When does diversity of user preferences improve outcomes in selfish routing?

Richard Cole
NYU

Thanasis Lianeas
UT Austin / NTUA

Evdokia Nikolova
UT Austin

Goal

Understand effect of **user diversity** on congestion, by studying resulting traffic assignment:

- Traffic congestion: many users choose same route
- Compare equilibrium **cost** of heterogeneous (diverse) user population to that of **comparable** homogeneous user population

Motivating example 1: risk-aversion

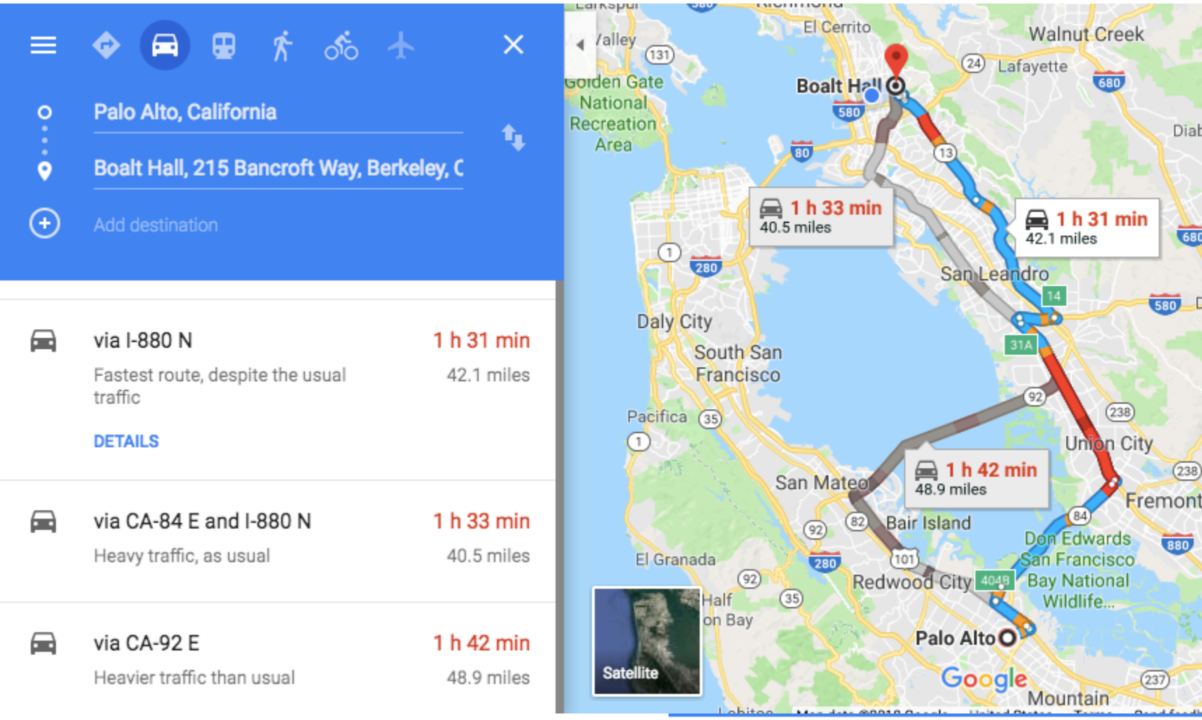

Users trade-off mean and variance of travel time

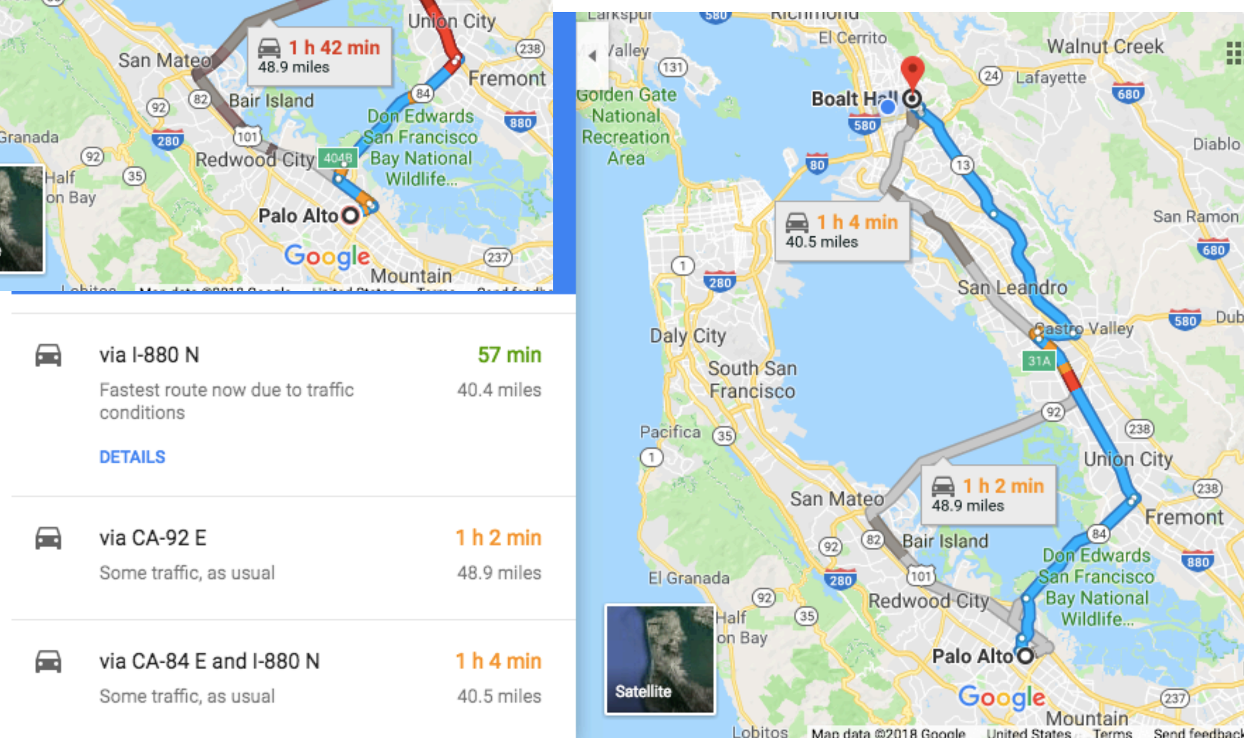

Motivating example 2: tolls

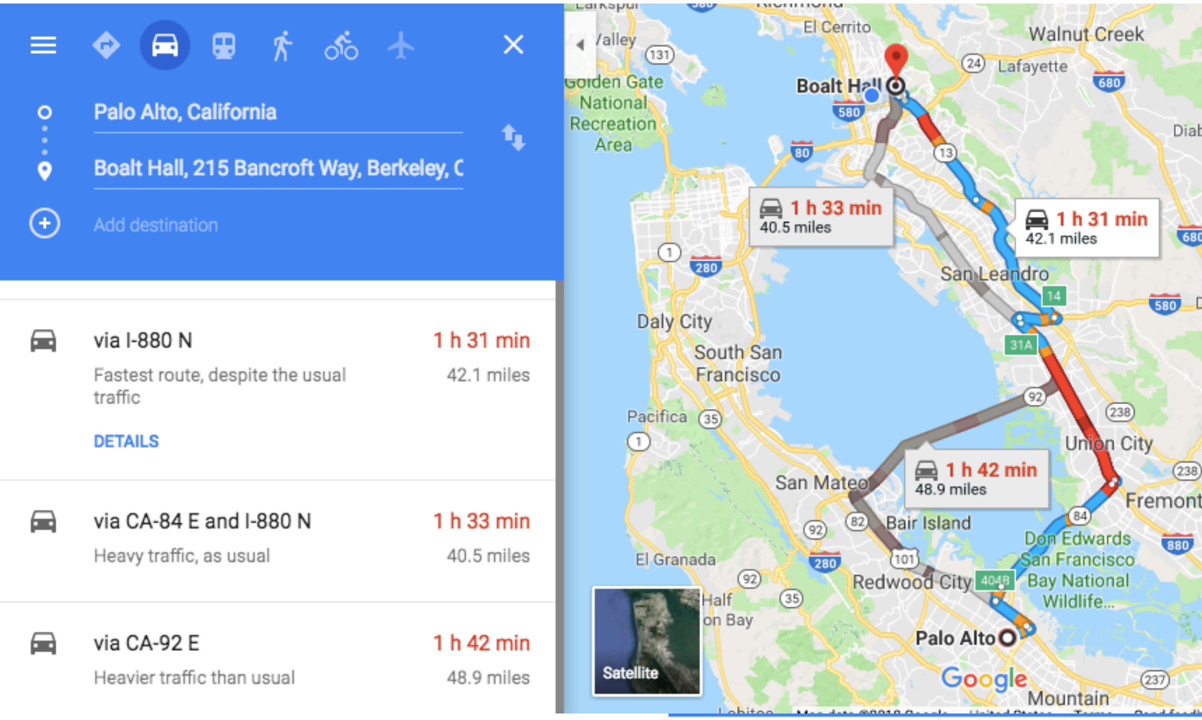

Users trade-off time and cost (tolls paid)

