

Improving the Price of Anarchy for Selfish Routing via Coordination Mechanisms

George Christodoulou*

Kurt Mehlhorn†

Evangelia Pyrga‡

Abstract

We reconsider the well-studied Selfish Routing game with affine latency functions. The Price of Anarchy for this class of games takes maximum value $4/3$; this maximum is attained already for a simple network of two parallel links, known as Pigou’s network. We improve upon the value $4/3$ by means of Coordination Mechanisms.

We increase the latency functions of the edges in the network, i.e., if $\ell_e(x)$ is the latency function of an edge e , we replace it by $\hat{\ell}_e(x)$ with $\ell_e(x) \leq \hat{\ell}_e(x)$ for all x . Then an adversary fixes a demand rate as input. The *engineered Price of Anarchy* of the mechanism is defined as the worst-case ratio of the Nash social cost in the modified network over the optimal social cost in the original network. Formally, if $\hat{C}_N(r)$ denotes the cost of the worst Nash flow in the modified network for rate r and $C_{opt}(r)$ denotes the cost of the optimal flow in the original network for the same rate then

$$ePoA = \max_{r \geq 0} \frac{\hat{C}_N(r)}{C_{opt}(r)}.$$

We first exhibit a simple coordination mechanism that achieves for any network of parallel links an engineered Price of Anarchy strictly less than $4/3$. For the case of two parallel links our basic mechanism gives $5/4 = 1.25$. Then, for the case of two parallel links, we describe an *optimal* mechanism; its engineered Price of Anarchy lies between 1.191 and 1.192.

1 Introduction

We reconsider the well-studied Selfish Routing game with affine cost functions and ask whether increasing the cost functions can reduce the cost of a Nash flow¹. In other words, the increased cost functions should induce a user behavior that reduces cost despite the fact that the cost is now determined by increased cost functions. We answer the question positively in the following sense. The Price of Anarchy, defined as the maximum ratio of Nash cost to optimal cost, is $4/3$ for this class of games. We show that increasing costs can reduce the price of anarchy to a value strictly below $4/3$ at least for the case of networks of parallel links. For a network of two parallel links, we reduce the price of anarchy to a value between 1.191 and 1.192 and prove that this is optimal. In order to state our results precisely, we need some definitions.

We consider single-commodity congestion games on networks, defined by a directed graph $G = (V, E)$, designated nodes $s, t \in V$, and a set $\ell = (\ell_e)_{e \in E}$ of non-decreasing, non-negative functions; ℓ_e is the latency function of edge $e \in E$. Let P be the set of all paths from s to t , and let $f(r)$ be a feasible s, t -flow routing r units of flow. For any $p \in P$, let $f_p(r)$ denote the amount of flow that $f(r)$ routes via path p . For ease of notation, when r is fixed and clear from context, we will write simply f, f_p instead

*University of Liverpool, United Kingdom. gchristo@liv.ac.uk

†Max-Planck-Institut für Informatik, Saarbrücken, Germany. mehlhorn@mpi-inf.mpg.de

‡Technische Universität München, Germany. pyrga@in.tum.de

¹A preliminary version of this work appeared in [8]




Figure 1: Pigou's network: We show the original network, the optimal flow and the Nash flow as a function of the rate r , respectively, the Price of Anarchy as a function of the rate ($PoA(r)$ is 1 for $r \leq 1/2$, then starts to grow until it reaches its maximum of $4/3$ at $r = 1$, and then decreases again and approaches 1 as r goes to infinity), and finally the modified latency functions. We obtain $ePoA(r) = 1$ for all r in the case of Pigou's network.

of $f(r), f_p(r)$. By definition, $\sum_{p \in P} f_p = r$. Similarly, for any edge $e \in E$, let f_e be the amount of flow going through e . We define the latency of p under flow f as $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$ and the cost of flow f as $C(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e)$ and use $C_{opt}(r)$ to denote the minimum cost of any flow of rate r . We will refer to such a minimum cost flow as an *optimal flow* (Opt). A feasible flow f that routes r units of flow from s to t is at *Nash (or Wardrop [29]) Equilibrium*² if for $p_1, p_2 \in P$ with $f_{p_1} > 0$, $\ell_{p_1}(f) \leq \ell_{p_2}(f)$. We use $C_N(r)$ to denote the maximum cost of a Nash flow for rate r . The *Price of Anarchy* (PoA) [23] (for demand r) is defined as

$$PoA(r) = \frac{C_N(r)}{C_{opt}(r)} \quad \text{and} \quad PoA = \max_{r>0} PoA(r).$$

PoA is bounded by $4/3$ in the case of affine latency functions $\ell_e(x) = a_e x + b_e$ with $a_e \geq 0$ and $b_e \geq 0$; see [28, 13]. The worst-case is already assumed for a simple network of two parallel links, known as Pigou's network; see Figure 1.

A *Coordination Mechanism*³ replaces the cost functions $(\ell_e)_{e \in E}$ by functions⁴ $\hat{\ell} = (\hat{\ell}_e)_{e \in E}$ such that $\hat{\ell}_e(x) \geq \ell_e(x)$ for all $x \geq 0$. Let $\hat{C}(f)$ be the cost of flow f when for each edge $e \in E$, $\hat{\ell}_e$ is used instead of ℓ_e and let $\hat{C}_N(r)$ be the maximum cost of a Nash flow of rate r for the modified latency functions. We define the *engineered Price of Anarchy* (for demand r) as

$$ePoA(r) = \frac{\hat{C}_N(r)}{C_{opt}(r)} \quad \text{and} \quad ePoA = \max_{r>0} ePoA(r).$$

We stress that the optimal cost refers to the original latency functions ℓ .

²This assumes continuity and monotonicity of the latency functions. For non-continuous functions, see the discussion later in this section.

³Technically, we consider *symmetric* coordination mechanisms in this work, as defined in [9] i.e., the latency modifications affect the users in a symmetric fashion.

⁴One can interpret the difference $\hat{\ell}_e - \ell_e$ as a flow-dependent toll imposed on the edge e .

Non-continuous Latency Functions: In the previous definition, as it will become clear in Section 2, it is important to allow non-continuous modified latencies. However, when we move from continuous to non-continuous latency functions, Wardrop equilibria do not always exist. Non-continuous functions have been studied by transport economists to model the effects of step-function congestion tolls and traffic lights. Several notions of equilibrium that handle discontinuities have been proposed in the literature⁵. The ones that are closer in spirit to Nash equilibria, are those proposed by Dafermos⁶ [15] and Berstein and Smith [4]. According to the Dafermos' [15] definition of *user optimization*, a flow is in equilibrium if no *sufficiently small* fraction of the users on any path, can decrease the latency they experience by switching to another path⁷. Berstein and Smith [4] introduced the concept of *User Equilibrium*, weakening further the Dafermos equilibrium, taking the fraction of the users to the limit approaching 0. The main idea of their definition is to capture the notion of the *individual commuter*, that is implicit in Wardrop's definition for continuous functions. The Dafermos equilibrium on the other hand is a stronger concept that captures the notion of coordinated deviations by *groups of commuters*.

We adopt the concept of User Equilibrium. Formally, we say that a feasible flow f that routes r units of flow from s to t is a User Equilibrium, iff for all $p_1, p_2 \in P$ with $f_{p_1} > 0$,

$$\ell_{p_1}(f) \leq \liminf_{\varepsilon \downarrow 0} \ell_{p_2}(f + \varepsilon \mathbf{1}_{p_2} - \varepsilon \mathbf{1}_{p_1}), \quad (1)$$

where $\mathbf{1}_p$ denotes the flow where only one unit passes along a path p .

