MS&E 246: Lecture 18
Network routing

Ramesh Johari

Network routing

- Last lecture: a model where N is finite
- Now: assume N is very large
- Formally: Represent the set of users as a *continuous interval*, $[0, B]$
- B represents the total amount of flow
- Each user is *infinitesimal*

  (also called *nonatomic)*

Network routing

We'll consider networks of *parallel links.*

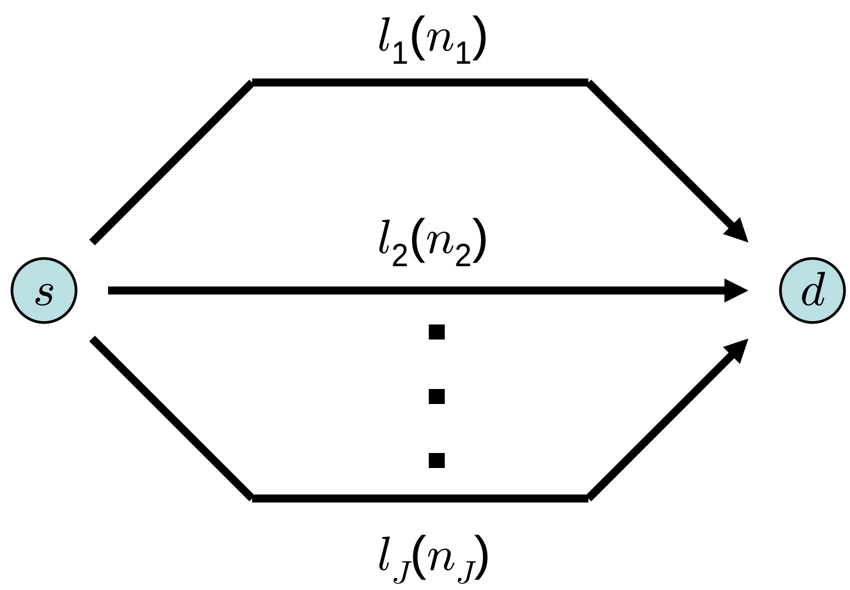

J = # of links; n_j = flow on link j

The network routing game

Strategy space of each user:

Links from s to d

Cost to each user:

Delay experienced on chosen link

The network routing game

Formally:

Delay on link j = $l_j(n_j) \geq 0$

(assume l_j is strictly increasing)

Strategy space = { 1, ..., J }

Payoff = $l_j(n_j)$ if link j chosen

The network routing game

What is a pure NE in this game?

Each infinitesimal user must have chosen the best path available, given links chosen by all others.

Formally: If $n_j > 0$, then $l_j(n_j) \leq l_k(n_k)$ for all $k \neq j$

Wardrop equilibrium

A *Wardrop equilibrium* $\mathbf{n}$ is a pure NE of this game:

(1) $\sum_j n_j = B$

(2) If $n_j > 0$, then $l_j(n_j) \leq l_k(n_k)$ for all $k \neq j$

Does a WE exist? Is it unique?

Wardrop equilibrium

Recall last lecture: we found a *potential*

$$V(\mathbf{p}) = \sum_{j \in J} \sum_{i=0}^{n_j(\mathbf{p})} l_j(i)$$

Intuition: replace sum by an integral.

Wardrop equilibrium

Theorem:

There exists a unique WE $\mathbf{n}$. It is the unique solution to:

$$\text{minimize} \quad \sum_{j=1}^{J} \int_0^{n_j} l_j(z)\,dz$$

$$\text{subject to} \quad \sum_{j=1}^{J} n_j = B$$

Wardrop equilibrium: proof

Define $V(\mathbf{n}) = \sum_{j=1}^{J} \int_0^{n_j} l_j(z)\, dz$

Suppose $\mathbf{n}$ is optimal, and $n_j > 0$.

Consider moving $\delta > 0$ units of flow to link $k \neq j$.