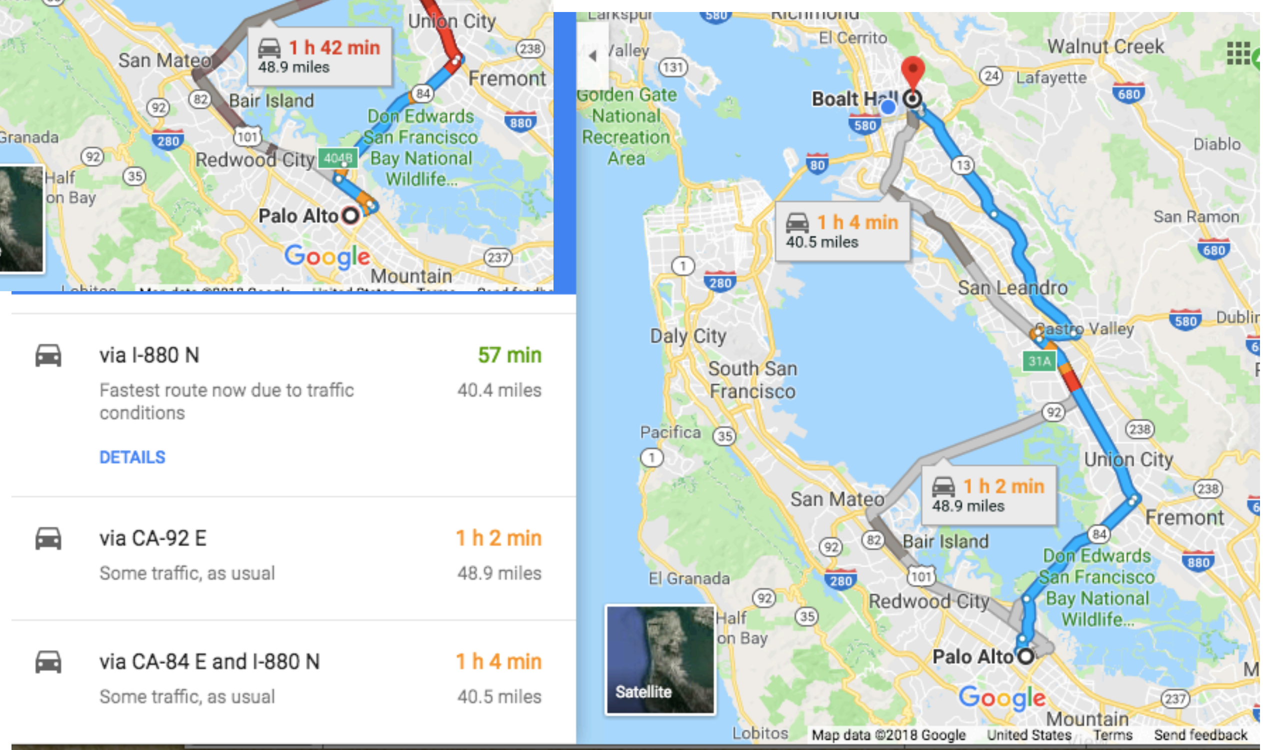

Evdokia Nikolova

Overview of results

- Does heterogeneity (diversity) of users reduce the cost of equilibrium? Users min (delay + r_i cost)
- Diversity helps if and only if the network is series-parallel for single origin-destination.
- Diversity helps if and only if the network is "block-matched" for multiple origin-destination pairs.

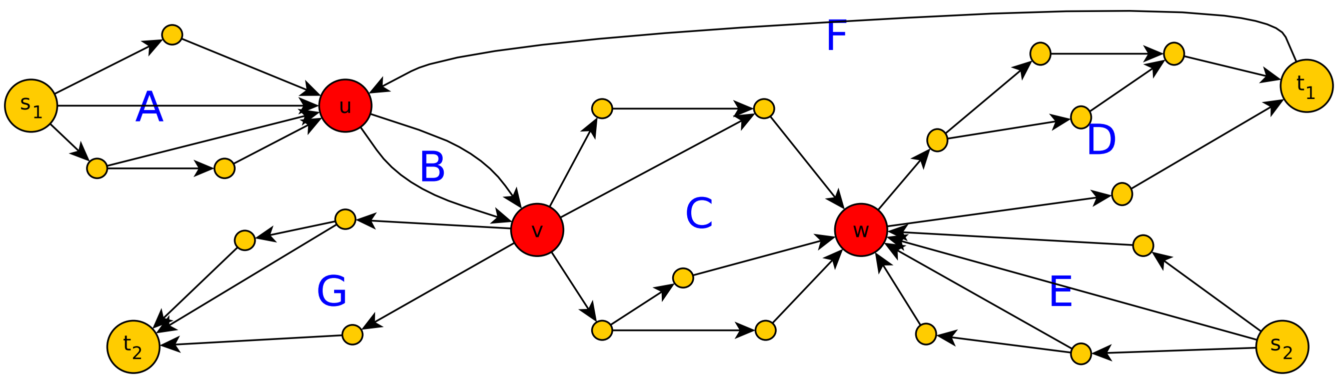

R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.

Model

- Directed graph $G = (V,E)$, multiple source-dest. pairs (s_k,t_k) with demand d_k (call this commodity k)
- Nonatomic players (*flow model)* choose feasible s-t paths
  Players' decisions: flow vector $x \in R^{|Paths|}$
- Edge delay $l_e(x_e)$ and "deviation" (toll) functions $\sigma_e(x_e)$
- Different player types tradeoff delay and deviation differently via diversity parameter r
- Players minimize delay plus deviation:

$$c^r_{path}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e) = \sum_{e \in path} \left(l_e(x_e) + r\sigma_e(x_e) \right)$$

Cost of flow

- Players minimize delay plus deviation:

$$c^{r}_{path}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e)$$

- What should be the cost of flow x?

