How much Can Taxes Help Selfish Routing

Ακαδημαϊκό Έτος 2010-2011

Παπαλεξίδης Νίκος

Introduction

- Selfish behavior in networks
- Nash equilibrium
- Pricing network edges – impose taxes
- Total cost (disutility) : Latency + Taxes
- Benefits of taxes on networks and complexity of computing optimal taxes

Nash equilibrium - flows

- Selfish routing – Each user routes itself on minimum-latency path, given the network congestion due to other users
- Nash equilibrium : Stable point in which no user has an incentive to unilaterally alter its strategy
- Nash flow : All traffic is routed on paths with minimum-possible latency
- Nash equilibrium is not optimal : Latency is not minimized
- Marginal taxes
- Total cost = Latency + Taxes

Routing model

- Traffic from source *s* to sink *t* in a graph *G(V,E)*
- Latency function $l_e(f)$ for each edge *e* *(function of the flow f)*
- Example :

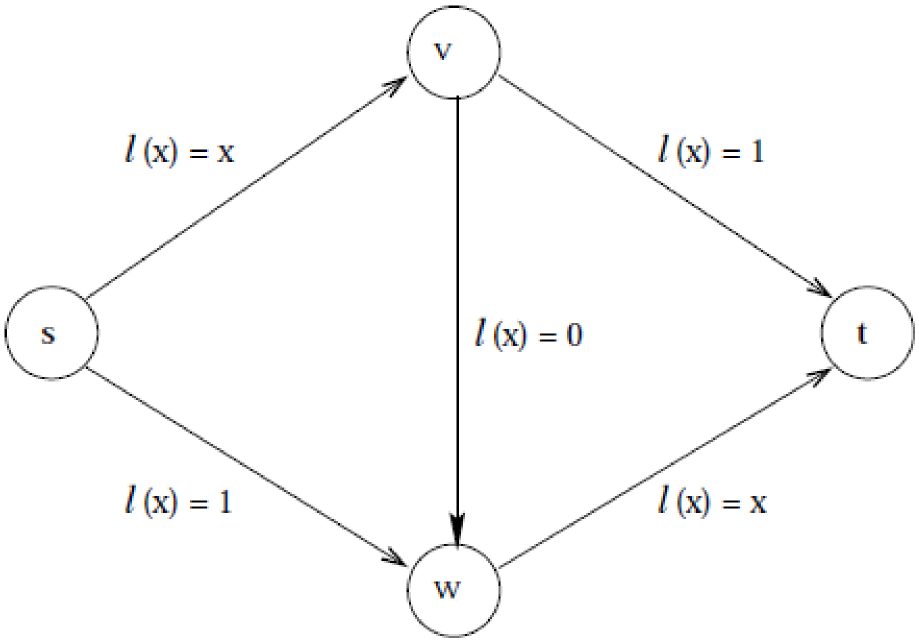

Braess paradox

Initial network

<u>delay = *1.5*</u>

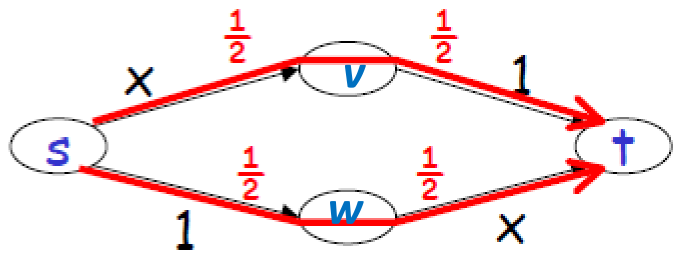

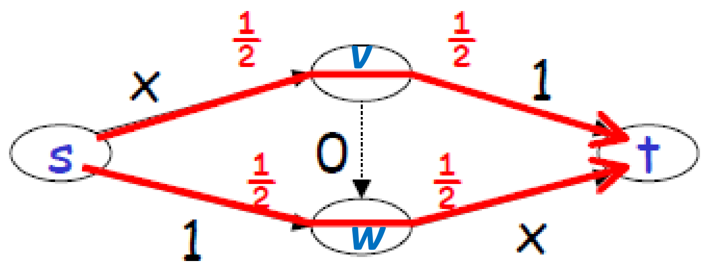

Extra edge *v-w* added

<u>delay = *2.0*</u>

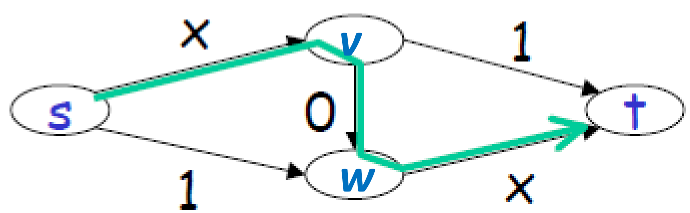

Improve situation with taxes

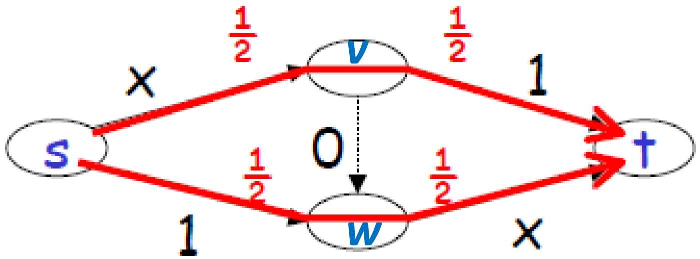

Tax = ½ on edge *v-w*

delay = *1.5*

tax + delay = *1.5*

Result :

No taxes paid

Latency improved (*1.5* from *2.0*)

Total cost = Latency + Taxes

Marginal taxes

Marginal cost pricing : Each user should pay a tax equal to the additional delay other users experience because of his presence

Marginal taxes on edges
s-v and *x-t* are ½ and
0 on other edges
<u>delay = *1.5*</u>

but :
<u>tax + delay = *2.0*</u>

Latency improved
Total cost did not improve

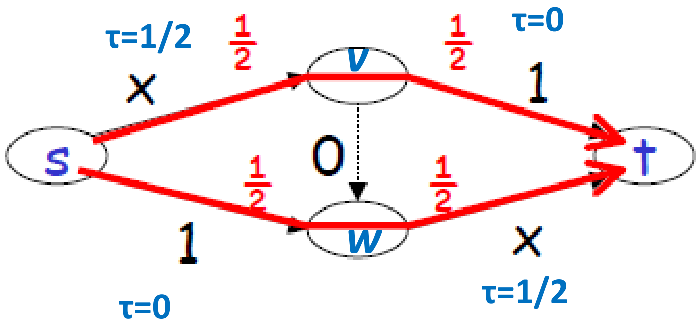

Problem - Questions

- Goal : Decrease the total cost (Latency + Taxes) using taxes

- Questions studied :
  - Are marginal taxes a good idea for minimizing Nash equilibrium?
  - Compare the efficiency of taxes with that of edge removal (note that a large edge tax removes the edge from the network)
