Complexity, Fairness, and the Price of Anarchy of the Maximum Latency Problem

Nicolás Stier Moses
MIT Operations Research Center
nstier@mit.edu

Joint work with J. Correa and A. Schulz

10th Conference on Integer Programming and Combinatorial Optimization

June 9, 2004

Outline

- Introduction
- Complexity and Approximations Results
- The Price of Anarchy and Related Results
- Summary

The Traffic Model

- Network $G = (N, A)$

- OD pairs (o_k, d_k) of rate r_k, $k \in K$

- $\mathcal{P}_k := \{$ paths from o_k to $d_k\}$, $\mathcal{P} := \cup \mathcal{P}_k$

- Nondecreasing latency functions

  $\ell_a : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for $a \in A$

  – belong to a given set $\mathcal{L}$ (e.g. polynomials)

- Travel time: $\ell_P(x) := \sum_{a \in P} \ell_a(x_a)$

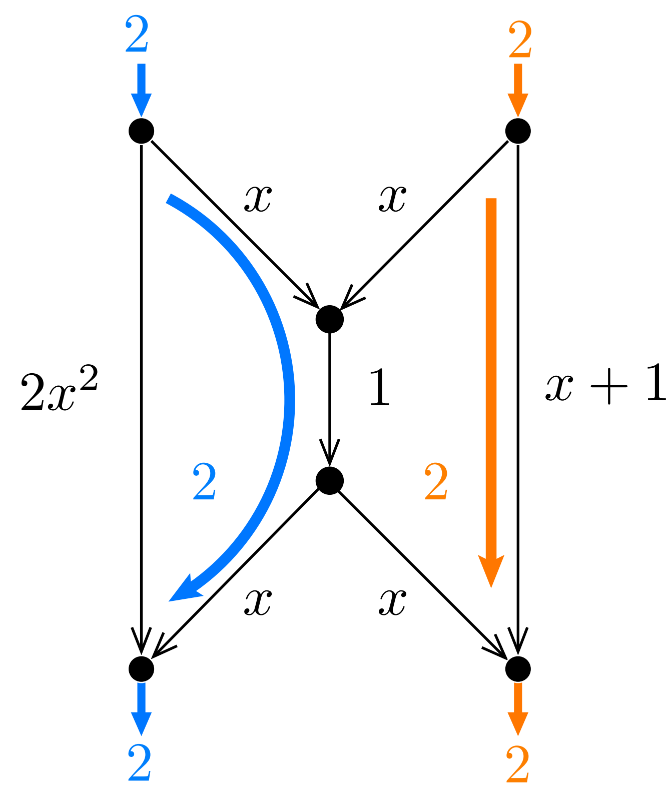

The Maximum Latency Problem

- The maximum travel time of a flow x is

$$L(x) := \max \ \{\ell_P(x) : P \in \mathcal{P}, x_P > 0\}$$

- A min-max flow f^{MM} is a minimizer of $L(x)$

- Static version of *quickest flow problem*, with load-dependent latencies — Ford & Fulkerson 1958; Köhler & Skutella 2002

- Applications:
  Computer Networks — Koutsoupias & Papadimitriou 1999
  Evacuation Problems — Hoppe & Tardos 2000

User Equilibrium

Definition: A flow f^{UE} is a **UE** if nobody can switch to a path with smaller travel time. Wardrop 1952

- **UE** exists and is essentially unique

  Beckmann, McGuire & Winsten 1956

- Unfairness $U(x) := \max_{k \in K} \left\{ \frac{\ell_Q(x)}{\ell_R(x)} : x_Q, x_R > 0, \;\; Q, R \in \mathcal{P}_k \right\}$

- $U(f^{\mathsf{UE}}) = 1$. We say **UE** is fair Wardrop 1952

System Optimum

- Average travel time $C(x) := \frac{\sum_{a \in A} \ell_a(x_a) x_a}{\sum_{k \in K} r_k}$