What is the change in V?

Wardrop equilibrium: proof

- The j'th term falls by: $l_j(n_j)\ \delta$
- The k'th term increases by: $l_k(n_k)\ \delta$

Since $\mathbf{n}$ is optimal, this change cannot reduce the cost.

So:

$$l_k(n_k)\ \delta - l_j(n_j)\ \delta \geq 0$$

Wardrop equilibrium: proof

- The j'th term falls by: $l_j(n_j)\ \delta$
- The k'th term increases by: $l_k(n_k)\ \delta$

Since $\mathbf{n}$ is optimal, this change cannot reduce the cost.

So:

$l_j(n_j) \leq l_k(n_k)$ for all $k \neq j$

Wardrop equilibrium: proof

Conclude:

Any optimal solution is a Wardrop equilibrium.

At least one optimal solution exists

(the feasible set is closed and bounded).

So at least one WE exists.

Wardrop equilibrium: proof

To check uniqueness:

Since l_j is strictly increasing,
each term of V is *strictly convex* in n_j
(strictly convex
$\Leftrightarrow$ for all $\alpha \in [0,1]$, x, x',
$f(\alpha x + (1 - \alpha) x') < \alpha f(x) + (1 - \alpha) f(x')$
$\Leftrightarrow$ second derivative is strictly positive)

Wardrop equilibrium: proof

Since l_j is strictly increasing,
each term of V is *strictly convex* in n_j

First implication:

Any WE is also an optimal solution.

(Second order conditions
automatically hold.)

Wardrop equilibrium: proof

Since l_j is strictly increasing,
each term of V is *strictly convex* in n_j

Second implication:

If $\mathbf{n} \neq \mathbf{n}$' are two solutions, then
($\mathbf{n}$ + $\mathbf{n}$')/2 is a feasible solution with
strictly lower cost.

So optimal solution (and hence WE)
is unique.

Total delay

We again expect Wardrop equilibria will not be Pareto efficient.

One way to find a Pareto efficient point:

minimize total delay.

Total delay = TD($\mathbf{n}$) = $\sum_{j=1}^{J} n_j \, l_j(n_j)$

Total delay

Let's assume l_j is affine for each j, i.e.,

$l_j(n_j) = a_j + b_j\, n_j, \quad a_j \geq 0,\ b_j \geq 0$

Then:

$\mathrm{TD}(\mathbf{n}) = \sum_{j=1}^{J} a_j\, n_j + b_j\, n_j^2$

Total delay

Let's assume l_j is affine for each j, i.e.,

$l_j(n_j) = a_j + b_j\, n_j, \quad a_j \geq 0,\ b_j \geq 0$

Then:

$\mathrm{TD}(\mathbf{n}) = \sum_{j=1}^{J} \int_0^{n_j} (a_j + 2\, b_j\, z)\, dz$

Total delay

Let's assume l_j is affine for each j, i.e.,

$l_j(n_j) = a_j + b_j\, n_j, \quad a_j \geq 0,\ b_j \geq 0$

Then:

$\mathrm{TD}(\mathbf{n}) = \sum_{j=1}^{J} \int_0^{n_j} (a_j + 2\, b_j\, z)\, \mathrm{d}z$

So a flow allocation that minimizes delay is a WE with respect to *different latencies:*

$m_j(n_j) = a_j + 2\, b_j\, n_j$

Pigovian taxes

In general, WE does not minimize total delay.

Pigovian taxes charge tolls to provide the right incentives to users.

(Named for the economist Pigou.)

Pigovian taxes

Suppose on link j,

a *toll* is charged, $T_j(n_j)$.

Also suppose that users value *time* and *money* identically.

Then a user of link j has payoff:

$l_j(n_j) + T_j(n_j)$

Pigovian taxes

Notice that if we choose:

$T_j(n_j) = b_j\, n_j$,

then

$l_j(n_j) + T_j(n_j) = m_j(n_j)$

Moral: A Nash equilibrium with these tolls minimizes total delay.

