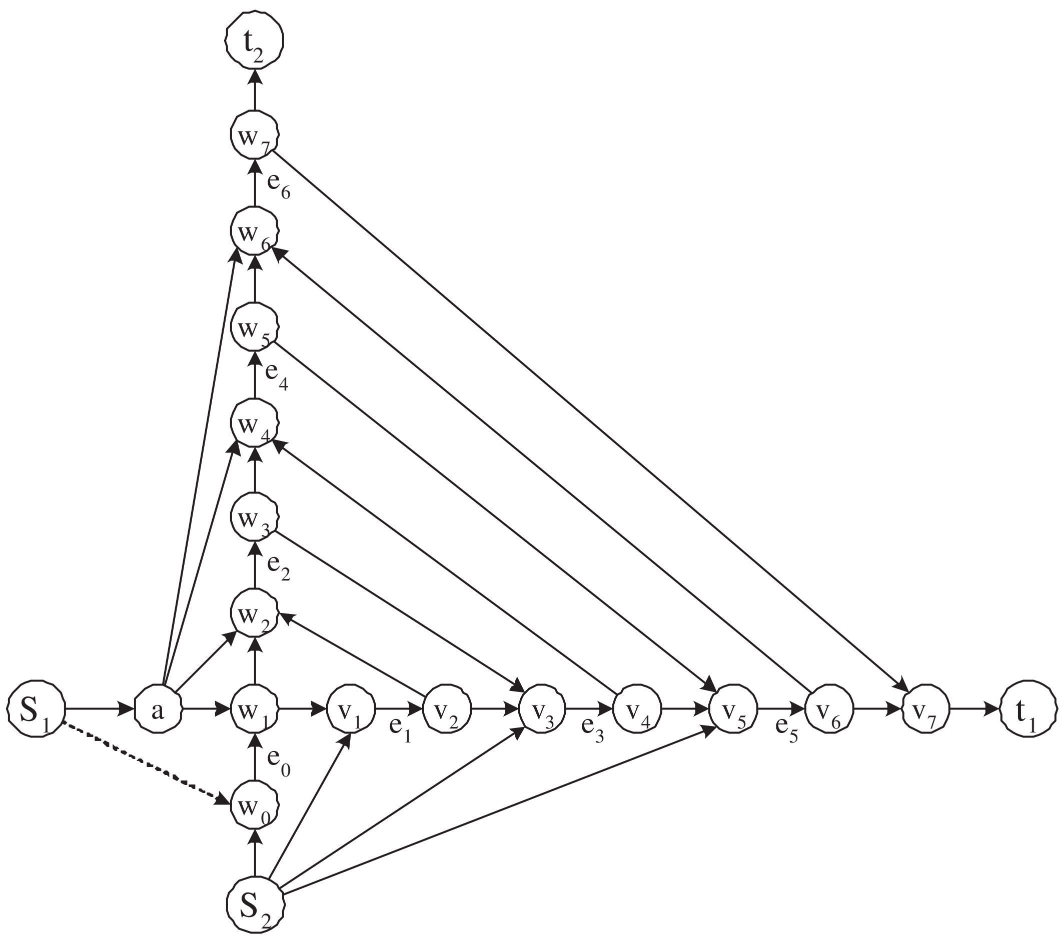

Figure 1: The users taking diagonal “shortcuts” results in delays that grow as a Fibonacci sequence.