Note that for continuous functions the above definition is identical to the Wardrop Equilibrium. One has to be careful when designing a Coordination Mechanism with discontinuous functions, because the existence of equilibria is not always guaranteed⁸. It is important to emphasize, that all the mechanisms that we suggest in this paper use both lower semicontinuous and regular⁹ latencies, and therefore User Equilibrium existence is guaranteed due to the theorem of [4]. Moreover, since our modified latencies are non-decreasing, all User Equilibria are also Dafermos-Sparrow equilibria. From now on, we refer to the User Equilibria as Nash Equilibria, or simply Nash flows.

Our Contribution: We demonstrate the possibility of reducing the Price of Anarchy for Selfish Routing via Coordination Mechanisms. We obtain the following results for networks of k parallel links.

- if original and modified latency functions are continuous, no improvement is possible, i.e., $ePoA \geq PoA$; see Section 2.
- for the case of affine cost functions, we describe a simple coordination mechanism that achieves an engineered Price of Anarchy strictly less than $4/3$; see Section 3. The functions $\hat{\ell}_e$ are of the form

$$\hat{\ell}_e(x) = \begin{cases} \ell_e(x) & \text{for } x \leq r_e \\ \infty & \text{for } x > r_e. \end{cases} \quad (2)$$

For the case of two parallel links, the mechanism gives $5/4$ (see Section 3.1), for Pigou's network it gives 1, see Figure 1.

- For the case of two parallel links with affine cost functions, we describe an *optimal*¹⁰ mechanism; its engineered Price of Anarchy lies between 1.191 and 1.192 (see Sections 4 and 5). It uses

⁵See [26, 24] for an excellent exposure of the relevant concepts, the relation among them, as well as for conditions that guarantee their existence.

⁶In [15], Dafermos weakened the original definition by [14] to make it closer to the concept of Nash Equilibrium.

⁷See Section 5 for a formal definition.

⁸See for example [16, 4] for examples where equilibria do not exist even for the simplest case of two parallel links and non-decreasing functions.

⁹See [4] for a definition of regular functions.

¹⁰The lower bound that we provide in Section 5 holds for all deterministic coordination mechanisms that use *non-decreasing modified latencies*, with respect to both notions of equilibrium described in the previous paragraph.

modified cost functions of the form

$$\hat{\ell}_e(x) = \begin{cases} \ell_e(x) & \text{for } x \leq r_e \text{ and } x \geq u_e \\ \ell_e(u_e) & \text{for } r_e < x < u_e. \end{cases} \quad (3)$$

The Price of Anarchy is a standard measure to quantify the effect of selfish behavior. There is a vast literature studying the Price of Anarchy for various models of selfish routing and scheduling problems (see [25]). We show that simple coordination mechanisms can reduce the Price of Anarchy for selfish routing games below the $4/3$ worst case for networks of parallel links and affine cost functions.

We believe that our arguments extend to more general cost functions, e.g., polynomial cost functions. However, the restriction to parallel links is crucial for our proof. We leave it as a major open problem to prove results for general networks or at least more general networks, e.g., series-parallel networks.

Implementation: We discuss the realization of the modified cost function in a simple traffic scenario where the driving speed on a link is a decreasing function of the flow on the link and hence the transit time is an increasing function. The step function in (3) can be realized by setting a speed limit corresponding to transit time $\ell_e(u_e)$ once the flow is above r_e . The functions in (2) can be approximately realized by access control. In any time unit only r_e items are allowed to enter the link. If the usage rate of the link is above r_e , the queue in front of the link will grow indefinitely and hence transit time will go to infinity.

Related Work: The concept of Coordination Mechanisms was introduced in (the conference version of) [9]. Coordination Mechanisms have been used to improve the Price of Anarchy in scheduling problems for parallel and related machines [9, 19, 22] as well as for unrelated machines [3, 6]; the objective is makespan minimization. Very recently, [10] considered as an objective the weighted sum of completion times. Truthful coordination mechanisms have been studied in [1, 7, 2].

Another very well-studied attempt to cope with selfish behavior is the introduction of taxes (tolls) on the edges of the network in selfish routing games [11, 18, 20, 21, 17, 5]. The disutility of a player is modified and equals her latency plus some toll for every edge that is used in her path. It is well known (see for example [11, 18, 20, 21]) that so-called marginal cost tolls, i.e., $\hat{\ell}_e(x) = \ell_e(x) + x\ell'_e(x)$, result in a Nash flow that is equal to the optimum flow for the original cost functions.¹¹ Roughgarden [27] seeks a subnetwork of a given network that has optimal Price of Anarchy for a given demand. [12] studies the question whether tolls can reduce the cost of a Nash equilibrium. They show that for networks with affine latencies, marginal cost pricing does not improve the cost of a flow at Nash equilibrium, as well as that the maximum possible benefit that one can get is no more than that of edge removal.

Discussion: The results of this paper are similar in spirit to the results discussed in the previous paragraph, but also very different. The above papers assume that taxes or tolls are determined with full knowledge of the demand rate r . Our coordination mechanisms must *a priori* decide on the modified latency functions *without knowledge of the demand*; it must determine the modified functions $\hat{\ell}$ and then an adversary selects the input rate r . More importantly, our target objectives are different; we want to minimize the ratio of the modified cost (taking into account the increase of the latencies) over the *original* optimal cost. Our simple strategy presented in Section 3 can be viewed as a generalization of link removal. Removal of a link reduces the capacity of the edge to zero, our simple strategy reduces the capacity to a threshold r_e .

¹¹It is important to observe that although the Nash flow is equal to the optimum flow, its cost with respect to the marginal cost function can be twice as large as its cost with respect to the original cost function. For Pigou's network, the marginal costs are $\hat{\ell}_1(x) = 2x$ and $\hat{\ell}_2(x) = 1$. The cost of a Nash flow of rate r with $r \leq 1/2$ is $2r^2$ with respect to marginal costs; the cost of the same flow with respect to the original cost functions is r^2 .

2 Continuous Latency Functions Yield No Improvement

The network in this section consists of k parallel links connecting s to t and the original latency functions are assumed to be continuous and non-decreasing. We show that substituting them by continuous functions brings no improvement.

