

Sensitivity of Wardrop Equilibria^{*}

Matthias Englert¹, Thomas Franke, and Lars Olbrich¹

¹Dept. of Computer Science, RWTH Aachen University, Germany
{englert,lars}@cs.rwth-aachen.de

Abstract. We study the sensitivity of equilibria in the well-known game theoretic traffic model due to Wardrop. We mostly consider single-commodity networks. Suppose, given a unit demand flow at Wardrop equilibrium, one increases the demand by ε or removes an edge carrying only an ε -fraction of flow. We study how the equilibrium responds to such an ε -change.

Our first surprising finding is that, even for linear latency functions, for every $\varepsilon > 0$, there are networks in which an ε -change causes every agent to change its path in order to recover equilibrium. Nevertheless, we can prove that, for general latency functions, the flow increase or decrease on every edge is at most ε .

Examining the latency at equilibrium, we concentrate on polynomial latency functions of degree at most p with nonnegative coefficients. We show that, even though the relative increase in the latency of an edge due to an ε -change in the demand can be unbounded, the path latency at equilibrium increases at most by a factor of $(1 + \varepsilon)^p$. The increase of the *price of anarchy* is shown to be upper bounded by the same factor. Both bounds are shown to be tight.

Let us remark that all our bounds are tight. For the multi-commodity case, we present examples showing that neither the change in edge flows nor the change in the path latency can be bounded.

1 Introduction


We analyze equilibria in the *Wardrop model* [15]. In this model we are given a network with load-dependent latency functions on the edges and a set of commodities, which is defined by source-sink pairs. For each commodity some demand (traffic flow) needs to be routed from the commodity's source to its sink. A common interpretation of the Wardrop model is that flow is controlled by an infinite number of selfish agents each of which carries an infinitesimal amount of flow. Each agent aims at minimizing its path latency. An allocation, in which no agent can improve its situation by unilaterally deviating from its current path is called *Wardrop equilibrium*.

Whereas the notion of equilibrium captures stability in closed systems, traffic is typically subject to external influences. Thus, from both the practical and the theoretical perspective it is a natural question, how equilibria respond to slight modifications of either the network topology or the traffic flow.

^{*} Supported by DFG grant WE 2842/1 and by the DFG GK/1298 “AlgoSyn”.

To analyze this issue, we suppose, we are given an equilibrium flow for unit demand and increase the demand by ε or remove an edge carrying only an ε -fraction of flow. How does the equilibrium responds to such an ε -change in terms of change in flow and latency?

Consider the classical network exhibiting Braess's Paradox [2]. Suppose a unit demand needs to be routed from node s to node t . At equilibrium all traffic follows the zig-zag-path. Increasing the demand by $0 < \varepsilon \leq 1$, the paths containing the dashed edges gain an ε -fraction of flow, whereas the zig-zag-path loses an ε -fraction.



Thus, in general, neither path flows nor edge flows at equilibrium are monotone functions of the demand. This observation is one of the reasons why studying the effects of changes might be intriguing.

Our findings for single-commodity networks are as follows. Allowing non-decreasing, continuous latency functions, we show that for every $\varepsilon > 0$,

- there are networks, in which after an ε -change every agent is forced to change its path in order to recover equilibrium and
- the flow increase or decrease on every edge, however, is at most ε for every network.

Thus, in contrast to our surprising finding of global instability of equilibrium flow, we can prove that edge flows are locally stable.

Examining the latency at equilibrium, we concentrate on polynomial latency functions of degree at most p with nonnegative coefficients. We show that, due to an ε -change in the demand,

- the path latency at equilibrium increases at most by a factor of $(1+\varepsilon)^p$ (even though the relative increase in the latency of an edge can be unbounded).

This result yields the same bound on the increase in the *Price of Anarchy*, as well.

All presented bounds are best possible.

For the multi-commodity case, we present examples for every $\varepsilon > 0$, showing that neither the change in edge flows nor the increase in the path latency can be bounded. This holds already for networks equipped with linear latency functions.

1.1 Related Work

The game theoretic traffic model considered in this paper was introduced by Wardrop [15]. Beckmann, McGuire, and Winston [1] observe that such an equilibrium flow is an optimal solution to a related convex program. They give existence and uniqueness results for traffic equilibria (see also [4] and [11]). Dafermos and Sparrow [4] show that the equilibrium state can be computed efficiently under some assumptions on the latency functions and many subsequent papers gave increasingly efficient methods for computing equilibria.