- A **SO** flow f^{SO} is an optimal solution to the following nonlinear multicommodity min-cost flow problem:

$$
\begin{array}{lll}
\min & C(x) & \\
\text{s.t.} & \displaystyle\sum_{P \ni a} x_P = x_a & \text{for all } a \in A \\
& \displaystyle\sum_{P \in \mathcal{P}_k} x_P = r_k & \text{for all } k \in K \\
& x_P \geq 0 & \text{for all } P \in \mathcal{P}
\end{array}
$$

An Example

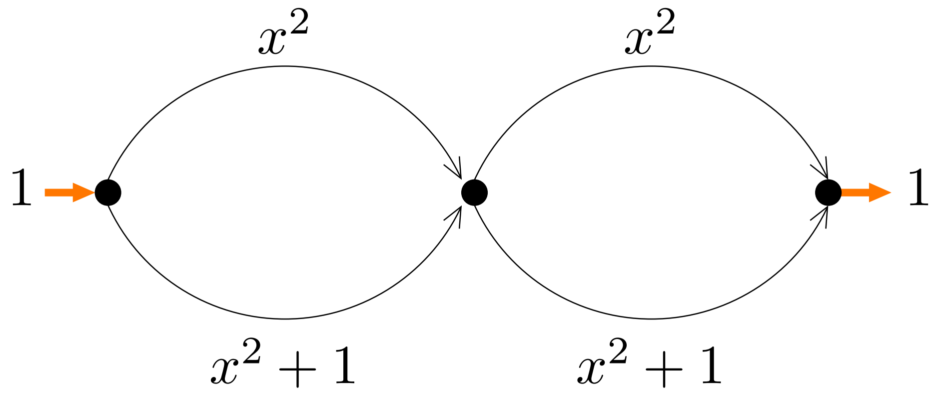

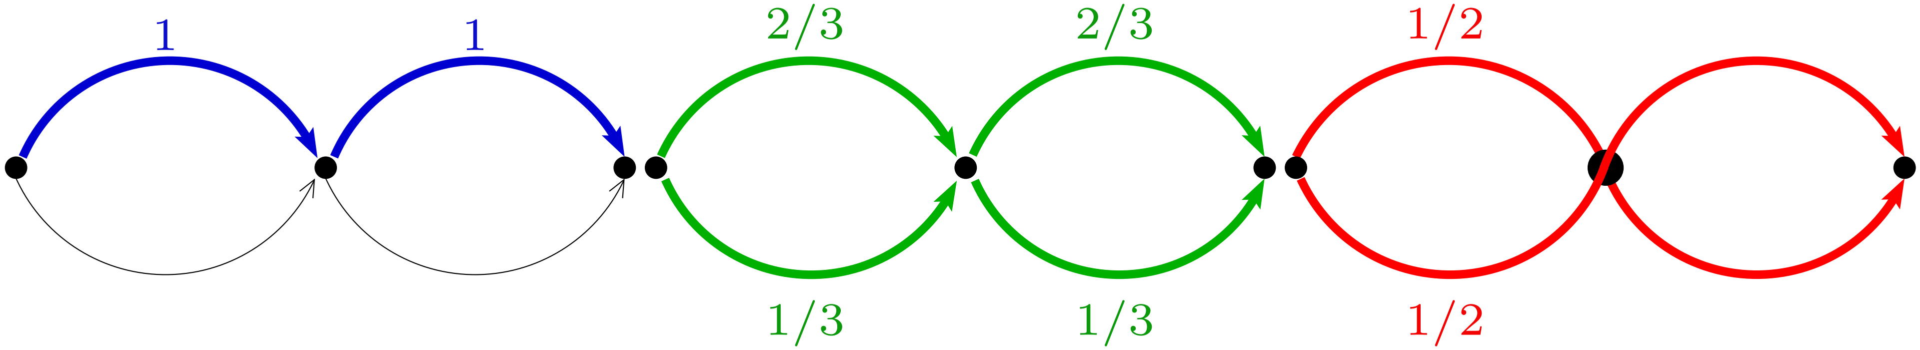

$U = 1$	$U \in \{7/4, 5/2\}$	$U = 1$
$C = 2$	$C = 4/3$	$C = 3/2$
$L = 2$	$L \in \{14/9, 20/9\}$	$L = 3/2$