Pigovian taxes

More generally:

If l_j is *strictly convex* and strictly increasing,

the *Pigovian tax* on link j is:

$$T_j(n_j) = n_j\, l_j'(n_j)$$

Efficiency: example

Example:

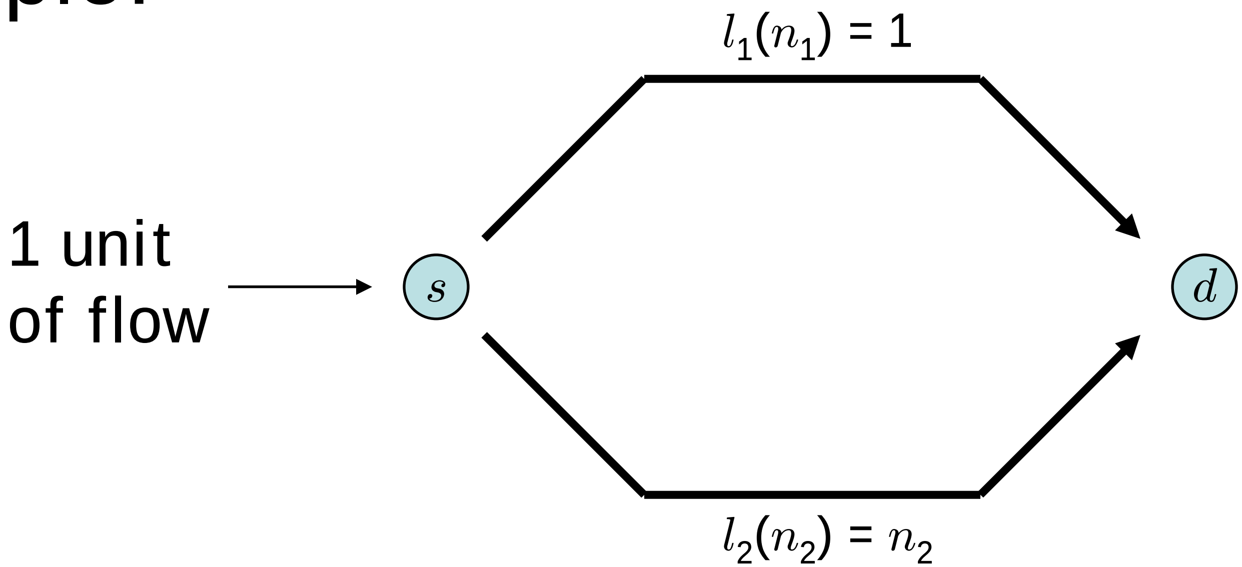

Efficiency: example

Wardrop equilibrium:

$n_1 = 0$, $n_2 = 1 \Rightarrow$

Total delay = 1

Minimizing total delay:

$n_1 = 1/2$, $n_2 = 1/2$, $\Rightarrow$

Total delay = 3/4

So WE cost is 4/3 higher than the optimum.

Efficiency

When l_j are affine,
the 1/3 increase in total delay is the worst possible (over all choices of l_j)

(Roughgarden and Tardos, 2002)

Other directions

More complex user models:

1. Users may not be infinitesimal
2. Users may have different preferences for money and time
3. Users may care about the variance of their delay

All these are much harder than WE.

Other directions

Partial optimization in the network:

ISPs frequently reroute traffic inside their own networks to improve performance.

How does this interact with selfish routing from source to destination?

Other directions

An ISP might set prices to maximize profit, not achieve minimum total delay.

How efficient is profit maximization?

Other directions

A service provider may wish to invest to improve the performance of existing links.

What is the equilibrium of the resulting game of investment, pricing, and routing?

Other directions

Service providers form contracts with each other to share traffic (e.g., peering).

What contracts should they use to:

-maximize profit?

-maximize network performance?