Another line of research examines the degradation of performance due to selfish behavior, called the Price of Anarchy [8, 11] and the inverse, the increase of the maximum latency incurred to an agent due to optimal routing [12].

Motivated by the discovery of Braess's Paradox [2] many similarly counter-intuitive and counterproductive traffic behavior have been discovered. Fisk [5] shows that considering multi-commodities the increase of one flow demand might decrease others path latencies at equilibrium. Hall [6] shows that the vector of path flows and the vector of the path latencies are continuous functions of the input demand. Furthermore, he proves that for single-commodity networks the path latency at equilibrium is a monotone function of the input demand. Dafermos and Nagurney [3] show that equilibrium flow pattern depend continuously upon the demands and (even non-separable) latency functions. More recently, Patriksson [9] gave a characterization for the existence of a directional derivative of the equilibrium solution. In [7] Joseffson and Patriksson show that while equilibrium edge costs are directionally differentiable, this does not hold for edge flows itself.

1.2 Outline

In Section 2, we introduce Wardrop's traffic model. In Section 3, we establish global instability of equilibrium flows and local stability of edge flows at equilibrium for general latency functions. For polynomial latency functions with nonnegative coefficients, we give a tight upper bound on the increase of the path latency at equilibrium due to an ε -change of the demand (Section 4). Subsequently, the same bound on the increase of the Price of Anarchy is derived. In Section 5, we briefly present some negative results for the multi-commodity case.

2 Wardrop's Traffic Model

We consider Wardrop's traffic model originally introduced in [15]. We are given a directed graph $G = (V, E)$ with non-decreasing, continuous latency functions $\ell = (\ell_e)_e$ with $\ell_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. Furthermore, we are given a set of commodities $[k] = \{1, \dots, k\}$ specified by source-sink pairs $(s_i, t_i) \in V \times V$ and flow demands d_i . The total demand is $d = \sum_{i \in [k]} d_i$. Let \mathcal{P}_i denote the admissible paths of commodity i , i.e., all paths connecting s_i and t_i , and let $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$. Let $(G, (d_i), \ell)$ denote an instance of the routing problem.

A non-negative path flow vector $(f_P)_{P \in \mathcal{P}}$ is *feasible* if it satisfies the flow demands $\sum_{P \in \mathcal{P}_i} f_P = d_i$ for all $i \in [k]$. We denote the set of all feasible flow vectors by \mathcal{F} . A path flow vector $(f_P)_{P \in \mathcal{P}}$ induces an edge flow vector $f = (f_{e,i})_{e \in E, i \in [k]}$ with $f_{e,i} = \sum_{P \in \mathcal{P}_i: e \in P} f_P$. The total flow on edge e is $f_e = \sum_{i \in [k]} f_{e,i}$. The latency of an edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of a path p is given by the sum of the edge latencies $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. The weighted average latency of commodity $i \in [k]$ is given by $L_i(f) = \sum_{e \in E} \ell_e(f_e) \cdot f_{e,i}$. Finally, the total cost of a flow is defined as $C(f) = \sum_{P \in \mathcal{P}} \ell_P(f_P) f_P$ and can be expressed as $C(f) = \sum_{e \in E} \ell_e(f_e) f_e$. We drop the argument f whenever it is clear from the context. Whenever we consider a single-commodity network, we further drop the index i .

A flow vector is considered stable when no fraction of the flow can improve its sustained latency by moving unilaterally to another path. Such a stable state is generally known as *Nash equilibrium*. In our model a flow is stable if and only if all used paths have the same minimal latency, whereas unused paths may have larger latency. We call such a flow *Wardrop equilibrium*.

Definition 1. *A feasible flow vector f is at Wardrop equilibrium if for every commodity $i \in [k]$ and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f) \leq \ell_{P_2}(f)$.*

It is well-known that Wardrop equilibria are exactly those allocations that minimize the following potential function introduced in [1]:

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} \ell_e(u) du .$$

The allocations in equilibrium do not only all have the same (optimal) potential but they also impose the same latencies on all edges. Thus, the path latencies $L_i = L_i(f)$ at equilibrium is uniquely determined. In this sense, the Wardrop equilibrium is essentially unique ([1], [4], [11]).