Computational Complexity

- A system optimum can be computed in polynomial time, if $C(x)$ is convex
- A user equilibrium can be computed in polynomial time via a system optimum with modified latencies

Beckmann et al. 1956

- Our Results:
  - The Maximum Latency Problem is NP-complete, even for s-t-networks with linear latencies
  - Computing a **UE** or a **SO** is a constant-factor approximation algorithm

Consequence of NP-hardness Proof

Corollary: If we only know a feasible flow $(x_a)_{a \in A}$, finding a flow decomposition $(y_P)_{P \in \mathcal{P}}$ of x with $L(y) \leq L(x)$ is NP-hard

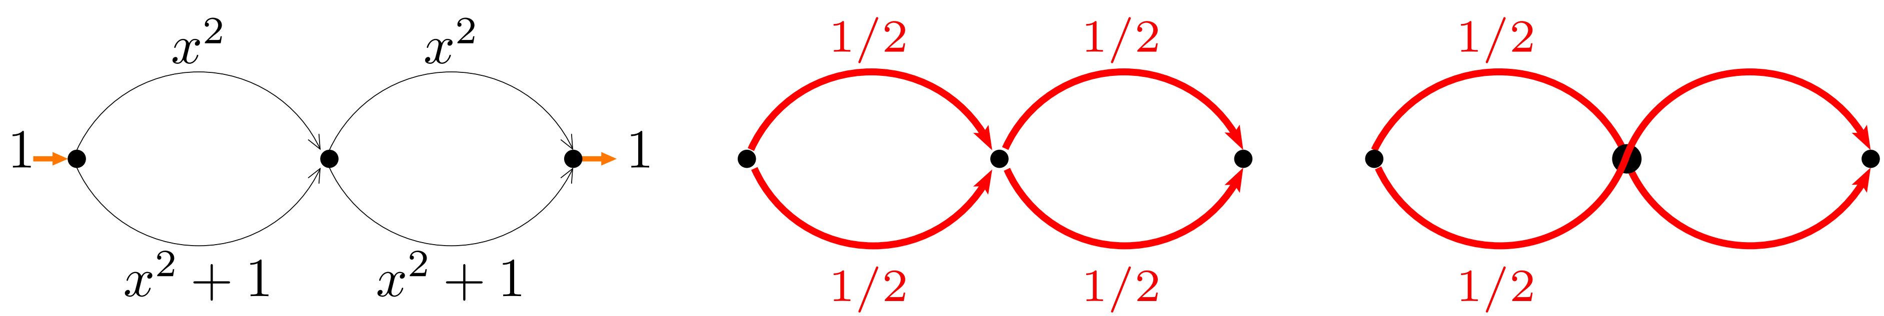

Preview of Results

(for s-t-networks with linear latencies)

	unfairness $U(x)$	avg latency $C(x)$	max latency $L(x)$
user equilibrium **UE**	1	$4/3$	$4/3$
system optimum **SO**	2	1	2
min-max flow **MM**	1	$4/3$	1

Worst-case bounds for $\dfrac{\text{col(row)}}{\text{optimal value for col}}$

Min-Max Flows are Fair with Linear Latencies

Theorem. In s-t-*networks* with *linear* latencies,

$$U(\textbf{MM}) \leq 1 \;\; U(\textbf{UE}) \qquad [U(\text{MM}) = 1]$$

Proof: Among all optimal flows, let f^{MM} be the one that uses the smallest number of paths $P_1, P_2, \ldots, P_u$.