  - Compute or approximate efficiently the optimal taxes

Model - definitions

- Simple Path P from s to t
- Flow f_p on each path P
- Flow f_e on edge e
- r : traffic rate
- $l_e(f_e)$: latency function on edge e
- $l_p(f)$: latency of a path P with respect to flow f
- $L(f)$: total latency $L(f) \equiv \sum_{P} l_P(f) f_P = \sum_{e \in E} l_e(f_e) f_e$
- Optimal or minimum-latency flow : minimizing $L(f)$
- τ_e : tax on edge e
- τ_P : total taxes on a path P
- $C(f,\tau)$: Cost of a flow f (Latency + Taxes) $C(f,\tau) \equiv \sum_{P} [l_P(f) + \tau_P] f_P = \sum_{e \in E} [l_e(f_e) + \tau_e] f_e$
- (G,r,l) instance of G
- $(G,r,l+\tau)$: instance with taxes

Flows at Nash equilibrium

- Proposition (2.6) : If f is at Nash equilibrium for $(G,r,l+\tau)$, then there is a constant $c \geq 0$ with $l_p(f) + \tau_P = c$. Moreover $C(f,\tau) = r \cdot c$ (all paths have the same latency +tax)

- Marginal tax τ_e :
$$\tau_e = f_e \cdot l'_e(f_e)$$

- **Proposition (3.1)** : (G,r,l) instance with latency function admitting minimum-latency flow f^*. If τ_e is marginal cost tax for edge e, f^* is at Nash equilibrium for $(G,r,l+\tau)$

  Meaning : Marginal taxes induce the minimum-latency flow as a flow at Nash equilibrium

- Effective way to minimize the total latency of a Nash flow with edge taxes

- How effective are marginal taxes if we account the total cost (latency+taxes)?

When do marginal costs help?

- Theorem 3.2 : *(G,r,l)* instance with linear latency function. Let f and f^{T} be Nash flows for *(G,r,l)* and *(G,r,l+τ)* respectively.

Then $C(f,0) \leq C(f^{T}, \tau)$

  - Meaning : Linear marginal taxes can not improve total cost
  - Same result for latency function $\alpha_e x^p + b_e$ with $\alpha_e, b_e \geq 0$ (fixed p)

Effectiveness of arbitrary taxes – Upper bounds

- *Price of anarchy (PA)*: Largest possible ratio between the total latency of a Nash flow and that of a minimum latency flow.
- Can be used as upper bound of the maximum-possible reduction in cost due to taxes.