1) Sum of first criterion only

   - In toll literature, cost is total travel time only

   - In risk-averse routing, cost is average travel time (meaningful for social planner who cares about long-term averages)

2) Total user cost (sum of both criteria)

   - Consistent with traditional definition of “social welfare” in Economics

- Both 1) and 2) are meaningful depending on application

Questions

- Players minimize delay plus deviation:

$$c^r_{path}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e)$$

- Two natural questions:

I. How does equilibrium cost of population with parameter r compare to equilibrium cost of population with parameter 0? (e.g., risk-averse vs risk-neutral people, or people who care about both time and money vs those who only care about time)

II. How does equilibrium cost of population with distribution of parameters D(r) compare to equilibrium cost of population with same average parameter $\bar{r}$?

Questions

- Players minimize delay plus deviation:

$$c_{path}^{r}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e)$$

- Two natural questions:

I. How does equilibrium cost of population with parameter r compare to equilibrium cost of population with parameter 0? (e.g., risk-averse vs risk-neutral people, or people who care about both time and money vs those who only care about time): We answer for flow cost = first criterion only *

II. How does equilibrium cost of population with distribution or parameters D(r) compare to equilibrium cost of population with same average parameter $\bar{r}$? We answer for cost = total user cost **

* T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

** R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.

Equilibrium definition

- Users select paths with minimum cost $c^r_{path}(x)$
- Definition: A flow x is at equilibrium if for every source-destination pair k and for every *path* with positive flow

$$c^r_{path}(x) \le c^r_{path'}(x), \quad \text{for every } \; path' \text{ and player type } r$$

- We call it a *heterogeneous equilibrium g* if there are different player types (with different r's)
- We call it a *homogeneous equilibrium f* if there is a single player type (same r)

Questions

- Players minimize delay plus deviation:

$$c^r_{path}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e)$$

- Two natural questions:

I. How does equilibrium cost of population with parameter r compare to equilibrium cost of population with parameter 0? (e.g., risk-averse vs risk-neutral people, or people who care about both time and money vs those who only care about time): We answer for flow cost = first criterion only *

II. How does equilibrium cost of population with distribution or parameters D(r) compare to equilibrium cost of population with same average parameter $\bar{r}$? We answer for cost = total user cost **

* T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

** R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.

Comparing equilibria with parameter r vs 0

Cost of Flow $C(x)$: sum of first criterion only

- e.g., although users are risk-averse, central planner is risk-neutral so C(x) is sum of *expected travel times*

Price of Risk Aversion (**PRA**): captures inefficiency introduced by users caring for second criterion vs not (e.g., risk averse vs risk-neutral)

$$\sup_{\substack{\text{problem}\\\text{instances}}} \frac{C(x^r)}{C(x^0)}$$

Homogeneous equilibrium with parameter r (Risk-averse equilibrium)

Homogeneous equilibrium with parameter r (Risk-neutral equilibrium)

T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

Price of Risk Aversion (PRA) for Arbitrary Latency Functions

Theorem: In a general graph, **PRA ≤ 1+ηrk**

- Here, η is a graph topology parameter:
  # forward subpaths in an alternating path [η ≤ ½|V|]
- k is the max $\sigma_e(x_e)/l_e(x_e)$ ratio at equilibrium x

Intuition:

- For 2-link networks: **PRA ≤ 1+1rk**
- For series-parallel networks: **PRA ≤ 1+1rk**
- For Braess networks: **PRA ≤ 1+2rk**

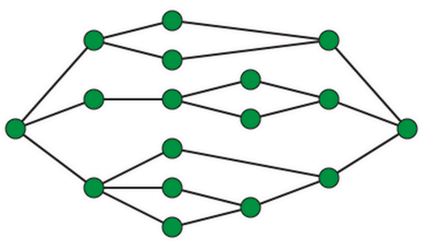

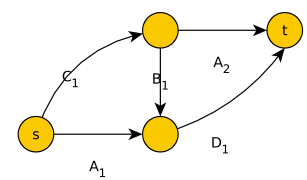

T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

Price of Risk Aversion (PRA) for Arbitrary Latency Functions

Theorem: In a general graph, **PRA** ≤ 1+ηrk

- Here, η is a graph topology parameter:
  # forward subpaths in an alternating path [η ≤ ½|V|]

Proof idea: Compare equilibria on an alternating path: forward edges have higher risk-neutral equilibrium flow, and backward edges have higher risk-averse equilibrium flow.