$$\begin{array}{lll} \min & z & \\ \text{s.t.} & \displaystyle\sum_{a \in P_i} \Big(q_a x_a + r_a \Big) \leq z & \text{for } i = 1, \ldots, u \\ & \displaystyle\sum_{i=1}^{u} x_{P_i} = r & \\ & x_{P_i} \geq 0 & \text{for } i = 1, \ldots, u \end{array}$$

Min-Max Flows are not Fair in General

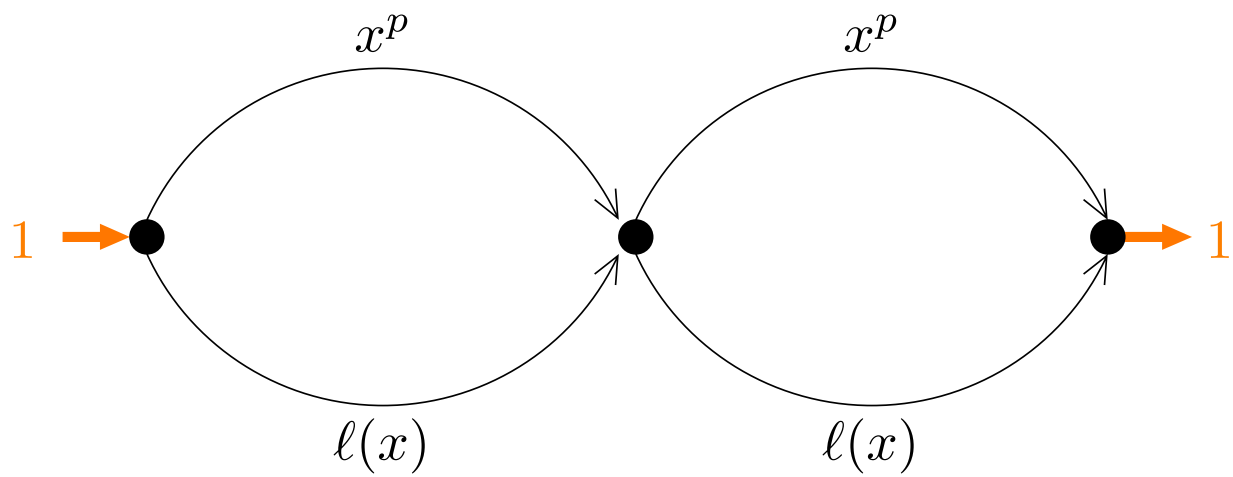

$$\ell(x) := (1+\varepsilon)^{p-1}x^p + 2 - \left(\frac{1+\varepsilon}{2+\varepsilon}\right)^{p-1} - \delta$$

- Min-max flow f^{MM} is: $\begin{cases} \frac{1}{2+\varepsilon} & \text{along "top-bottom" and "bottom-top"} \\ \frac{\varepsilon}{2+\varepsilon} & \text{along "top-top"} \end{cases}$