3 Sensitivity of Equilibrium Flows

For most of the paper we concentrate on the single-commodity case. First, for any given $\varepsilon > 0$, we present a network with linear latency functions, in which every agent needs to change its current path to recover equilibrium. Then we prove, that due to ε -changes the flow on every edge does not change by more than ε .

3.1 Instability of Equilibria: Every Agent Needs to Move

In [14] Roughgarden uses the *generalized Braess graphs* to show, that the path latency at equilibrium can arbitrarily decrease by removing several edges from a network. Our definition of B_k differs from the definition in [14] in the non-constant latency functions.




Fig. 1. Having unit demand, the solid paths in $B_{k=3}$ carry $1/3$ of flow each and the dashed edges carry zero flow. After increasing the demand by $(1 + \varepsilon) = (1 + 1/3)$, the solid paths lose all their flow and the paths containing the dashed edges gain flow of $(1 + \varepsilon)/(k + 1) = 1/3$ each.

Definition 2. For every $k \in \mathbb{N}$, let $B_k = (V_k, E_k)$ be the graph with $V_k = \{s, v_1, \dots, v_k, w_1, \dots, w_k, t\}$ and $E_k = \{(s, v_i), (v_i, w_i), (w_i, t) : 1 \leq i \leq k\} \cup \{(v_i, w_{i-1}) : 2 \leq i \leq k\} \cup \{(s, w_k)\} \cup \{(v_1, t)\}$. Let B_k be equipped with the following latency functions.

- $\ell_{v_i, w_i}^k(x) = 0$ and $\ell_{s, v_{k-i+1}}^k(x) = \ell_{w_i, t}^k(x) = i \cdot k \cdot x$ for $1 \leq i \leq k$,
- $\ell_{v_i, w_{i-1}}^k(x) = 1$ for $2 \leq i \leq k$ and
- $\ell_{s, w_k}^k(x) = \ell_{v_1, t}^k(x) = 1$.

Let B_k be called the k th Braess graph.

Let $\varepsilon > 0$ and consider the instance $(B_{\lceil 1/\varepsilon \rceil}, 1, \ell)$.

Let $(P_1, \dots, P_{2k+1})^T = (P_{s, w_k, t}, P_{s, v_k, w_k, t}, P_{s, v_k, w_{k-1}, t}, P_{s, v_{k-1}, w_{k-1}, t}, \dots, P_{s, v_1, t})^T$ denote the corresponding path vector. The equilibrium flow is described by the vector (f_{P_j}) of path flows

$$f_{P_j} = \begin{cases} 0 & \text{for } j = 1, 3, \dots, 2k + 1 \\ 1/k & \text{for } j = 2, 4, \dots, 2k \end{cases}$$

summing up to $\sum_P f_P = \sum_{j=1}^{2k+1} f_{P_j} = 1$.

All paths have path length $\ell_P(f) = k + 1$ and since any unilateral deviation strictly increases the sustained latency, the edge flows in equilibrium are unique (Figure 1).

Increasing the demand by $(1 + \varepsilon)$, the equilibrium flow vector becomes (f'_{P_j}) with

$$f'_{P_j} = \begin{cases} (1 + \varepsilon)/(k + 1) & \text{for } j = 1, 3, \dots, 2k + 1 \\ 0 & \text{for } j = 2, 4, \dots, 2k \end{cases}$$

summing up to $\sum_P f'_P = \sum_{j=1}^{2k+1} f'_{P_j} = 1 + \varepsilon$. The path latency can easily be computed to be $1 + \frac{k^2(1+\varepsilon)}{k+1}$.

Note that the path flow decomposition in equilibrium does not need to be unique. Nevertheless, we have uniqueness in B_k .

Definition 3. An edge $e \in E$ carrying flow of at most ε is called ε -edge.

Theorem 1. Let $\varepsilon > 0$ and consider $(B_{\lceil \frac{1}{\varepsilon} \rceil}, 1, \ell)$. Then, increasing the flow by ε causes the entire demand to be redistributed to recover a Wardrop equilibrium, i.e., every agent is forced to change its path. Adding another edge to the network, one can achieve the same result for the removal of an ε -edge.