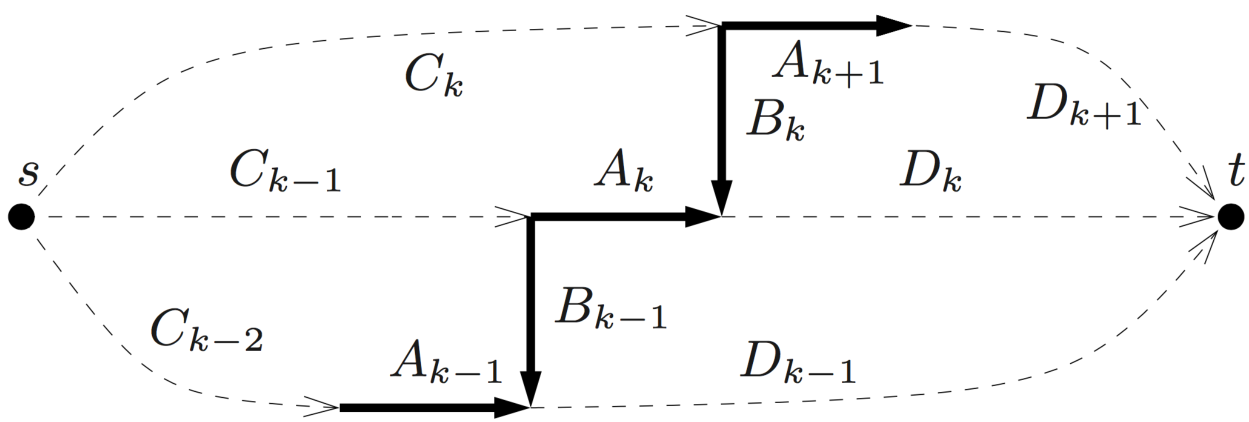

Price of Risk Aversion

- In graphs with general $l_e(x_e)$, $\sigma_e(x_e)$ functions where users minimize $l(x) + r\,\sigma(x)$,

Cost(Risk-averse eq.) ≤ (1+ηrk) Cost(Risk-neutral eq.)

- η=1 for series-parallel graphs, η=2 for Braess graph, η≤ |V|/2 for a general graph

- Alternative bound with respect to latency functions:

Cost(Risk-averse eq.) ≤ (1+rk) POA Cost(Risk-neutral eq.)

- Open: extend to nonlinear combination of criteria.

T. Lianeas, E. Nikolova, N. Stier-Moses. *Mathematics of Operations Research, forthcoming*

Questions

- Players minimize delay plus deviation:

$$c^r_{path}(x) = \sum_{e \in path} l_e(x_e) + r \sum_{e \in path} \sigma_e(x_e)$$

- Two natural questions:

I. How does equilibrium cost of population with parameter r compare to equilibrium cost of population with parameter 0? (e.g., risk-averse vs risk-neutral people, or people who care about both time and money vs those who only care about time): We answer for flow cost = first criterion only *

II. How does equilibrium cost of population with distribution or parameters D(r) compare to equilibrium cost of population with same average parameter $\bar{r}$? We answer for cost = total user cost **

* T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

** R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.

Heterogeneous vs Homogeneous Equilibrium

- We compare the cost of a heterogeneous equilibrium to that of an "averaged" homogeneous equilibrium
- For given commodity, there is d_i flow with parameter r_i so the average diversity parameter is $r = \sum_i d_i r_i$
- Compare equilibrium cost: (total demand $d = \Sigma\, d_i$)

For *heterogeneous equilibrium g*: $C^{ht}(g) = \Sigma_i\, d_i\, c^{ri}(g)$

For *homogeneous equilibrium f*: $C^{hm}(f) = d\, c^{r}(f)$

Network topologies

- Series-parallel networks (SPN):

  Inductive definition:

  1) Simplest SPN is single edge
  2) Connect 2 SPN in series or parallel

SPN

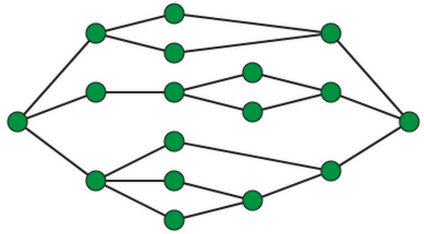

Not SPN

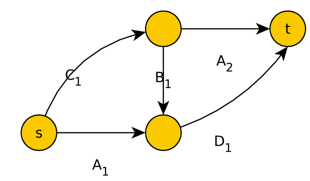

- Block representation of series-parallel networks

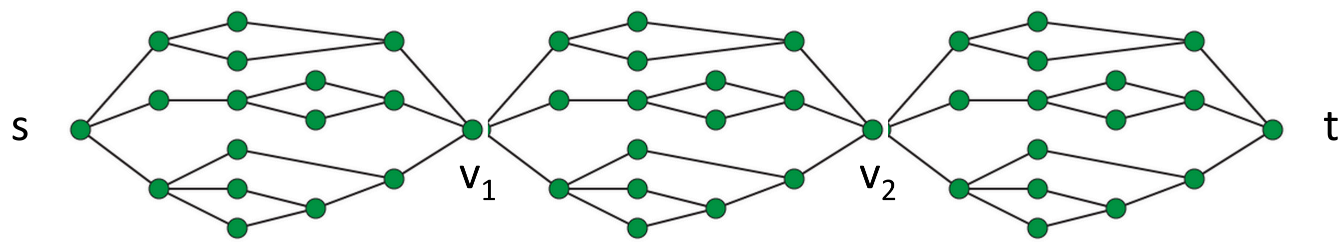

- Block matching networks (for multiple commodities)