Lemma 1. *Assume that the original functions ℓ_e are continuous and non-decreasing. Consider some modified latency functions $\hat{\ell}$ and some rate r for which there is a Nash Equilibrium flow \hat{f} such that the latency function $\hat{\ell}_i$ is continuous at $\hat{f}_i(r)$ for all $1 \leq i \leq k$. Then $ePoA(r) \geq PoA(r)$.*

Proof. It is enough to show that $\hat{C}_N(r) \geq C_N(r)$. Let f be a Nash flow for rate r and the original cost functions. If $f = \hat{f}$, the claim is obvious. If $\hat{f} \neq f$, there must be a j with $\hat{f}_j(r) > f_j(r)$. The local continuity of $\hat{\ell}_i$ at $\hat{f}_i(r)$, implies that $\hat{\ell}_i(\hat{f}_i(r)) = \hat{\ell}_i(\hat{f}_i'(r))$, for all $i, i' \leq k$ such that $\hat{f}_i(r), \hat{f}_i'(r) > 0$. Therefore,

$$\hat{C}_N(r) = \hat{C}(\hat{f}(r)) = \sum_{i=1}^k \hat{f}_i(r) \hat{\ell}_i(\hat{f}_i(r)) = r \cdot \hat{\ell}_j(\hat{f}_j(r)) \geq r \cdot \ell_j(\hat{f}_j(r)) \geq r \cdot \ell_j(f_j(r))$$

since $\hat{\ell}_j(x) \geq \ell_j(x)$ for all x and ℓ_j is non-decreasing. Since f is a Nash flow we have $\ell_i(f_i(r)) \leq \ell_j(f_j(r))$ for any i with $f_i(r) > 0$. Thus

$$C_N(r) = \sum_{i=1}^k f_i(r) \ell_i(f_i(r)) \leq r \cdot \ell_j(f_j(r)).$$

□

3 A Simple Coordination Mechanism

Let $\ell_i(x) = a_i x + b_i = (x + \gamma_i)/\lambda_i$ be the latency function of the i -th link, $1 \leq i \leq k$. We call λ_i the *efficiency* of the link. We order the links in order of increasing b -value and assume $b_1 < b_2 < \dots < b_k$ as two links with the same b -value may be combined (by adding their efficiencies). We may also assume $a_i > 0$ for $i < k$; if $a_i = 0$, links $i + 1$ and higher will never be used. We say that a link is *used* if it carries positive flow. The following theorem summarizes some basic facts about optimal flows and Nash flows; it is proved by straightforward calculations.¹² We state the theorem for the case that a_k is positive. The theorem is readily extended to the case $a_k = 0$ by letting a_k go to zero and determining the limit values. We will only use the theorem in situations, where $a_k > 0$.

Theorem 1. *Let $b_1 < b_2 < \dots < b_k$ and $\lambda_i \in \mathfrak{R}$ for all i . Let $\Lambda_j = \sum_{i \leq j} \lambda_i$ and $\Gamma_j = \sum_{i \leq j} \gamma_i$. Consider a fixed rate r and let f_i^* and f_i^N , $1 \leq i \leq k$, be the optimal flow and the Nash flow for rate r respectively. Let*

$$r_j = \sum_{1 \leq i < j} (b_{i+1} - b_i) \Lambda_i.$$

Then Nash uses link j for $r > r_j$ and Opt uses link j for $r > r_j/2$. If Opt uses exactly j links at rate r then

$$f_i^* = \frac{r \lambda_i}{\Lambda_j} + \delta_i/2, \quad \text{where } \delta_i = \frac{\Gamma_j \lambda_i}{\Lambda_j} - \gamma_i$$

¹²In a Nash flow all used links have the same latency. Thus, if j links are used at rate r and f_i^N is the flow on the i -th link, then $a_1 f_1^N + b_1 = \dots = a_j f_j^N + b_j \leq b_{j+1}$ and $r = f_1^N + \dots + f_j^N$. The values for r_j and f_i^N follow from this. Similarly, in an optimal flow all used links have the same marginal costs.

and

$$C_{opt}(r) = \frac{1}{\Lambda_j} (r^2 + \Gamma_j r) - \sum_{i \leq j} \frac{\delta_i^2}{4\lambda_i} = \frac{1}{\Lambda_j} (r^2 + \Gamma_j r) - C_j, \text{ where } C_j = \left(\sum_{i \leq h \leq j} (b_h - b_i)^2 \lambda_h \lambda_i \right) / (4\Lambda_j).$$

If Nash uses exactly j links at rate r then

$$f_i^N = \frac{r\lambda_i}{\Lambda_j} + \delta_i \quad \text{and} \quad C_N(r) = \frac{1}{\Lambda_j} (r^2 + \Gamma_j r).$$

If $s < r$ and Opt uses exactly j links at s and r then

$$C_{opt}(r) = C_{opt}(s) + \frac{1}{\Lambda_j} ((r-s)^2 + (\Gamma_j + 2s)(r-s)).$$

If $s < r$ and Nash uses exactly j links at s and r then

$$C_N(r) = C_N(s) + \frac{1}{\Lambda_j} ((r-s)^2 + (\Gamma_j + 2s)(r-s)).$$

Finally, $\Gamma_j + r_j = b_j \Lambda_j$ and $\Gamma_{j-1} + r_j = b_j \Lambda_{j-1}$.

We next define our simple coordination mechanism. It is governed by parameters R_1, R_2, \dots, R_{k-1} ; $R_i \geq 2$ for all i . We call the j -th link *super-efficient* if $\lambda_j > R_{j-1} \Lambda_{j-1}$. In Pigou's network, the second link is super-efficient for any choice of R_1 since $\lambda_2 = \infty$ and $\lambda_1 = 1$. Super-efficient links are the cause of high Price of Anarchy. Observe that Opt starts using the j -th link at rate $r_j/2$ and Nash starts using it at rate r_j . If the j -th link is super-efficient, Opt will send a significant fraction of the total flow across the j -th link and this will result in a high Price of Anarchy. Our coordination mechanism induces the Nash flow to use super-efficient links earlier. The latency functions $\hat{\ell}_i$ are defined as follows: $\hat{\ell}_i = \ell_i$ if there is no super-efficient link $j > i$; in particular the latency function of the highest link (= link k) is unchanged. Otherwise, we choose a threshold value T_i (see below) and set $\hat{\ell}_i(x) = \ell_i(x)$ for $x \leq T_i$ and $\hat{\ell}_i(x) = \infty$ for $x > T_i$. The threshold values are chosen so that the following behavior results. We call this behavior *modified Nash (MN)*.

Assume that Opt uses h links, i.e., $r_h/2 \leq r \leq r_{h+1}/2$. If $\lambda_{i+1} \leq R_i \Lambda_i$ for all $i, 1 \leq i < h$, MN behaves like Nash. Otherwise, let j be minimal such that link $j+1$ is super-efficient; MN changes its behavior at rate $r_{j+1}/2$. More precisely, it freezes the flow across the first j links at their current values when the total flow is equal to $r_{j+1}/2$ and routes any additional flow across links $j+1$ to k . The thresholds for the lower links are chosen in such a way that this freezing effect takes place. The additional flow is routed by using the strategy recursively. In other words, let $j_1 + 1, \dots, j_t + 1$ be the indices of the super-efficient links. Then MN changes behavior at rates $r_{j_i+1}/2$. At this rate the flow across links 1 to j_i is frozen and additional flow is routed across the higher links.

We use $C_{MN}(r) = \hat{C}_N^{R_1, \dots, R_{k-1}}(r)$ to denote the cost of MN at rate r when operated with parameters R_1 to R_{k-1} . Then $ePoA(r) = C_{MN}(r)/C_{opt}(r)$. For the analysis of MN we use the following strategy. We first investigate the benign case when there is no super-efficient link. In the benign case, MN behaves like Nash and the worst case bound of $4/3$ on the PoA can never be attained. More precisely, we will exhibit a function $B(R_1, \dots, R_{k-1})$ which is smaller than $4/3$ for all choices of the R_i 's and will prove $C_{MN}(r) \leq B(R_1, \dots, R_{k-1})C_{opt}(r)$. We then investigate the non-benign case. We will derive a recurrence relation for

$$ePoA(R_1, \dots, R_{k-1}) = \max_r \frac{\hat{C}_N^{R_1, \dots, R_{k-1}}(r)}{C_{opt}(r)}$$

In the case of a single link, i.e., $k = 1$, MN behaves like Nash which in turn is equal to Opt. Thus $ePoA() = 1$. The coming subsections are devoted to the analysis of two links and more than two links, respectively.