- Prop. 4.1 : Linear latency functions, *PA=4/3* (see example Braess Paradox)
- Prop. 4.2 : Latency functions polynomial with degree at most *p* , *PA → Θ(p/logp)*

- Theorem 4.5 : *(G,r,l)* and *(G,r,l+τ)* instances with *f* and f^{τ} Nash flows respectively. Then

$$L(f) \leq \left\lfloor \frac{n}{2} \right\rfloor \cdot C(f^{\tau}, \tau)$$

Comparing Taxes with Edge Removal

- <u>Networks with linear latency functions</u>

- Theorem 5.1 : An instance with linear latency functions admits an optimal set of taxes that is $0/\infty$

Meaning : - Taxes in linear latency networks are equivalent with edge removal with respect to the maximum reduction of the Nash flow (= *4/3*)

- Taxes in these networks can not improve the Nash flow more than the removal of some edges

Comparing Taxes with Edge Removal

- <u>Networks with general latency functions</u>

- Theorem 5.2 : For each integer *n≥2*, there is an instance *(G,r,l)* with $c(H,r,l)=\left\lfloor \frac{n}{2} \right\rfloor$ for all subgraphs *H* of *G* but $c(H,r,l+\tau)=1$ for some tax *τ ≥ 0*.

Meaning : - Taxes in general latency networks can improve the Nash flow by a $\left\lfloor \frac{n}{2} \right\rfloor$ factor beyond what is achievable by removing edges.

- Removing edges can improve the Nash flow by a $\left\lfloor \frac{n}{2} \right\rfloor$ factor ([1])
- Combined taxes+edge removal cannot improve more than $\left\lfloor \frac{n}{2} \right\rfloor$ (due to Theorem 4.5)