Proof. For the path flow vector (f_{P_j}) and (f'_{P_j}) it holds, that, $f_{P_j} = 0 \Leftrightarrow f'_{P_j} > 0$ and $f_{P_j} > 0 \Leftrightarrow f'_{P_j} = 0$. For the second assertion, simply simulate a demand increase by directly connecting source s with sink t and choose the latency function, such that (s, t) carries an ε -fraction of flow. Then remove this edge. \square

Let us remark that under mild conditions on the latency functions Theorem 1 can easily be transferred to optimal flows, i.e., flows minimizing the total cost. This is since optimal flows are Wardrop equilibria with respect to the so-called *marginal cost functions* $h_e(x) = (x \cdot \ell_e(x))' = \ell_e(x) + x \cdot \ell'_e(x)$, if $x \cdot \ell_e(x)$ are differentiable, convex functions for $e \in E$ (see [1]). Thus, it is sufficient to change the linear latency functions in $B_{\lceil \frac{1}{\varepsilon} \rceil}$.

3.2 Edge Flows Are Locally Stable

Let $f, f' \in \mathcal{F}$ be feasible flows for demands $d \leq d'$ and let $\Delta(f, f')$ denote the difference of f' and f ,

$$(\Delta(f, f'))_e = f'_e - f_e, \forall e \in E .$$

An edge e is *positive* (with respect to f' and f), if $f'_e - f_e > 0$, and *negative* if $f'_e - f_e < 0$. A path is positive (or negative), if all its edges are positive (or negative). Observe that the flow conservation property holds for the difference of two network flows.

Definition 4. A closed path consisting of flow carrying edges is called an alternating flow cycle.

Lemma 1. Let f denote an equilibrium flow for an instance $(G, 1, \ell)$ with non-decreasing, continuous latency functions. Then there is an equilibrium flow f' for $(G, 1 + \varepsilon, \ell)$, such that $\Delta(f, f')$ does not contain an alternating flow cycle.

Proof. Let f' denote an equilibrium flow for $(G, 1 + \varepsilon, \ell)$. Assume there is an alternating flow cycle C in $\Delta(f, f')$. Since we can assume both equilibrium flows to be cycle free, we can assume that the alternating flow cycle C contains positive and negative edges. C can thus be divided into positive and negative path segments, $C = p_1 n_1 p_2 \dots n_k$, where p_i denotes a sequence of positive edges and n_i denotes a sequence of negative edges. Let u_i be the first node of p_i and denote the last node of n_i by v_i . Thus, there are two paths from u_1 to v_k in C (Figure 2).




Fig. 2. An alternating flow cycle in $\Delta(f, f')$. Solid paths are positive, the dashed paths are negative. Thus, f certainly uses the dashed paths and possibly the solid paths and $f_e > f'_e$. For f' the converse holds.

For $u, v \in V$, let $\ell(u, v)$ denote the minimum path latency from u to v under f . For $u = s$, simply write $\ell(v)$. For f' , write $\ell'(u, v)$ and $\ell'(v)$.

There are two facts we will make consistently use of. Since in equilibrium the flow routes only on shortest paths, we have

$$\ell(v) \leq \ell(u) + \ell(u, v) \text{ for any } u, v \in V \text{ ,} \tag{1}$$

and

$$\ell(v) = \ell(u) + \ell(u, v) \tag{2}$$

if there is a flow carrying path between s and v containing u . We show, that assuming f and f' being at equilibrium yields $\ell'(u_1, v_k) = \ell(u_1, v_k)$. On one hand, since n_k connects u_1 with v_k and there is more flow on every edge of n_k under f than under f' we have

$$\ell'(u_1, v_k) \leq \sum_{e \in n_k} \ell_e(f'_e) \leq \sum_{e \in n_k} \ell_e(f_e) = \ell(u_1, v_k) \text{ .}$$

For the reverse direction, we show $\ell'(v_k) \geq \ell'(u_1) + \ell(u_1, v_k)$, since then $\ell(u_1, v_k) \leq \ell'(v_k) - \ell'(u_1) \leq \ell'(u_1, v_k)$.