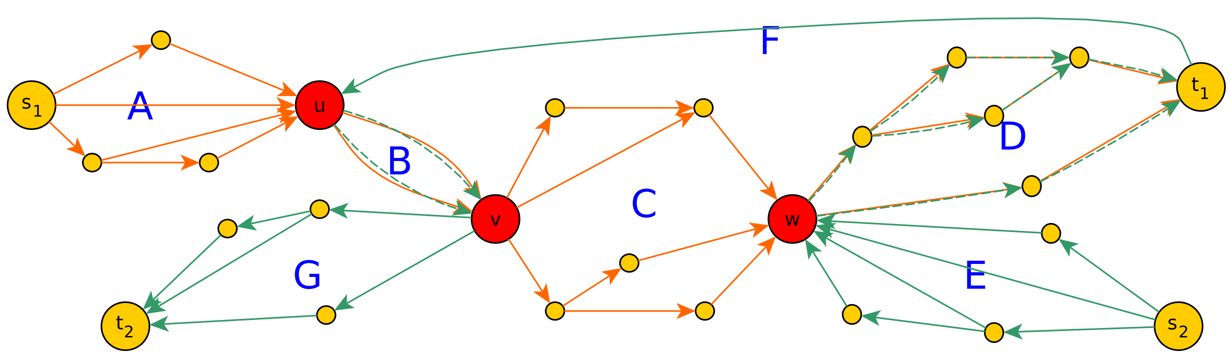

Single commodity: sufficiency

- If we have a series-parallel network, then diversity helps, i.e. $C^{ht}(g) \le C^{hm}(f)$.

- Key lemma: there exists a path P used by f s.t. $c_p(g) \le c_p(f)$, for any diversity parameter r_i.
  - Proof by induction on series-parallel structure of graph.
- Then there exists a path P used by f s.t. $c_p(g) \le c_p(f)$, for the average diversity parameter r.

$$C^{ht}(g) \le \sum_i d_i \left(\sum_{e \in P} l_e(g_e) + r_i \sum_{e \in P} \sigma_e(g_e) \right) = l_P(g) + r\sigma_P(g) \le l_P(f) + r\sigma_P(f) = C^{hm}(f)$$

Single commodity: necessity

- If diversity always helps, then the network must be series-parallel.
- Key lemma: For any strictly heterogeneous demand on the Braess graph, there exist edge functions s.t. $C^{ht}(g) > C^{hm}(f)$.

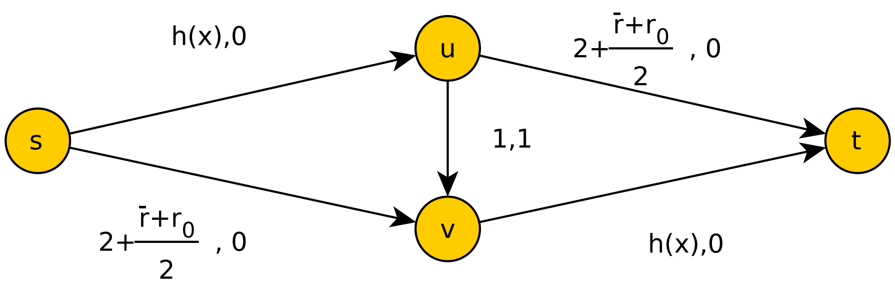

- Similarly, for any strictly heterogeneous demand on a general non-series-parallel graph, we embed the Braess construction above.

Multi-commodity: sufficiency

- If we have a block-matching network, then for any instance with average-respecting demand, diversity helps, i.e. $C^{ht}(g) \leq C^{hm}(f)$.

- Proof follows from single commodity result:
- Every commodity is routed along a series-parallel network, hence $C^{ht}(g) \leq C^{hm}(f)$ for that commodity.
- Summing up over all commodities gives results.

Multi-commodity: necessity

- Consider a multi-commodity network G. If diversity helps for every instance with average-respecting demand i.e., $C^{ht}(g) \leq C^{hm}(f)$, then G must be a block-matching network.

- Example of multi-commodity network (non-block matching) where diversity hurts:

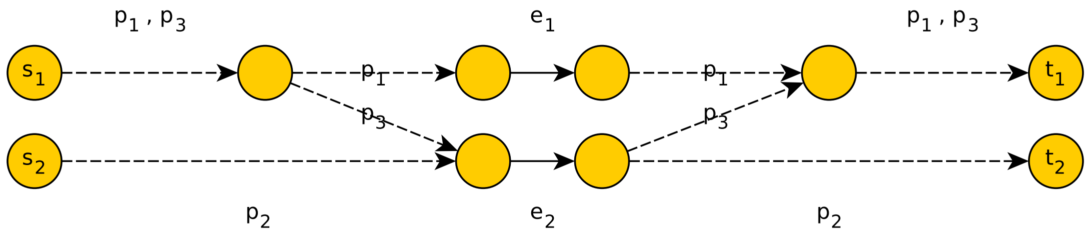

- Idea for theorem proof: by contradiction, embedding above example in a general multi-commodity network.