3.1 Two Links

The modified algorithm is determined by a parameter $R > 1$. If $\lambda_2 \leq R\lambda_1$, modified Nash is identical to Nash. If $\lambda_2 > R\lambda_1$, the modified algorithm freezes the flow across the first link at $r_2/2$ once it reaches this level. In Pigou's network we have $\ell_1(x) = x$ and $\ell_2(x) = 1$. Thus $\lambda_2 = \infty$. The modified cost functions are $\hat{\ell}_2(x) = \ell_2(x)$ and $\hat{\ell}_1(x) = x$ for $x \leq r_2/2 = 1/2$ and $\hat{\ell}_1(x) = \infty$ for $x > 1/2$. The Nash flow with respect to the modified cost function is identical to the optimum flow in the original network and $\hat{C}_N(f^*) = C(f^*)$. Thus $ePoA = 1$ for Pigou's network.

Theorem 2. *For the case of two links, $ePoA \leq \max(1 + 1/R, (4 + 4R)/(4 + 3R)$. In particular $ePoA = 5/4$ for $R = 4$.*

Proof. Consider first the benign case $\lambda_2 \leq R\lambda_1$. There are three regimes: for $r \leq r_2/2$, Opt and Nash behave identically. For $r_2/2 \leq r \leq r_2$, Opt uses both links and Nash uses only the first link, and for $r \geq r_2$, Opt and Nash use both links. $PoA(r)$ is increasing for $r \leq r_2$ and decreasing for $r \geq r_2$. The worst case is at $r = r_2$. Then $PoA(r_2) = C_N(r_2)/C_{opt}(r_2) = C_N(r_2)/(C_N(r_2) - C_2) = 1/(1 - C_2/C_N(r_2))$. We upper-bound $C_2/C_N(r_2)$. Recall that $r_2 = (b_2 - b_1)\lambda_1$, $r_2 + \Gamma_1 = b_2\lambda_1$ and $C_N(r_2) = (1/\lambda_1)(r_2^2 + \Gamma_1 r_2)$. We obtain

$$\frac{C_2}{C_N(r_2)} = \frac{(b_2 - b_1)^2 \lambda_1 \lambda_2}{4\Lambda_2(1/\lambda_1)(r_2^2 + \Gamma_1 r_2)} = \frac{(b_2 - b_1)^2 \lambda_1 \lambda_2}{4\Lambda_2(1/\lambda_1)(b_2 - b_1)\lambda_1 b_2 \lambda_1} \leq \frac{\lambda_2}{4\Lambda_2} \leq \frac{1}{4(1 + 1/R)}.$$

Thus $PoA(r) \leq B(R) := 1/(1 - R/(4(R + 1))) = (4 + 4R)/(4 + 3R)$.

We come to the case $\lambda_2 > R\lambda_1$: There are two regimes: for $r \leq r_2/2$, Opt and MN behave identically. For $r > r_2/2$, Opt uses both links and MN routes $r_2/2$ over the first link and $r - r_2/2$ over the second link. Thus for $r \geq r_2/2$:

$$ePoA(r) = \frac{C_N(r)}{C_{opt}(r)} = \frac{C_{opt}(r_2/2) + \frac{(r-r_2/2)^2}{\lambda_2} + b_2(r-r_2/2)}{C_{opt}(r_2/2) + \frac{(r-r_2/2)^2}{\Lambda_2} + b_2(r-r_2/2)} \leq \frac{\Lambda_2}{\lambda_2} \leq 1 + 1/R.$$

□

3.2 Many Links

As already mentioned, we distinguish cases. We first study the benign case $\lambda_{i+1} \leq R_i \Lambda_i$ for all i , $1 \leq i < k$, and then deal with the non-benign case.

The Benign Case: We assume $\lambda_{i+1} \leq R_i \Lambda_i$ for all i , $1 \leq i < k$. Then MN behaves like Nash. We will show $ePoA \leq B(R_1, \dots, R_{k-1}) < 4/3$; here B stands for benign case or base case. Our proof strategy is as follows; we will first show (Lemma 2) that for the i -th link the ratio of Nash flow to optimal flow is bounded by $2\Lambda_k/(\Lambda_i + \Lambda_k)$. This ratio is never more than two; in the benign case, it is bounded away from two. We will then use this fact to derive a bound on the Price of Anarchy (Lemma 4).

Lemma 2. *Let h be the number of links that Opt is using. Then*

$$\frac{f_i^N}{f_i^*} \leq \frac{2\Lambda_h}{\Lambda_i + \Lambda_h}$$

for $i \leq h$. If $\lambda_{j+1} \leq R_j \Lambda_j$ for all j , then

$$\frac{2\Lambda_h}{\Lambda_i + \Lambda_h} \leq \frac{2P}{P+1},$$

where $P := R_1 \cdot \prod_{1 < i < k} (1 + R_i)$.

Proof. For $i > j$, the Nash flow on the i -th link is zero and the claim is obvious. For $i \leq j$, we can write the Nash and the optimal flow through link i as

$$f_i^N = r\lambda_i/\Lambda_j + (\Gamma_j\lambda_i/\Lambda_j - \gamma_i) \quad \text{and} \quad f_i^* = r\lambda_i/\Lambda_h + (\Gamma_h\lambda_i/\Lambda_h - \gamma_i)/2$$

Therefore their ratio as a function of r is

$$F(r) = \frac{f_i^N}{f_i^*} = \frac{\Lambda_h}{\Lambda_j} \cdot \frac{2r + 2\Gamma_j - 2b_i\Lambda_j}{2r + \Gamma_h - b_i\Lambda_h}$$

The sign of the derivative $F'(r)$ is equal to the sign of $\Gamma_h - b_i\Lambda_h - 2\Gamma_j + 2b_i\Lambda_j$ and hence constant. Thus $F(r)$ attains its maximum either for r_j or for r_{j+1} . We have

$$\begin{aligned} F(r_{j+1}) &\leq \frac{\Lambda_h}{\Lambda_j} \cdot \frac{2r_{j+1} + 2\Gamma_j - 2b_i\Lambda_j}{2r_{j+1} + \Gamma_h - b_i\Lambda_h} = \frac{\Lambda_h}{\Lambda_j} \cdot \frac{2(b_{j+1} - b_i)\Lambda_j}{2b_{j+1}\Lambda_{j+1} - 2\Gamma_{j+1} + \Gamma_h - b_i\Lambda_h} \\ &= \frac{2(b_{j+1} - b_i)\Lambda_h}{\sum_{g \leq j+1} (2b_{j+1} - 2b_g)\lambda_g + \sum_{g \leq h} (b_g - b_i)\lambda_g} = \frac{2(b_{j+1} - b_i)\Lambda_h}{\sum_{g \leq j} (2b_{j+1} - b_g - b_i)\lambda_g + \sum_{j < g \leq h} (b_g - b_i)\lambda_g} \\ &= \frac{2(b_{j+1} - b_i)\Lambda_h}{\sum_{g \leq i} (2b_{j+1} - b_g - b_i)\lambda_g + \sum_{i < g \leq j} (2b_{j+1} - b_g - b_i)\lambda_g + \sum_{j < g \leq h} (b_g - b_i)\lambda_g} \\ &\leq \frac{2(b_{j+1} - b_i)\Lambda_h}{\sum_{g \leq i} 2(b_{j+1} - b_i)\lambda_g + \sum_{i < g \leq h} (b_{j+1} - b_i)\lambda_g} \\ &= \frac{2\Lambda_h}{\sum_{g \leq i} 2\lambda_g + \sum_{i < g \leq h} \lambda_g} = \frac{2\Lambda_h}{\Lambda_i + \Lambda_h} \end{aligned}$$