In the following, we repeatedly make use of equations (1) and (2).

$$\begin{aligned} \ell'(v_k) &= \ell'(u_k) + \ell'(u_k, v_k) \geq \ell'(v_{k-1}) - \ell'(u_k, v_{k-1}) + \ell'(u_k, v_k) \\ &= \ell'(u_{k-1}) + \ell'(u_{k-1}, v_{k-1}) - \ell'(u_k, v_{k-1}) + \ell'(u_k, v_k) \\ &\geq \ell'(u_1) + \sum_{i=1}^k \ell'(u_i, v_i) - \sum_{i=2}^k \ell'(u_i, v_{i-1}) \\ &\geq \ell'(u_1) + \sum_{i=1}^k \ell(u_i, v_i) - \sum_{i=2}^k \ell(u_i, v_{i-1}) \\ &\geq \ell'(u_1) + \sum_{i=1}^k (\ell(v_i) - \ell(u_i)) - \sum_{i=2}^k (\ell(v_{i-1}) - \ell(u_i)) \\ &= \ell'(u_1) - \ell(u_1) + \ell(v_k) = \ell'(u_1) + \ell(u_1, v_k) \text{ .} \end{aligned}$$

The third inequality is valid since f and f' route only on shortest paths. Explicitly, $\ell'(u_i, v_i) = \sum_{e \in p_i} \ell_e(f'_e) \geq \sum_{e \in p_i} \ell_e(f_e) \geq \ell(u_i, v_i)$ for each $i \in [k]$ and $\ell'(u_i, v_{i-1}) \leq \sum_{e \in n_i} \ell_e(f'_e) \leq \sum_{e \in n_i} \ell_e(f_e) = \ell(u_i, v_{i-1})$ for each $i \in \{2, \dots, k\}$. Thus, $\ell'(u_1, v_k) = \ell(u_1, v_k)$. We deduce that the latency on every edge $e \in n_k$ does not change due to the flow change. Since the same analysis can be conducted for any path segment p_i and n_i , the latency of both paths on C connecting two arbitrary nodes remains unchanged. Therefore, by removing the bottleneck edge flow in C no edge latency is affected and the alternating flow cycle is eliminated. We may remove the set of alternating flow cycles in any order. Adding f to the altered difference, one gets the desired equilibrium flow for demand $1 + \varepsilon$. \square

Thus, $(\Delta(f, f'))$ can be assumed a network flow of volume ε , when edges are allowed to be traversed in both directions. We can now state the following theorem.

Theorem 2. *Let f denote an equilibrium flow for an instance $(G, 1, \ell)$ with non-decreasing, continuous latency functions.*

- *Then there is an equilibrium flow f' for $(G, 1 + \varepsilon, \ell)$, such that $|(\Delta(f, f'))_e| \leq \varepsilon$ for all $e \in E$.*
- *Consider an ε -edge (u, v) in G . There is an equilibrium flow f' for $(G' = (V, E - \{(u, v)\}), 1, \ell)$, such that $|(\Delta(f, f'))_e| \leq \varepsilon$ for all $e \in E$.*

Proof. Since the difference of f and f' can be assumed alternating flow cycle free, it constitutes a network flow of volume ε . To show the second assertion, let a single ε -edge (u, v) be removed. With the same argumentation as in Lemma 1, we can exclude alternating flow cycles in $(\Delta(f, f'))$ that do not include (u, v) . Due to the flow conservation property for every node $u \neq w \neq v$, $(\Delta(f, f'))$ is a network flow from u to v of volume ε . \square

Note, that since every edge gains or loses at most ε flow (Theorem 2), with respect to the number of paths $B_{\lceil \frac{1}{\varepsilon} \rceil}$ is a minimal example exhibiting global instability.

4 Stability of the Path Latency

The latency increase at equilibrium due to a demand increase clearly depends on the latency functions. Considering polynomials with nonnegative coefficients, the maximal degree is the critical parameter. Note, that the results in this section do not trivially result from Theorem 2, since the relative flow increase on an edge might be unbounded.