Comparing Taxes with Edge Removal

Linear latency

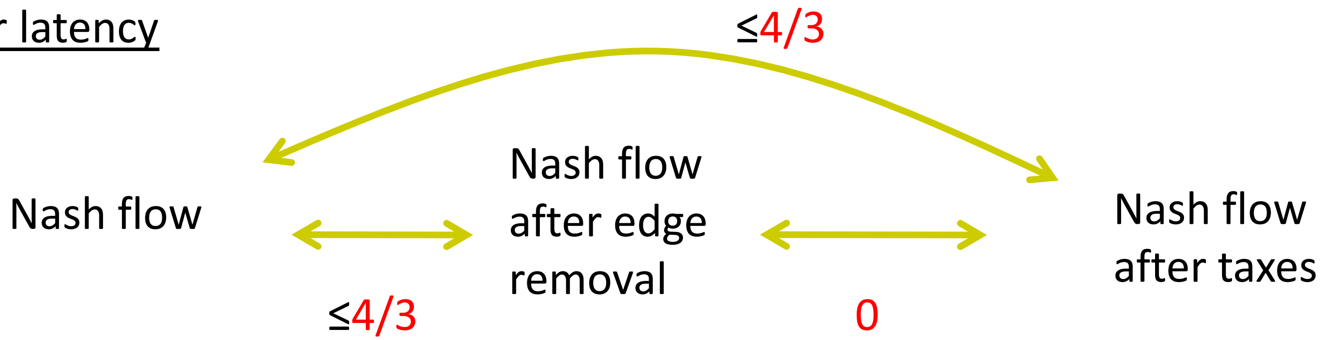

General latency

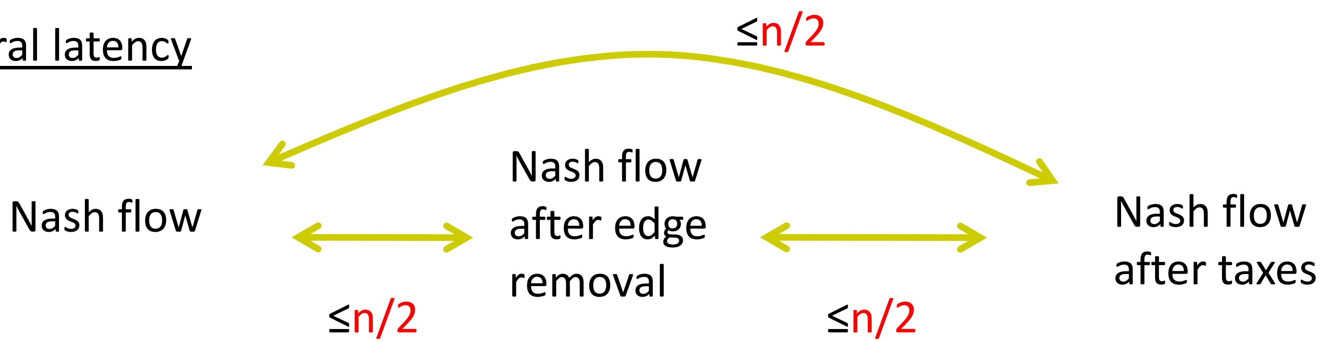

Comparing Taxes with Edge Removal

- Examples : Braess Graphs

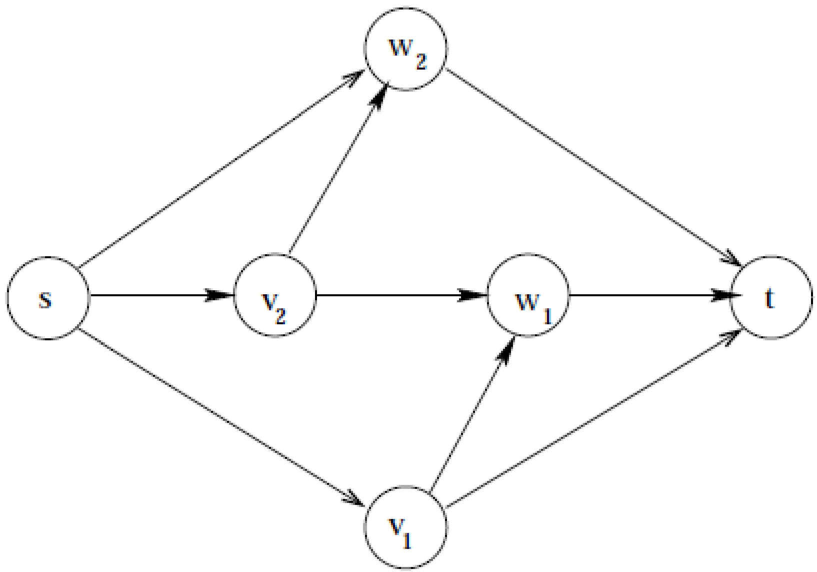

(a) B^2

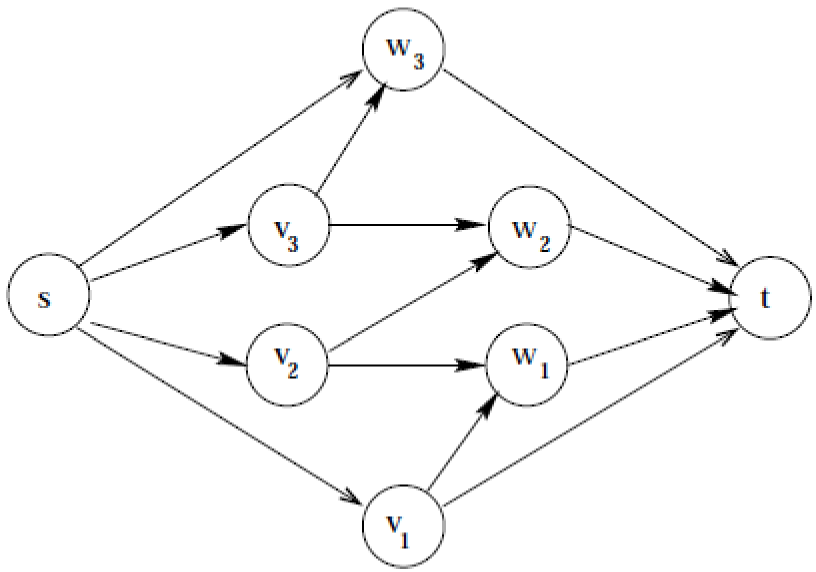

(b) B^3

Complexity of Computing Optimal Taxes

Trivial algorithm : Assign all edges zero taxes

Approximation factor

- Linear latency : *4/3*
- Polynomial latency functions with degree *p : Θ(p/logp)*
- General latency functions : $\left\lfloor \frac{n}{2} \right\rfloor$
- Theorem 6.2 : Unless *P=NP*, (*ε>0*), the problem of computing optimal taxes has no approximation algorithm with factor
  - *(4/3-ε)* for linear latency networks
  - *o(p/logp)* for polynomial with degree p networks
  - $O(n^{1-\varepsilon})$ for general networks

Conclusion

Problem Studied	Linear Latency function	General latency function
Can marginal taxes help?	No	Yes
Maximum benefit of taxes	4/3	n/2
Taxes better than network design (edge removal)	No	Yes
Approximability of optimal taxes	4/3	$O(n^{1-\varepsilon})$

References

- [1] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science, pages 472–481,* 2001.
- [2] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM, 49(2):236–259, 2002.* Preliminary version in *FOCS ’00.*
- [3] R.Cole, Y.Dodis and T.Roughgarden. How Much Can Taxes Help Selfish Routing? EC ‘03, pages 98-107.