and

$$\begin{aligned} F(r_j) &\leq \frac{\Lambda_h}{\Lambda_j} \cdot \frac{2r_j + 2\Gamma_j - 2b_i\Lambda_j}{2r_j + \Gamma_h - b_i\Lambda_h} = \frac{\Lambda_h}{\Lambda_j} \cdot \frac{2(b_j - b_i)\Lambda_j}{2b_j\Lambda_j - 2\Gamma_j + \Gamma_h - b_i\Lambda_h} \\ &= \frac{2(b_j - b_i)\Lambda_h}{\sum_{g \leq j} (2b_j - 2b_g)\lambda_g + \sum_{g \leq h} (b_g - b_i)\lambda_g} = \frac{2(b_j - b_i)\Lambda_h}{\sum_{g \leq j} (2b_j - b_g - b_i)\lambda_g + \sum_{j < g \leq h} (b_g - b_i)\lambda_g} \\ &= \frac{2(b_j - b_i)\Lambda_h}{\sum_{g \leq i} (2b_j - b_g - b_i)\lambda_g + \sum_{i < g \leq j} (2b_j - b_g - b_i)\lambda_g + \sum_{j < g \leq h} (b_g - b_i)\lambda_g} \\ &\leq \frac{2(b_j - b_i)\Lambda_h}{\sum_{g \leq i} 2(b_j - b_i)\lambda_g + \sum_{i < g \leq h} (b_j - b_i)\lambda_g} = \frac{2\Lambda_h}{\sum_{g \leq i} 2\lambda_g + \sum_{i < g \leq h} \lambda_g} = \frac{2\Lambda_h}{\Lambda_i + \Lambda_h} \end{aligned}$$

If $\lambda_{j+1} \leq R_j\Lambda_j$ for all j , then $\Lambda_{j+1} = \lambda_j + \Lambda_j \leq (1 + R_j)\Lambda_j$ for all j and hence $\Lambda_h \leq \Lambda_k \leq P\Lambda_1$. \square

Lemma 3. For any reals μ , α , and β with $1 \leq \mu \leq 2$ and $\alpha/\beta \leq \mu$, $\beta\alpha \leq \frac{\mu-1}{\mu^2}\alpha^2 + \beta^2$.

Proof. We may assume $\beta \geq 0$. If $\beta = 0$, there is nothing to show. So assume $\beta > 0$ and let $\alpha/\beta = \delta\mu$ for some $\delta \leq 1$. We need to show (divide the target inequality by β^2) $\delta\mu \leq (\mu - 1)\delta^2 + 1$ or equivalently $\mu\delta(1 - \delta) \leq (1 - \delta)(1 + \delta)$. This inequality holds for $\delta \leq 1$ and $\mu \leq 2$. \square

Lemma 4. If $f_i^N/f_i^* \leq \mu \leq 2$ for all i , then $PoA \leq \mu^2/(\mu^2 - \mu + 1)$. If $\lambda_j \leq R_j\Lambda_j$ for all j , then

$$PoA \leq B(R_1, \dots, R_{k-1}) := \frac{4P^2}{3P^2 + 1},$$

where $P = R_1 \cdot \prod_{1 < j < k} (1 + R_j)$.

Proof. Assume that Nash uses j links and let L be the common latency of the links used by Nash. Then $L = a_i f_i^N + b_i$ for $i \leq j$ and $L \leq b_i = a_i f_i^N + b_i$ for $i > j$. Thus

$$\begin{aligned} C_N(r) &= Lr = \sum_i L f_i^* \leq \sum_i (a_i f_i^N + b_i) f_i^* \leq \frac{\mu - 1}{\mu^2} \sum_i a_i (f_i^N)^2 + \sum_i (a_i (f_i^*)^2 + b_i f_i) \\ &\leq \frac{\mu - 1}{\mu^2} C_N(r) + C_{opt}(r) \end{aligned}$$

and hence $PoA \leq \mu^2 / (\mu^2 - \mu + 1)$. If $\lambda_j \leq R_j \Lambda_j$ for all j , we may use $\mu = 2P / (P + 1)$ and obtain $PoA \leq 4P^2 / (3P^2 + 1)$. \square

The General Case: We come to the case where $\lambda_{i+1} \geq R_i \Lambda_i$ for some i . Let j be the smallest such i . For $r \leq r_{j+1}/2$, MN and Opt use only links 1 to j and we are in the benign case. Hence $ePoA$ is bounded by $B(R_1, \dots, R_{j-1}) < 4/3$. MN routes the flow exceeding $r_{j+1}/2$ exclusively on higher links.

Lemma 5. *MN does not use links before Opt.*

Proof. Consider any $h > j + 1$. MN starts to use link h at $s_h = r_{j+1}/2 + \sum_{j+1 \leq i < h} (b_{i+1} - b_i)(\Lambda_i - \Lambda_j)$ and Opt starts to use it at $r_h/2 = r_{j+1}/2 + \sum_{j+1 \leq i < h} (b_{i+1} - b_i)\Lambda_i/2$. We have $s_h \geq r_h/2$ since $\Lambda_i - \Lambda_j \geq \Lambda_i/2$ for $i > j$. \square

We need to bound the cost of MN in terms of the cost of Opt. In order to do so, we introduce an intermediate flow Mopt (modified optimum) that we can readily relate to MN and to Opt. Mopt uses links 1 to j to route $r_{j+1}/2$ and routes $f = r - r_{j+1}/2$ optimally across links $j + 1$ to k . Let f_i^* and f_i^m be the optimal flows and the flows of Mopt, respectively, at rate r . Let $r_s = \sum_{i \leq j} f_i^* \geq r_{j+1}/2$ be the total flow routed across the first j links in the optimal flow (the subscript s stands for small) and let

$$t = \frac{r - r_{j+1}/2}{r - r_s}$$

We will show $t \leq 1 + 1/R_j$ below. We next relate the cost of Mopt on links $j + 1$ to k to the cost of Opt on these links. To this end we scale the optimal flow on these links by a factor of t , i.e., we consider the following flow across links $j + 1$ to k : on link i , $j + 1 \leq i \leq k$, it routes $t \cdot f_i^*$. The total flow on the high links (= links $j + 1$ to k) is $r - r_{j+1}/2$ and hence Mopt incurs at most the cost of this flow on its high links. Thus

$$\sum_{i > j} \ell_i(f_i^m) f_i^m \leq \sum_{i > j} \ell_i(t f_i^*) t f_i^* \leq t^2 \left(\sum_{i > j} \ell_i(f_i^*) f_i^* \right).$$

The cost of MN on the high links is at most $ePoA(R_{j+1}, \dots, R_{k-1})$ times this cost by the induction hypothesis. We can now bound the cost of MN as follows: (Kurt changes this on September 1st; I left the old version in the document; the old version comes first, the new version comes then).