Theorem 3. *Let f and f' be equilibrium flows for instances $(G, 1, \ell)$ and $(G, 1 + \varepsilon, \ell)$ with polynomial latency functions of degree at most p with nonnegative coefficients. Let L and L' denote the corresponding path latencies. Then $L' \leq (1 + \varepsilon)^p \cdot L$.*

Proof. Due to a scaling argument it is sufficient to consider monic monomials as latency functions. For equilibrium flows f and f' we have

$$L = \sum_{P \in \mathcal{P}} f_P \ell_P(f) = \sum_e f_e \ell_e(f_e) \text{ and } (1 + \varepsilon) \cdot L' = \sum_e f'_e \ell_e(f'_e) ,$$

and we want to show that $\sum_e f'^{p_e+1} \leq (1 + \varepsilon)^{p+1} \sum_e f^{p_e+1}$, where $\ell_e(x) = x^{p_e}$. Since equilibrium flows f and f' minimize the potential function

$$\Phi(x) = \sum_e \int_0^{x_e} \ell_e(u) du$$

over feasible flows x of volume 1 and $(1 + \varepsilon)$, respectively, it holds that

$$(1 + \varepsilon)^{p+1} \cdot \Phi(f) = (1 + \varepsilon)^{p+1} \cdot \sum_e \frac{1}{p_e + 1} f_e^{p_e+1} \leq \sum_e \frac{(1 + \varepsilon)^{p-p_e}}{p_e + 1} f'^{p_e+1} , \quad (\text{A})$$

and similarly,

$$\Phi(f') = \sum_e \frac{1}{p_e + 1} f'^{p_e+1} \leq \sum_e \frac{(1 + \varepsilon)^{p_e+1}}{p_e + 1} f_e^{p_e+1} . \quad (\text{B})$$

For contradiction, assume

$$(1 + \varepsilon)^{p+1} \sum_e f_e^{p_e+1} < \sum_e f'^{p_e+1} . \quad (\text{C})$$

Calculating $p \cdot (\text{A}) + (p + (p + 1)((1 + \varepsilon)^p - 1)) \cdot (\text{B}) + ((1 + \varepsilon)^p - 1) \cdot (\text{C})$ yields

$$\sum_{k=0}^p c_k \sum_{p_e=k} f_e^{p_e+1} < \sum_{k=0}^p c'_k \sum_{p_e=k} f'^{p_e+1} , \quad (3)$$

with

$$c_k = p \cdot \frac{(1 + \varepsilon)^{p+1}}{k + 1} - ((p + 1)(1 + \varepsilon)^p - 1) \cdot \frac{(1 + \varepsilon)^{k+1}}{k + 1} + ((1 + \varepsilon)^p - 1) \cdot (1 + \varepsilon)^{p+1}$$

and

$$c'_k = p \cdot \frac{(1 + \varepsilon)^{p-k}}{k + 1} - ((p + 1)(1 + \varepsilon)^p - 1) \cdot \frac{1}{k + 1} + ((1 + \varepsilon)^p - 1) .$$

In the following we show that $c'_k \leq 0$ for $0 \leq k \leq p$. Analogous arguments can be used to show $c_k \geq 0$. Hence, we have a contradiction to equation (3).

For any $0 \leq k \leq p$ and $\varepsilon = 0$, we have $c'_k = 0$. We show that c'_k is monotonically decreasing in ε (for $\varepsilon \geq 0$). The derivative of c'_k with respect to $(1 + \varepsilon)$ is

$$\frac{\partial c'_k}{\partial(1 + \varepsilon)} = p \cdot (p - k) \cdot \frac{(1 + \varepsilon)^{p-k-1}}{k + 1} - p \cdot (p + 1) \frac{(1 + \varepsilon)^{p-1}}{k + 1} + p \cdot (1 + \varepsilon)^{p-1} .$$

Thus, it is sufficient to show that

$$\frac{1}{(1+\varepsilon)^{p-k-1}} \cdot \frac{\partial c'_k}{\partial(1+\varepsilon)} = p \cdot (p-k) \cdot \frac{1}{k+1} - p \cdot (p+1) \frac{(1+\varepsilon)^k}{k+1} + p \cdot (1+\varepsilon)^k \leq 0 .$$

For $\varepsilon = 0$, the left hand side equals 0. It remains to show that $\frac{1}{(1+\varepsilon)^{p-k-1}} \cdot \frac{\partial c'_k}{\partial(1+\varepsilon)}$ is monotonically decreasing in ε (for $\varepsilon \geq 0$). This is the case since

$$\frac{\partial\left(\frac{1}{(1+\varepsilon)^{p-k-1}} \cdot \frac{\partial c'_k}{\partial(1+\varepsilon)}\right)}{\partial(1+\varepsilon)} = \frac{(k-p) \cdot p \cdot k}{k+1} \cdot (1+\varepsilon)^{k-1} \leq 0$$

and the proof is complete. \square

The bound is tight, as shown by the network consisting of two nodes connected by an edge, equipped with the latency function $\ell(x) = x^p$. Allowing negative coefficients, the relative increase obviously can be unbounded.