- $U(f^{\mathsf{MM}}) \approx 2^p$

System Optima may not be Fair

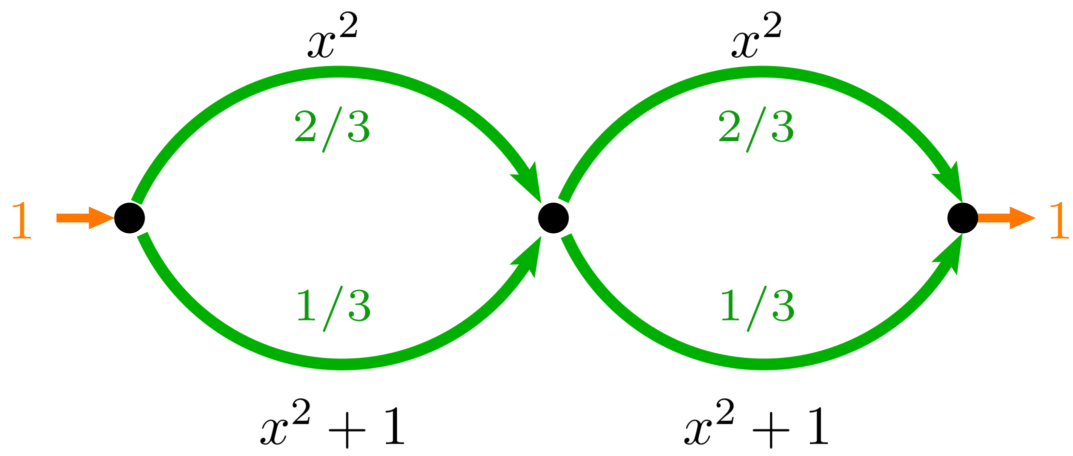

$U(f^{\mathsf{SO}}) \in \{7/4, 5/2\}$

- A **SO** is a **UE** if latencies are $\ell_a^*(x) := \ell_a(x) + x\,\ell_a'(x)$

Beckmann et al. 1956

- $\gamma(\mathcal{L}) := \sup\,\{\ell^*(x)/\ell(x) : \ell \in \mathcal{L}, x \in \mathbb{R}_+\}$ Roughgarden 2002

- Example: γ(polynomials of degree p) $= p+1$

System Optima may not be Fair

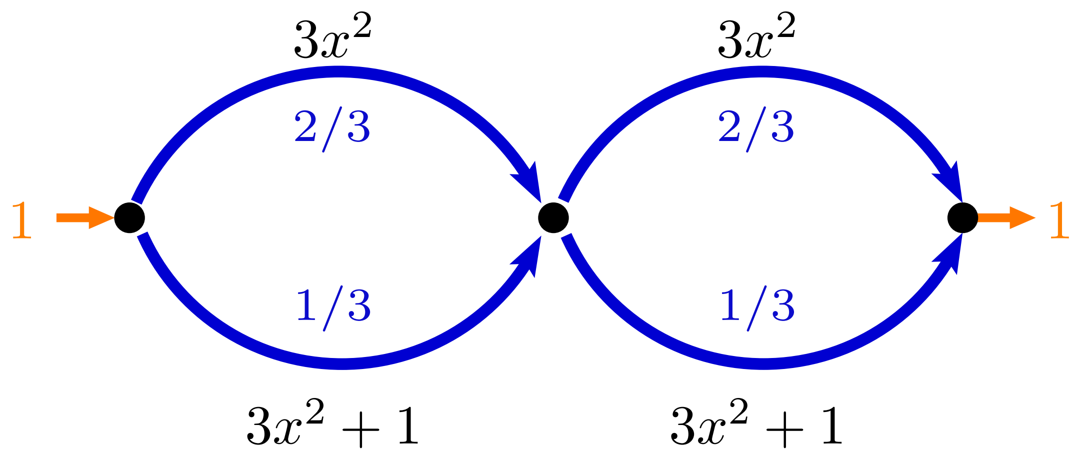

- A **SO** is a **UE** if latencies are $\ell_a^*(x) := \ell_a(x) + x\,\ell_a'(x)$

Beckmann et al. 1956

- $\gamma(\mathcal{L}) := \sup\{\ell^*(x)/\ell(x) : \ell \in \mathcal{L}, x \in \mathbb{R}_+\}$ Roughgarden 2002

- Example: γ(polynomials of degree p) $= p + 1$

System Optima are not too Unfair

Theorem. In networks with latencies drawn from a family of *differentiable* latencies $\mathcal{L}$,

$$U(\textbf{SO}) \leq \gamma(\mathcal{L})\ U(\textbf{UE}) \qquad [= \gamma(\mathcal{L})]$$

Proof:

- $\ell_a(x) \leq \ell_a^*(x) \leq \gamma(\mathcal{L})\, \ell_a(x)$, for all a
- Latencies $\ell_P^*(f^{\mathsf{SO}})$ constant for $P \in \mathcal{P}_k : f_P^{\mathsf{SO}} > 0$
- For $Q, R \in \mathcal{P}_k : f_Q^{\mathsf{SO}}, f_R^{\mathsf{SO}} > 0,\ \ \ell_Q(f^{\mathsf{SO}})/\ell_R(f^{\mathsf{SO}}) \leq \gamma(\mathcal{L})$

Unfairness of System Optima in s-t-Networks

Corollary. Consider an s-t-*network* with latencies drawn from a family of *differentiable* latencies $\mathcal{L}$. For any flow x,

$$L(\textbf{SO}) \leq \gamma(\mathcal{L})\ L(x)$$

Proof: $C(f^{\mathsf{SO}}) \leq C(x) \;\Rightarrow\; \min\{\ell_P(f^{\mathsf{SO}}) : P \in \mathcal{P}_k, f_P^{\mathsf{SO}} > 0\} \leq L(x)$

□

- Implies that for s-t-*networks* with latencies drawn from a family of *differentiable* latencies $\mathcal{L}$:

$$L(\textbf{SO}) \leq \gamma(\mathcal{L})\, L(\textbf{MM})$$

Reminder

(results for s-t-networks with linear latencies)

		unfairness $U(x)$	avg latency $C(x)$	max latency $L(x)$
user equilibrium	**UE**	1	$4/3$	$4/3$
system optimum	**SO**	2	1	2
min-max flow	**MM**	1	$4/3$	1

Worst-case bounds for $\dfrac{\text{col(row)}}{\text{optimal value for col}}$

The Price of Anarchy for Avg Latency

Theorem. In networks with *linear* latencies, Roughgarden & Tardos 2002

$$C(\mathbf{UE}) \leq \tfrac{4}{3} C(\mathbf{SO})$$

Proof: Correa, Schulz & St. 2003

$$C^{tot}(f^{\mathsf{UE}}) = \sum_P f_P^{\mathsf{UE}} \ell_P(f^{\mathsf{UE}}) \leq \sum_P x_P \ell_P(f^{\mathsf{UE}}) = \sum_a x_a(q_a f_a^{\mathsf{UE}} + r_a)$$

$$\leq \sum_a x_a(q_a x_a + r_a) + \tfrac{1}{4} \sum_a (f_a^{\mathsf{UE}})^2 q_a$$

$$\leq C^{tot}(x) + \tfrac{1}{4} C^{tot}(f^{\mathsf{UE}}),$$

because $(x_a - f_a^{\mathsf{UE}}/2)^2 \geq 0$.

The Price of Anarchy for Avg Latency II

Theorem. In networks with latencies drawn from a family of *continuous* latencies $\mathcal{L}$,

$$C(\textbf{UE}) \leq \alpha(\mathcal{L})\ C(\textbf{SO})$$

- Example: α(polynomials of degree p) $= \Omega(p/\ln p)$
- In essence, same proof as in previous slide
- Guarantee was previously known to be $\alpha(\mathcal{L})$ for differentiable latencies satisfying that $x\ell(x)$ is convex   Roughgarden 2003

Consequences of Avg Latency Result

Min-Max Flows Approximate the Avg Latency

Corollary. In s-t-*networks* with latencies drawn from a family of *continuous* latencies $\mathcal{L}$,

$$C(\textbf{MM}) \leq \alpha(\mathcal{L})\ \ C(\textbf{SO})$$

Proof:

$$C(\textbf{MM}) \leq L(\textbf{MM}) \leq L(\textbf{UE}) = C(\textbf{UE}) \leq \alpha(\mathcal{L})\, C(\textbf{SO})$$