$$\begin{aligned} C_{MN}(r) &= C_N(r_{j+1}/2) + C_{MN}(\text{flow } f \text{ across links } j + 1 \text{ to } k) \\ &\leq B(R_1, \dots, R_{j-1}) C_{opt}(r_{j+1}/2) + t^2 ePoA(R_{j+1}, \dots, R_{k-1}) \left(\sum_{i > j} \ell_i(f_i^*) f_i^* \right) \\ &\leq B(R_1, \dots, R_{j-1}) C_{opt}(r_s) + t^2 ePoA(R_{j+1}, \dots, R_{k-1}) \left(\sum_{i > j} \ell_i(f_i^*) f_i^* \right) \\ &\leq \max(B(R_1, \dots, R_{j-1}), t^2 ePoA(R_{j+1}, \dots, R_{k-1})) C_{opt}(r) \end{aligned}$$

$$\begin{aligned}
C_{MN}(r) &= C_N(r_{j+1}/2) + C_{MN}(\text{flow } f \text{ across links } j+1 \text{ to } k) \\
&\leq B(R_1, \dots, R_{j-1})C_{opt}(r_{j+1}/2) + t^2 ePoA(R_{j+1}, \dots, R_{k-1}) \left(\sum_{i>j} \ell_i(f_i^*)f_i^* \right) \\
&\leq B(R_1, \dots, R_{j-1}) \left(\sum_{i\leq j} \ell_i(f_i^*)f_i^* \right) + t^2 ePoA(R_{j+1}, \dots, R_{k-1}) \left(\sum_{i>j} \ell_i(f_i^*)f_i^* \right) \\
&\leq \max(B(R_1, \dots, R_{j-1}), t^2 ePoA(R_{j+1}, \dots, R_{k-1}))C_{opt}(r)
\end{aligned}$$

Lemma 6. $t \leq 1 + 1/R_j$.

Proof. Assume that Opt uses h links where $j+1 \leq h \leq k$. Then $r_h/2 \leq r \leq r_{h+1}/2$. Let $r = r_h/2 + \delta$. According to Theorem 1, $f_i^* = r\lambda_i/\Lambda_h + (\Gamma_h\lambda_j/\Lambda_h - \gamma_i)/2$ and hence

$$r_s = \left(\frac{r_h}{2} + \delta \right) \frac{\Lambda_j}{\Lambda_h} + \frac{1}{2} \left(\frac{\Gamma_h\Lambda_j}{\Lambda_h} - \Gamma_j \right).$$

Since $\Gamma_h + r_h = b_h\Lambda_h$ and $\Gamma_j + r_j = b_j\Lambda_j$ (see Theorem 1), this simplifies to

$$r_s = \frac{\Lambda_j\delta}{\Lambda_h} + \frac{b_h\Lambda_h - \Gamma_h}{2} \frac{\Lambda_j}{\Lambda_h} + \frac{1}{2} \left(\frac{\Gamma_h\Lambda_j}{\Lambda_h} - b_j\Lambda_j + r_j \right) = \frac{\Lambda_j\delta}{\Lambda_h} + \frac{1}{2} ((b_h - b_j)\Lambda_j + r_j) = \frac{\Lambda_j\delta}{\Lambda_h} + r_s^*,$$

where $r_s^* = \frac{1}{2} ((b_h - b_j)\Lambda_j + r_j)$. We can now bound t .

$$t = \frac{r - r_{j+1}/2}{r - r_s} = \frac{r_h/2 + \delta - r_{j+1}/2}{r_h/2 + \delta - r_s^* - (\Lambda_j\delta)/\Lambda_h} \leq \max \left(\frac{r_h/2 - r_{j+1}/2}{r_h/2 - r_s^*}, \frac{1}{(1 - \frac{\Lambda_j}{\Lambda_h})} \right).$$

Next observe that

$$\begin{aligned}
\frac{r_h/2 - r_{j+1}/2}{r_h/2 - r_s^*} &= \frac{\sum_{j+1 \leq i < h} (b_{i+1} - b_i)\Lambda_i}{(\sum_{j \leq i < h} (b_{i+1} - b_i)\Lambda_i) - (b_h - b_j)\Lambda_j} = \frac{\sum_{j+1 \leq i < h} (b_{i+1} - b_i)\Lambda_i}{\sum_{j+1 \leq i < h} (b_{i+1} - b_i)(\Lambda_i - \Lambda_j)} \\
&\leq \max_{j+1 \leq i < h} \frac{\Lambda_i}{\Lambda_i - \Lambda_j} = \frac{\Lambda_{j+1}}{\Lambda_{j+1} - \Lambda_j} = \frac{\Lambda_j + \lambda_{j+1}}{\lambda_{j+1}} \leq 1 + \frac{1}{R_j}.
\end{aligned}$$

The second term in the upper bound for t is also bounded by this quantity. \square

We summarize the discussion.

Lemma 7. For every k and every j with $1 \leq j < k$. If $\lambda_{j+1} > R_j\Lambda_j$ and $\lambda_i \leq R_i\Lambda_{i+1}$ for $i < j$ then

$$ePoA(R_1, \dots, R_{k-1}) \leq \max \left(B(R_1, \dots, R_{j-1}), \left(1 + \frac{1}{R_j} \right)^2 ePoA(R_{j+1}, \dots, R_{k-1}) \right).$$

We are now ready for our main theorem.

Theorem 3. For any k , there is a choice of the parameters R_1 to R_{k-1} such that the engineered Price of Anarchy with these parameters is strictly less than $4/3$.

Proof. We define R_{k-1} , then R_{k-2} , and so on. We set $R_{k-1} = 8$. Then $ePoA(R_{k-1}) = 5/4$ and $(1 + 1/R_{k-1})^2 ePoA() = (9/8)^2 < 4/3$. Assume now that we have already defined R_{k-1} down to R_{i+1} so that $ePoA(R_{i+1}, \dots, R_{k-1}) < 4/3$ and $(1 + 1/R_j)^2 ePoA(R_{j+1}, \dots, R_{k-1}) < 4/3$ for $j \geq i + 1$. We next define R_i . We have

$$ePoA(R_i, \dots, R_{k-1}) \leq \max \left(\begin{array}{l} B(R_i, \dots, R_{k-1}), \\ \max_{j: i \leq j < k} \left(B(R_i, \dots, R_{j-1}), \left(1 + \frac{1}{R_j}\right)^2 ePoA(R_{j+1}, \dots, R_{k-1}) \right) \end{array} \right),$$

where the first line covers the benign case and the second line covers the non-benign case. We choose R_i such that $(1 + 1/R_i)^2 ePoA(R_{i+1}, \dots, R_{k-1}) < 4/3$. Then, $B(R_i, \dots, R_k) < 4/3$ and $B(R_i, \dots, j-1) < 4/3$ by Lemma 4 and the induction step is complete. \square

4 An Improved Mechanism for the Case of Two Links

In this section we present a mechanism which achieves $ePoA = 1.192$ for a network that consists of two parallel links. The ratio $C_N(r)/C_{opt}(r)$ is maximal for $r = r_2$. At this rate Nash still uses only the first link and Opt uses both links. In order to avoid this maximum ratio (if larger than 1.192), we force MN to use the second link earlier by increasing the latency of the first link after some rate x_1 , $r_2/2 \leq x_1 \leq r_2$ to a value above b_2 . In the preceding section, we increased the latency to ∞ . In this way, we avoided a bad ratio at r_2 , but paid a price for very large rates. The idea for the improved construction, is to increase the latency to a finite value. This will avoid the bad ratio, but also allows MN to use both links for large rates. In particular, we obtain the following result.