4.1 Increase of the Price of Anarchy

The Price of Anarchy quantifies the degradation of performance due to selfish behavior.

Definition 5. For an instance (G, d, ℓ) with equilibrium flow f and optimal flow f^* the Price of Anarchy is defined as $\frac{C(f)}{C(f^*)}$.

In [13] the Price of Anarchy is shown to be asymptotically $\Theta\left(\frac{p}{\ln p}\right)$ for polynomial latency functions of degree at most p with nonnegative coefficients.

Corollary 4. Let ρ and ρ' denote the Price of Anarchy for instances $(G, 1, \ell)$ and $(G, 1+\varepsilon, \ell)$ with polynomial latency functions of degree at most p with nonnegative coefficients. Then $\rho' \leq (1+\varepsilon)^p \cdot \rho$.

Proof. Let \bar{L}_d denote the average path latency for an optimal flow in (G, d, ℓ) . Let C_{opt} , C'_{opt} , C^* and C'^* denote the costs of an optimal flow and an equilibrium flow, respectively. Then $\rho = C^*/C_{\text{opt}}$ and $\rho' = C'^*/C'_{\text{opt}}$. Since $C_{\text{opt}} = 1 \cdot \bar{L}_1$ and $C'_{\text{opt}} = (1+\varepsilon) \cdot \bar{L}_{1+\varepsilon}$, we have

$$(1+\varepsilon) \cdot C_{\text{opt}} = (1+\varepsilon) \cdot \bar{L}_1 \leq (1+\varepsilon) \cdot \bar{L}_{1+\varepsilon} = C'_{\text{opt}} ,$$

since the average latency is clearly monotone in the demand. Thus, the increase of the Price of Anarchy can be bounded by

$$\frac{\rho'}{\rho} = \frac{C'^*/C'_{\text{opt}}}{C^*/C_{\text{opt}}} = \frac{L' \cdot (1+\varepsilon) \cdot C_{\text{opt}}}{L \cdot C'_{\text{opt}}} \leq \frac{L \cdot (1+\varepsilon)^p \cdot (1+\varepsilon) \cdot C_{\text{opt}}}{L \cdot C_{\text{opt}} \cdot (1+\varepsilon)} = (1+\varepsilon)^p ,$$

where the inequality is due to Theorem 3. \square

This upper bound is tight in the following sense: There is a network family $(G, d, \ell(p))$, such that $\lim_p \frac{\rho'/\rho}{(1+\varepsilon)^p} = 1$ for every $\varepsilon > 0$. This holds for mildly modified instances of Pigou's example [10]. Assume two nodes to be connected via two edges equipped with latency functions $\ell_1(x) = x^p$ and $\ell_2(x) = (1 + \varepsilon)^p$. We calculate $C^* = 1, C'^* = (1 + \varepsilon)^{p+1}, C_{\text{opt}} = \frac{(1+\varepsilon)^{p+1}}{(p+1)^{(p+1)/p}} + (1 - \frac{1+\varepsilon}{(p+1)^{1/p}})(1 + \varepsilon)^p$, and $C'_{\text{opt}} = \frac{(1+\varepsilon)^{p+1}}{(p+1)^{(p+1)/p}} + (1 + \varepsilon - \frac{1+\varepsilon}{(p+1)^{1/p}})(1 + \varepsilon)^p$. Thus, we have

$$\frac{\rho'}{\rho} = (1 + \varepsilon)^p \cdot \left(1 - \frac{(p + 1)^{1/p} \varepsilon p}{(p + 1)^{(p+2)/p} - p(p + 1)^{1/p}} \right),$$

and it holds that $\lim_p \frac{\rho'/\rho}{(1+\varepsilon)^p} = 1$ for every fixed $\varepsilon > 0$.