Consequences of Previous Result II

Corollary. In s-t-*networks* with latencies drawn from a family of *continuous* latencies $\mathcal{L}$, Weitz 2001

$$L(\textbf{UE}) \leq \alpha(\mathcal{L}) \; L(\textbf{MM})$$

The Price of Anarchy is Small for Max Latency

The Price of Anarchy for Max Latency

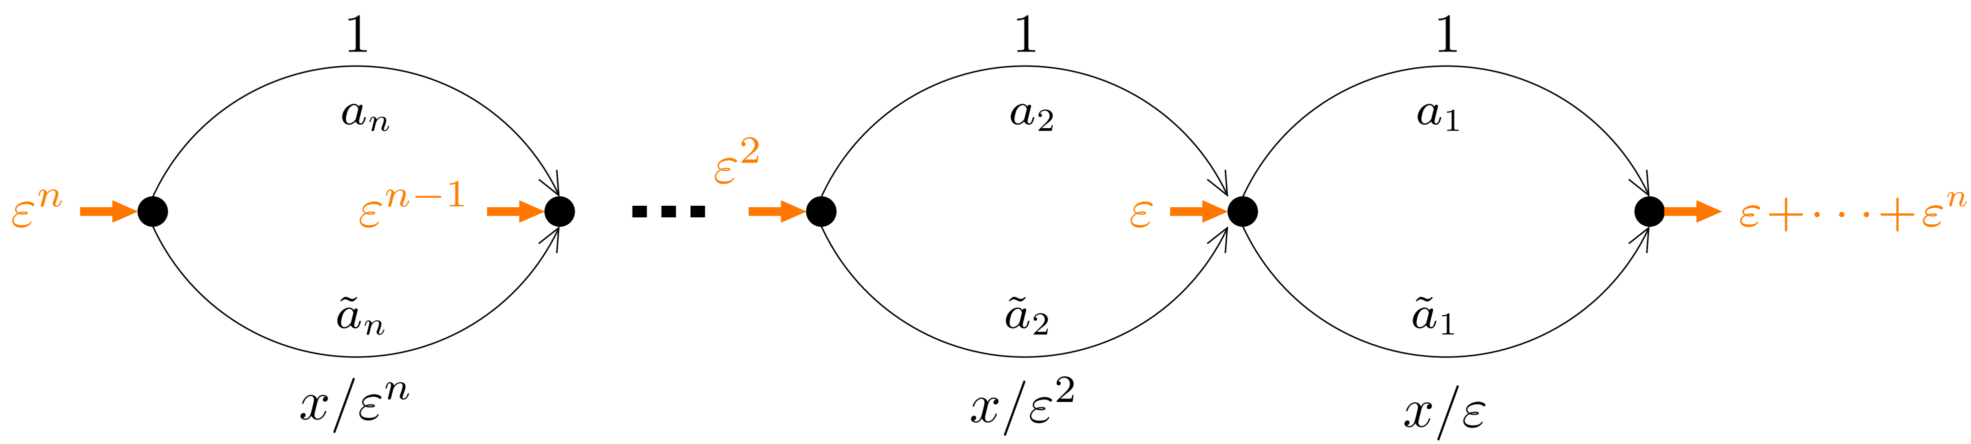

- **UE** uses paths $\tilde{a}_i, a_{i-1}, \ldots, a_1$ $\Rightarrow$ $L(\textbf{UE}) = n$
- Consider flow x that uses paths $a_i, \tilde{a}_{i-1}, \ldots, \tilde{a}_1$ $\Rightarrow$ $L(x) \xrightarrow{\varepsilon \to 0} 1$

Price of Anarchy Unbounded with Multiple Sources

Summary of Results

(for s-t-networks with latencies drawn from $\mathcal{L}$)

	unfairness $U(x)$	avg latency $C(x)$	max latency $L(x)$
user equilibrium **UE**	1	$\alpha(\mathcal{L})$	$\alpha(\mathcal{L})$
system optimum **SO**	$\gamma(\mathcal{L})$	1	$\gamma(\mathcal{L})$
min-max flow **MM**	$?$	$\alpha(\mathcal{L})$	1

All the bounds are **tight**