Multi-commodity: necessity

- Consider a multi-commodity network G. If diversity helps for every instance with average-respecting demand i.e., $C^{ht}(g) \leq C^{hm}(f)$, then G must be a block-matching network.

- By single-commodity necessity theorem, we know that sub-network for each commodity must be series-parallel.
- Remains to show that for any block B of commodity 1 and block D of commodity 2, either E(B)=E(D) or B and D do not share edges.
- Suppose the contrary, namely B and D share a common edge but (w.l.o.g.) B has an edge that is not in D.
- We'll construct edge delay and deviation functions (using previous example) such that diversity hurts, reaching a contradiction.

Multi-commodity: necessity

- Consider a multi-commodity network G. If diversity helps for every instance with average-respecting demand i.e., $C^{ht}(g) \leq C^{hm}(f)$, then G must be a block-matching network.

- By single-commodity necessity theorem, we know that sub-network for each commodity must be series-parallel.
- Remains to show that for any block B of commodity 1 and block D of commodity 2, either E(B)=E(D) or B and D do not share edges.
- Lemma 1: Let P be a simple s_2-t_2 path in G_2 that shares an edge with block B. The first edge on P in B departs from the start node of B.
- Lemma 2: All simple s_2-t_2 paths of G_2 that share an edge with block B reach the starting node of B before any of its internal nodes.

Multi-commodity: necessity

- Consider a multi-commodity network G. If diversity helps for every instance with average-respecting demand i.e., $C^{ht}(g) \leq C^{hm}(f)$, then G must be a block-matching network.

- Lemma 1: Let P be a simple s_2-t_2 path in G_2 that shares an edge with block B. The first edge on P in B departs from the start node of B.

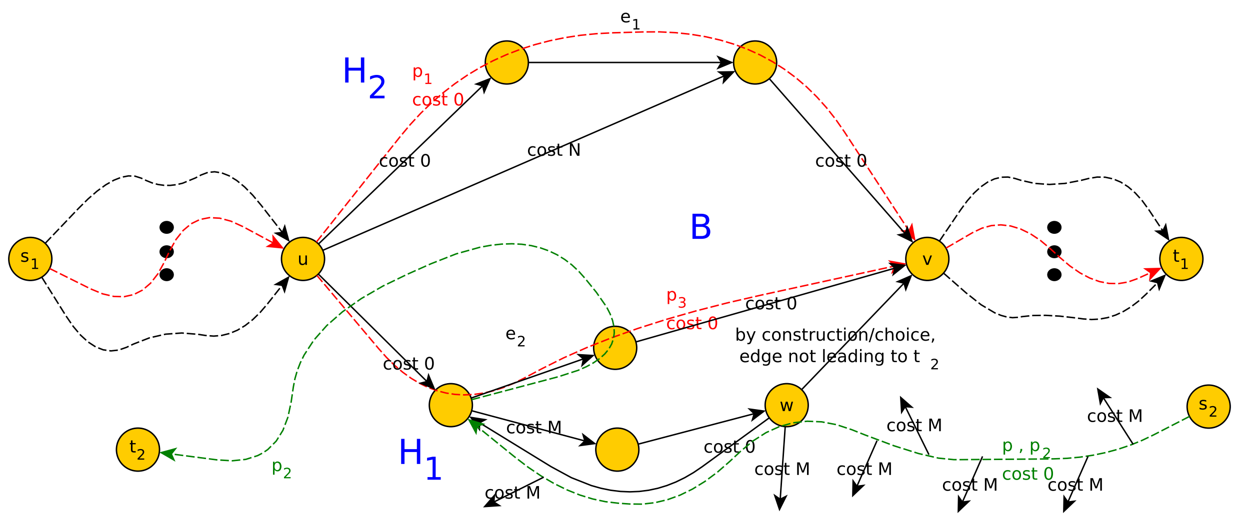

Multi-commodity: necessity

- Consider a multi-commodity network G. If diversity helps for every instance with average-respecting demand i.e., $C^{ht}(g) \leq C^{hm}(f)$, then G must be a block-ma

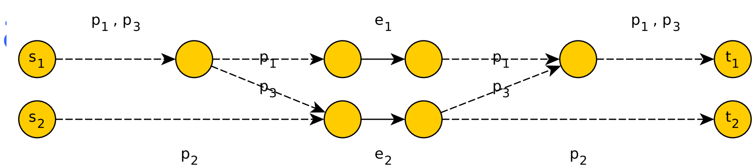

- Lemma 1: Let P be a simple s_2-t_2 path in G_2 that shares an edge with block B. The first edge on P in B departs from the start node of B.