5 Instability in Multi-commodity Networks

There are no analogous results to Theorem 2 and 3 for the multi-commodity case. Figure 3 shows a network with two commodities, with both demands being 1, in which after increasing the demand of the second commodity or both demands by ε , the entire demand of the first commodity needs to be shifted to a single edge to recover an equilibrium state. If a single ε -edge is being removed, other edges might also lose an arbitrary fraction of the commodity's demand.




Fig. 3. (left) Unlabeled edges cause no latency. Assume there are $2 \cdot \lceil \frac{1}{\varepsilon} \rceil - 1$ many edges on the unique path connecting s_2 with t_2 . For $d_1 = d_2 = 1$, the flow demand of commodity 1 is uniformly spread over all $\lceil 1/\varepsilon \rceil$ paths using one edge on the path connecting s_2 and t_2 . After increasing d_2 by ε , we have $f_{(s_1, t_1)} = 1$. (right) For $d_1 = 1$ and $d_2 = k$, the path latency of the first commodity multiplicatively increases by $1 + k \cdot \varepsilon$ if both demands are increased by a factor of $(1 + \varepsilon)$.

Figure 3 also shows a network with 2 commodities. (Insisting on unit demands, one can split commodity 2 into k small commodities.) Increasing the demands by ε the path latency of commodity 1 increases by a factor of $1 + k \cdot \varepsilon$. Simple examples exhibit an even higher increase.

6 Open Problems

Suppose, given a unit demand flow at Wardrop equilibrium, one removes an edge carrying only an ε -fraction of flow. How does the path latency change after

recovering equilibrium? Considering a network with two parallel edges, one gets a lower bound of $\frac{1}{(1-\varepsilon)^p}$. Is this bound tight?

Furthermore, we believe that our bound on the increase of the path latency holds for a broader class of latency functions, namely for latency functions with bounded *elasticity*.

What can be said about the sensitivity of equilibria in related models? For instance, are analogous results possible in atomic games, where every agents control some non-negligible amount of flow each?

Acknowledgments

We wish to thank Alexander Skopalik for many insightful discussions.

References

- [1] Beckmann, M., McGuire, C.B., Winston, C.B.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
- [2] Braess, D.: Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258–268 (1968)
- [3] Dafermos, S., Nagurney, A.: Sensitivity Analysis for the Asymmetric Network Equilibrium Problem. Mathematical Programming 28, 174–184 (1984)
- [4] Dafermos, S., Sparrow, F.T.: The Traffic Assignment Problem for a General Network. Journal of Research of the National Bureau of Standards 73(2), 91–118 (1969)
- [5] Fisk, C.: More Paradoxes in the Equilibrium Assignment Problem. Transportation Research, Series B 13(4), 305–309 (1979)
- [6] Hall, M.A.: Properties of the Equilibrium State in Transportation Networks. Transportation Science 12, 208–216 (1978)
- [7] Josefsson, M., Patriksson, M.: Sensitivity Analysis of Separable Traffic Equilibria with Application to Bilevel Optimization in Network Design. Transportation Research Series B 41(1), 4–31 (2007)
- [8] Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
- [9] Patriksson, M.: Sensitivity Analysis of Traffic Equilibria. Transportation Research 38, 258–281 (2004)
- [10] Pigou, A.C.: The Economics of Welfare. Macmillan, Basingstoke (1920)
- [11] Roughgarden, T., Tardos, É.: How Bad is Selfish Routing. Journal of the ACM 49(2), 236–259 (2002)
- [12] Roughgarden, T.: How Unfair is Selfish Routing. In: Roughgarden, T. (ed.) Proc. of the 13th Annual Symposium on Discrete Algorithms (SODA), pp. 203–204 (2002)
- [13] Roughgarden, T.: The Price of Anarchy is Independent of the Network Topology. In: Proc. of the 34th Annual Symposium on Theory of Computing Discrete Algorithms (STOC), pp. 428–437 (2002)
- [14] Roughgarden, T.: On the Severity of Braess’s Paradox: Designing Networks for Selfish Users is Hard. Journal of Computer and System Sciences 72(5), 922–953 (2004)
- [15] Wardrop, J.G.: Some Theoretical Aspects of Road Traffic Research. In: Proc. of the Institute of Civil Engineers Pt. II, pp. 325–378 (1952)