DS 574 Algorithmic Mechanism Design	**Congestion Pricing Worksheet**
Fall 2026	Prof. Kira Goldner

Selfish Routing

You're driving from Boston to Cambridge and have two routes to choose from. The highway takes one hour regardless of how many drivers use it. The bridge is faster when it is uncrowded, but slows down as more drivers choose it: if a fraction x of the traffic takes the bridge, crossing takes x hours. We normalize the total traffic to one unit, so a fraction x is also an amount of traffic x. Each driver wants to get to Cambridge as quickly as possible.

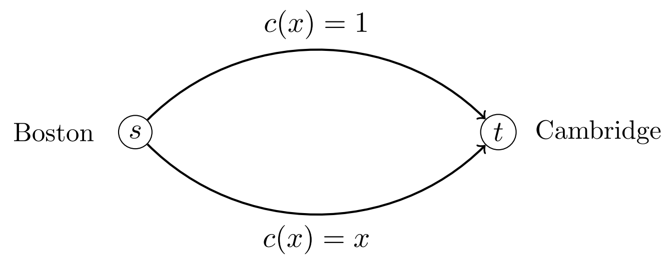

When we defined Nash equilibria, we asked whether any player could benefit by changing their own action. Here a driver's action is a route. An equilibrium means that no driver can get to Cambridge faster by switching routes.

We model a large population in which each driver is an infinitesimal part of the traffic. One driver switching routes does not change either route's travel time. This is called *nonatomic* traffic. Thus a driver simply compares the travel times on the available routes at the current traffic levels.

A driver's objective. Each driver chooses their own route to minimize their own travel time, given the routes everyone else is taking. If fraction x is on the bridge, the driver compares x hours on the bridge with 1 hour on the highway. They choose a route, not the overall traffic split x. Their objective does not include the time other drivers spend on the road.

Definition 1. An *equilibrium flow* is a routing of traffic in which no driver can reduce their travel time by changing routes. Every route that is used must be a fastest available route: used routes have equal travel times, and an unused route cannot be faster.

A *flow* just describes how the traffic is split among routes. In this example, putting x on the bridge leaves $1 - x$ on the highway. Write $C(f)$ for the total time spent driving under a flow f.

Pigou's example. Where will drivers go in equilibrium?

Claim 1. In Pigou's example, the unique equilibrium flow routes all of the traffic onto the bottom (bridge) edge. Every driver's travel time is 1, and $C(f) = 1$.

Proof.

The designer's objective. The designer chooses how to split the traffic among routes to minimize the total time everyone spends driving. Unlike an individual driver, the designer can reroute a positive fraction of the traffic and accounts for how that changes travel times for everyone. An *optimal flow* is a traffic split that minimizes this total, whether or not drivers would choose to follow the assignment.

Write $C(f)$ for the *total travel time*, or social cost, of a flow f. To compute it, multiply the amount of traffic on each road by the time spent on that road, then add. Since total traffic is one unit, total and average travel time coincide.

In the highway-and-bridge example, the designer chooses $x \in [0, 1]$, the fraction assigned to the bridge, to minimize

$$C(x) = \underbrace{x \cdot x}_{\text{time spent by bridge drivers}} + \underbrace{(1 - x) \cdot 1}_{\text{time spent by highway drivers}}.$$

Sending fewer drivers over the bridge makes it faster for those who remain, but requires more drivers to spend an hour on the highway.

Claim 2. The optimal flow in Pigou's example is

Proof.

How much worse can selfish routing be than the best assignment of routes? We measure this using the price of anarchy.

Definition 2 (Koutsoupias–Papadimitriou '99)**.** The *price of anarchy* of an instance is

$$\text{PoA} = \frac{C(\text{worst equilibrium flow})}{C(\text{optimal flow})}.$$

What is the price of anarchy?

Braess's Paradox

Could adding a road make traffic worse? A planner could always leave the new road unused, so the best possible routing cannot get worse. But drivers choose their own routes. A new shortcut may change what they choose.

Consider two routes, each with one road whose travel time equals its traffic and one road that takes an hour. We still have one unit of traffic. First find the equilibrium without the middle road; then add a zero-cost road from v to w.

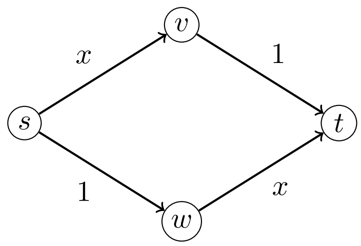

Before

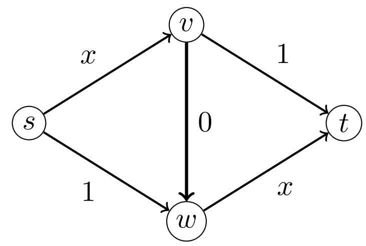

After

On a network, a route may use several roads. Write f_e for the traffic on road e and $c_e(f_e)$ for its travel time. The cost $c_P(f)$ of a route P is the sum of its roads' travel times. Total travel time is

$$C(f) = \sum_e \underbrace{f_e}_{\text{traffic on road } e} \underbrace{c_e(f_e)}_{\text{time per driver on } e} .$$

A driver using several roads contributes the time spent on each; adding road by road counts that driver's whole trip.

The highlighted roads show the three available routes. We use P_1, P_2, P_3 to refer to them in the calculations.

P_1: upper route

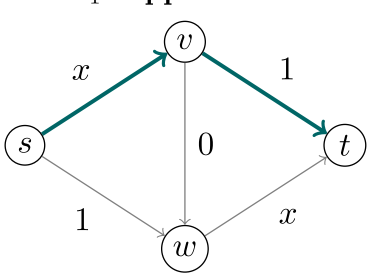

P_2: lower route

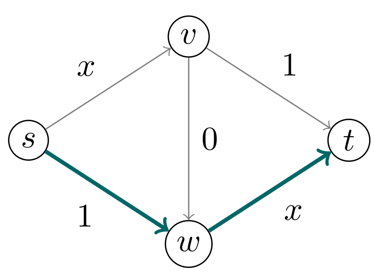

P_3: shortcut

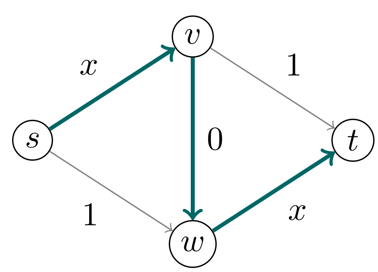

Claim 3. In the "before" network, the unique equilibrium flow is

Proof.

Theorem 1 (Braess's paradox)**.** *In the "after" network, the unique equilibrium flow is*

Proof.

Why is the old equilibrium no longer an equilibrium?

Claim 4. The optimal flow in the "after" network has total cost

What is the price of anarchy after adding the road?

Both examples gave a price of anarchy of . Could a larger network be worse, even if each road still has a linear travel-time function plus a constant?

Theorem 2 (Roughgarden–Tardos '02)**.** *Consider any selfish routing instance—any directed graph, any number of source-sink pairs, any traffic rates—in which every cost function is* affine*, i.e. of the form $c_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$. Then PoA $\leq 4/3$. Pigou's example shows this bound is tight.*

The general affine bound is stated without proof. Pigou and Braess attain it; other affine instances can have a smaller price of anarchy.

Interventions

In the optimal split for the highway and bridge, drivers on the highway would rather take the bridge. How could we make them willing to stay on the highway?