Theorem 4. *There is a mechanism for a network of two parallel links that achieves $ePoA = 1.192$.*

Proof. Recall that the price of anarchy of Nash is $(4 + 4R)/(4 + 3R)$ where $R = a_1/a_2$. Let R_0 be such that $(4 + 4R_0)/(4 + 3R_0) = 1.192$. Then $R_0 = 96/53$. We only need to consider the case $R > R_0$. The latency function of the second link is unchanged and the latency function of the first link is changed into

$$\widehat{\ell}_1(x) = \begin{cases} \ell_1(x), & x \leq x_1 \\ \ell_1(x_2), & x_1 < x \leq x_2 \\ \ell_1(x) & x > x_2. \end{cases} \quad (4)$$

where x_1 and x_2 satisfy $r_2/2 \leq x_1 \leq r_2 \leq x_2$ and will be fixed later. In words, when either the flow in the first link does not exceed x_1 , or is larger than x_2 , the network remains unchanged. However, when the flow in the first link is between these two values, the mechanism increases the latency of this link to $\ell_1(x_2)$. Let r^* be such that

$$\ell_2(r^* - x_1) = \ell_1(x_2).$$

We will fix x_1 and x_2 such that $r^* \geq r_2$.

What is the effect of this modification? For $r \leq r_2/2$, Opt and MN are the same and $ePoA(r) = 1$. For $r_2/2 \leq r \leq x_1$, MN behaves like Nash and $ePoA(r)$ increases. At $r = x_1$, MN starts to use the second link. MN will route any additional flow on the second link until $r = r^*$. At $r = r^*$, MN routes x_1 on the first link and $r^* - x_1$ on the second link. Beyond r^* , MN routes additional flow on the first link until the flow on the first link has grown to x_2 . This is the case at $r^{**} = r^* - x_1 + x_2$. For $r \geq r^{**}$, MN behaves like Nash.

Figure 2 shows the graph of $ePoA(r)$. We have $ePoA(r) = 1$ for $r \leq r_2/2$. For $r_2/2 \leq r \leq x_1$, $ePoA(r)$ increases to

$$ePoA(x_1) = \frac{a_1 x_1^2 + b_1 x_1}{C_{opt}(x_1)}.$$




Figure 2: The engineered price of anarchy for the construction of Section 4.

For $x_1 \leq r \leq r^*$, $ePoA(r)$ is convex. It will first decrease and reach the value one (this assumes that r^* is big enough) at the rate where Opt routes x_1 on the first link; after this rate it will increase again. At r^* , $ePoA$ has a discontinuity because at r^* MN routes x_1 on the first link for a cost of $\ell_1(x_1)x_1$ and at $r^* + \varepsilon$ it routes $x_1 + \varepsilon$ on the first link for a cost of $\ell_1(x_2)(x_1 + \varepsilon)$. Thus

$$\lim_{r \rightarrow r^*+} \frac{C_{MN}(r)}{C_{opt}(r)} = \lim_{r \rightarrow r^*+} \frac{\ell_1(x_2)r}{C_{opt}(r)} = \frac{\ell_1(x_2)r^*}{C_{opt}(r^*)} = \frac{\ell_2(r^* - x_1)r^*}{C_{opt}(r^*)}.$$

For $r \geq r^*$, $ePoA(r)$ decreases. Thus

$$ePoA = \max \left(\frac{a_1 x_1^2 + b_1 x_1}{C_{opt}(x_1)}, \frac{\ell_2(r^* - x_1)r^*}{C_{opt}(r^*)} \right). \quad (5)$$

It remains to show that $x_1 \leq r_2$ and $r^* \geq r_2$ can be chosen¹³ such that the right-hand side is at most 1.192. By Theorem 1,

$$C_{opt}(r) = b_1 r + \frac{a_1}{1+R} (r^2 + R r_2 r - R r_2^2/4),$$

for $r \geq r_2/2$ and $R = a_1/a_2$. Also $\ell_2(r^* - x_1) = a_2(r^* - x_1) + b_2 = a_2(r^* - x_1) + b_1 + a_1 r_2$.

We first determine the maximum $x_1 \leq r_2$ such that $a_1 x_1^2 + b_1 x_1 / C_{opt}(x_1) \leq 1.192$ for all b_1 . Since $a_1 x_1^2 + b_1 x_1 / C_{opt}(x_1)$ is decreasing in b_1 , this x_1 is determined for $b_1 = 0$. It follows that $\alpha = x_1/r_2$ is defined by the equation

$$\frac{4(R+1)\alpha^2}{4\alpha R - R + 4\alpha^2} = 1.192. \quad (6)$$

For $R \geq R_0$, this equation has a unique solution $\alpha_0 \in [1/2, 1]$, namely

$$\alpha_0 = \frac{1}{2} \cdot \frac{149R + 2\sqrt{894R(R+1)}}{125R - 24}.$$

We turn to the second term in equation (5). For $r^* > r_2$, it is a decreasing function of b_1 . Substituting $b_1 = 0$ into the second term and setting $\beta = r^*/r_2$ yields after some computation

$$ePoA_2 = \frac{4\beta(R+1)(\beta - \alpha + R)}{R(4\beta^2 + 4\beta R - R)}. \quad (7)$$

¹³The optimal choice for x_1 and r^* is such that both terms are equal and as small as possible. We were unable to solve the resulting system explicitly. We will prove in the next section that the mechanism defined by these optimal choices of the parameters x_1 and r^* is optimal.

For fixed $\alpha = \alpha_0$ and any $R \geq R_0$, $ePoA_2$ is minimized for $\beta = \beta_0 = \frac{R + \sqrt{R} \sqrt{R + 4\alpha_0(R - \alpha_0)}}{4\alpha_0}$. For $R \geq R_0$, one can prove that $\beta_0 \geq 1$, as needed. Substituting α_0 and β_0 into $ePoA_2$ yields a function of R . It is easy to see, using the derivative, that the maximum value of this function for $R \geq R_0$ is at most 1.192. \square

5 A Lower Bound for the Case of Two Links

We prove that the construction of the previous section is optimal among the class of deterministic mechanisms that guarantee the existence of an equilibrium for every rate $r > 0$ and that use non-decreasing¹⁴ latency functions. For the above mechanisms we show that $ePoA \geq 1.191$.

As in the preceding sections, we use $\hat{\ell}$ to denote the modified latency functions. As mentioned above, we are making two assumptions about the $\hat{\ell}$'s: an equilibrium flow must exist for every rate r , and $\hat{\ell}_i$ is non-decreasing, i.e., if $x < x'$, then $\hat{\ell}_i(x) \leq \hat{\ell}_i(x')$, for $i = 1, 2$. It is worthwhile to recall the equilibrium conditions for general latency functions (as given by Dafermos-Sparrow [14]): if (x, y) is an equilibrium for rate $r = x + y$, then $\hat{\ell}_2(y') \geq \hat{\ell}_1(x)$ for $y' \in (y, r]$ (otherwise $y' - y$ amount of flow would move from the first link to the second) and $\hat{\ell}_1(x') \geq \hat{\ell}_2(y)$ for $x' \in (x, r]$ (otherwise, $x' - x$ amount of flow would move from the second link to the first). Since we assume our functions to be monotone, the condition $\hat{\ell}_2(y') \geq \hat{\ell}_1(x)$ for $y' \in (y, r]$ is equivalent to $\liminf_{y' \downarrow y} \hat{\ell}_2(y') \geq \hat{\ell}_1(x)$ provided that $y < r$ (or equivalently, $x > 0$). Since we are discussing a network of two parallel links, the latter condition is in turn equivalent to (1).