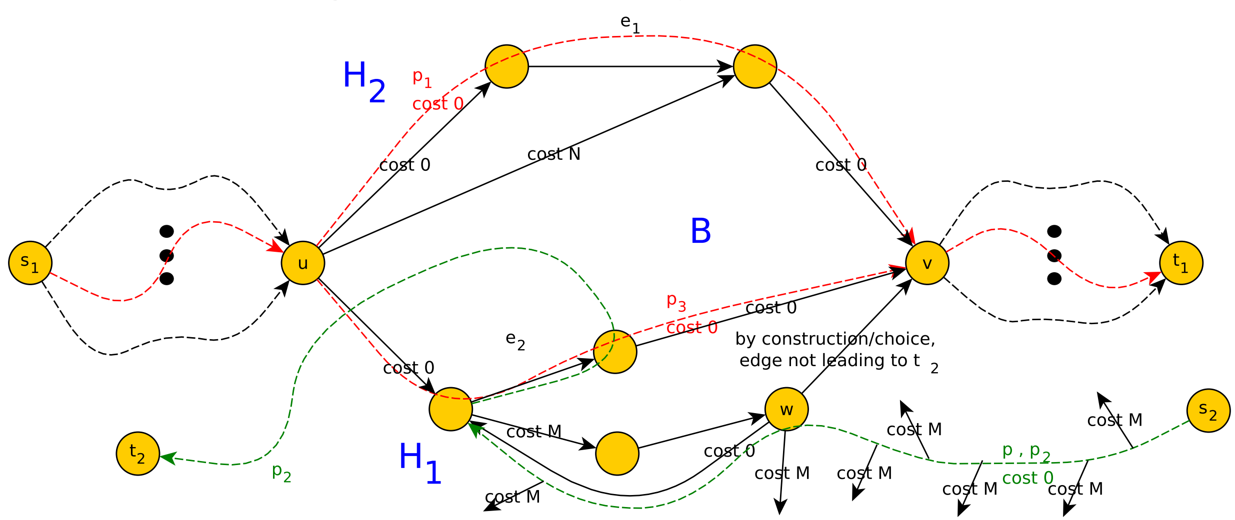

Related work

- Classic routing games:
  - Wardrop'52, Beckmann et al. '56, … surveys in Nisan et al. '07, Correa & Stier-Moses'11
- Risk-averse routing:
  - a few references in transportation (but not too many), Ordóñez & Stier-Moses'10, Nie'11, Angelidakis-Fotakis-Lianeas'13, Cominetti-Torico'13, Meir-Parkes'15, …
- Tolls with heterogeneous users:
  - Cole-Dodis-Roughgarden'03, Fleischer-Jain-Mahdian'04, Fleischer'05, Karakostas-Kolliopoulos'05, …
- Other related selfish routing models:
  - Kleer-Schäfer'16-'17, Fotakis-Spirakis '08, Acemoglu-Makhdoumi-Malekian-Ozdaglar'16, Meir-Parkes'14-'18,…

* T. Lianeas, E. Nikolova, N. Stier-Moses. *Math of OR, forthcoming*

** R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.

Summary

- Does heterogeneity (diversity) of users reduce the cost of equilibrium? Users min (delay + r_i cost)
- Diversity helps if and only if the network is series-parallel for single origin-destination.
- Diversity helps if and only if the network is "block-matched" for multiple origin-destination pairs.

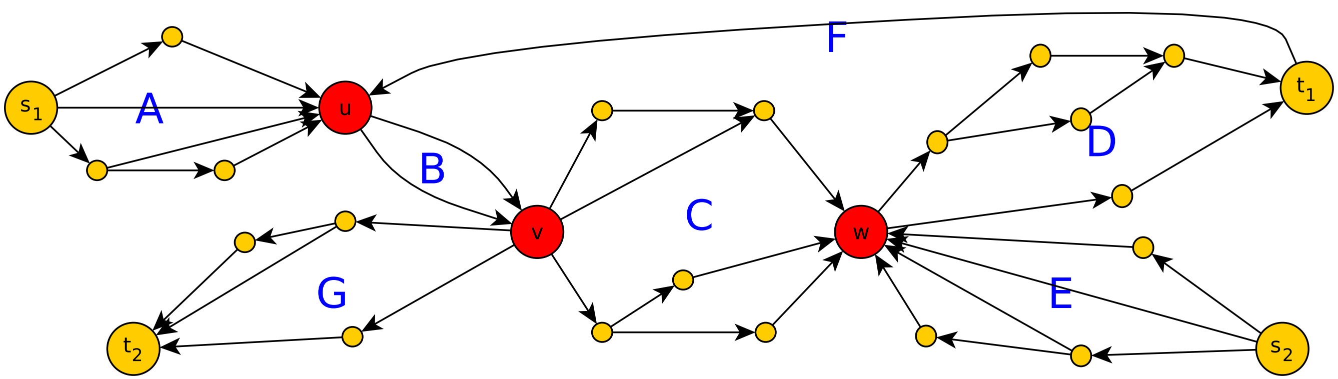

R. Cole, T. Lianeas, E. Nikolova. *IJCAI* 2018.