Theorem 5. *The construction of Section 4 is optimal and $ePoA \geq 1.191$.*

Proof. We analyze a network with latency functions $\ell_1(x) = x$ and $\ell_2(x) = x/R + 1 = (x + R)/R$, $2 \leq R \leq 4$, and derive a lower bound as a function of the parameter R ; the restriction $2 \leq R \leq 4$ will become clear below. In a second step we choose R so as to maximize the lower bound; the optimal choice is $R = R^* \approx 2.1$. For $r \leq 1/2$, Opt uses one link and $C_{opt}(r) = r^2$, and for $r \geq 1/2$, Opt uses two links and $C_{opt}(r) = (r^2 + Rr - R/4)/(1 + R)$; $C_{opt}(1) = (3R + 4)/(4R + 4)$ and $C_N(1) = 1$. Thus $PoA = PoA(1) = (4R + 4)/(3R + 4)$. For $R \geq 2$, we have $PoA(1) \geq 12/10 = 1.2$.

Let $(x_1, 1 - x_1)$ be an equilibrium flow for rate 1 and let

$$r^* = \inf\{r; \text{there is an equilibrium flow } (x, r - x) \text{ for MN with } x > x_1\};$$

$r^* = \infty$ if there is no equilibrium flow (x, y) with $x > x_1$. The equilibrium conditions for flow $(x_1, 1 - x_1)$ imply

$$\hat{\ell}_1(x') \geq \hat{\ell}_2(1 - x_1) \geq \ell_2(1 - x_1) \geq 1 \text{ for } x_1 < x' \leq 1. \quad (8)$$

Lemma 8. *If $\hat{\ell}_1(x_1) \geq 1$ or $r^* = \infty$ or $r^* \leq 1$, $ePoA \geq 1.2$.*

Proof. If $\hat{\ell}_1(x_1) \geq 1$, we have $C_{MN}(1) \geq 1$ and hence $ePoA \geq PoA(1) \geq 1.2$.

If $r^* = \infty$, $ePoA(\infty) \geq 1 + 1/R$. For $R \leq 4$, this is at least 1.25.

If $r^* < 1$, there is an equilibrium flow (x, y) with $x > x_1$ and $r = x + y < 1$. Then $\hat{\ell}_1(x) \geq 1$ by inequality (8). Also $\hat{\ell}_2(y) \geq 1$. Thus $C_{MN}(r) \geq r$ and hence $ePoA(r) \geq r/C_{opt}(r)$. For $r \leq 1$, we have

$$\frac{r}{C_{opt}(r)} = \frac{r(1 + R)}{r^2 + Rr - R/4} \geq 1 + \frac{R/4}{r^2 + Rr - R/4} \geq 1 + \frac{R/4}{1 + R - R/4} = \frac{4 + 4R}{4 + 3R} = PoA(1) \geq 1.2.$$

\square

¹⁴It remains open whether similar arguments can be applied for showing the lower bound for non-monotone mechanisms with respect to User Equilibria.

In the light of the Lemma above, we proceed under the assumption $\hat{\ell}_1(x_1) < 1$, and hence $x_1 < 1$, and $1 < r^* < \infty$. Then $(x_1, 0)$ is an equilibrium flow, since $\hat{\ell}_1(x_1) < 1 \leq \hat{\ell}_2(y')$ for $0 < y' \leq x_1$. Thus

$$ePoA(x_1) \geq \frac{x_1^2}{C_{opt}(x_1)}. \quad (9)$$

By definition of r^* , MN routes at most x_1 on the first link for any rate $r < r^*$ and for any $\varepsilon > 0$ there is an $r < r^* + \varepsilon$ such that $(x, r-x)$ with $x > x_1$ is an equilibrium flow for MN.

For $r < r^*$, any equilibrium flow $(x, r-x)$ has $x \leq x_1$. Thus, for $x' \in (x, r] \supseteq (x_1, r]$, $\hat{\ell}_1(x') \geq \hat{\ell}_2(r-x) \geq \ell_2(r-x) \geq \ell_2(r-x_1)$. Since this inequality holds for any $r < r^*$, we have

$$\hat{\ell}_1(x') \geq \ell_2(r^* - x_1) \quad \text{for } x' \in (x_1, r^*). \quad (10)$$

For $\varepsilon > 0$, let

$$F_\varepsilon = \{(x, y); (x, y) \text{ is an equilibrium flow with } r^* \leq x + y \leq r^* + \varepsilon \text{ and } x > x_1\}.$$

Observe that F_ε is non-empty by definition of r^* .

Lemma 9. *If for arbitrarily small $\varepsilon > 0$, there is a $(x, y) \in F_\varepsilon$ with $x \geq r^*$, then $ePoA \geq PoA(1) \geq 1.2$.*

Proof. Let $r = x + y$. Then $ePoA \geq r^2 / C_{opt}(r)$. Since this inequality holds for arbitrarily small ε , $ePoA \geq (r^*)^2 / C_{opt}(r^*) \geq PoA(1) = 1.2$. \square

We proceed under the assumption that there is an $\varepsilon_0 > 0$ such that F_{ε_0} contains no pair (x, y) with $x \geq r^*$.

Lemma 10. *If for arbitrarily small $\varepsilon \in (0, \varepsilon_0)$, F_ε contains either a pair $(x, r^* - x_1)$ or pairs (x, y) and (u, v) with $y \neq v$, then $ePoA \geq \ell_2(r^* - x_1)r^* / OPT(r^*)$.*

Proof. Assume first that F_ε contains a pair $(x, r^* - x_1)$ and let $r = x + r^* - x_1$. Then

$$C_{MN}(r) = \hat{\ell}_1(x)x + \hat{\ell}_2(r^* - x_1)(r^* - x_1) \geq \ell_2(r^* - x_1)r$$

since $\hat{\ell}_1(x) \geq \ell_2(r^* - x_1)$ by (10).

Assume next that F_ε contains pairs (x, y) and (u, v) with $y \neq v$. Then $\hat{\ell}_1(x) \geq \ell_2(r^* - x_1)$ and $\hat{\ell}_1(u) \geq \ell_2(r^* - x_1)$ by (10). We may assume, $y > v$. Let $r = x + y$. Since (u, v) is an equilibrium $\hat{\ell}_2(y') \geq \hat{\ell}_1(u)$ for $y' \in (v, u+v)$ and hence $\hat{\ell}_2(y) \geq \hat{\ell}_1(u)$. Thus

$$C_{MN}(r) = \hat{\ell}_1(x)x + \hat{\ell}_2(y)y \geq \ell_2(r^* - x_1)r.$$

We have now shown that $C_{MN}(r) \geq \ell_2(r^* - x_1)r$ for r 's greater than r^* and arbitrarily close to r^* . Thus $ePoA \geq \ell_2(r^* - x_1)r^* / OPT(r^*)$. \square

We proceed under the assumption that there is an $\varepsilon_0 > 0$ such that F_{ε_0} contains no pair (x, y) with $x \geq r^*$, no pair $(x, r^* - x_1)$ and no two pairs with distinct second coordinate. In other words, there is a $y_0 < r^* - x_1$ such that all pairs in F_{ε_0} have second coordinate equal to y_0 .

Let $(x_0, y_0) \in F_{\varepsilon_0}$. Then $y_0 < r^* - x_1$. Let (x, y) be an equilibrium for rate $r = (r^* + x_1 + y_0)/2$. Then $r = (2r^* + y_0 - (r^* - x_1))/2 < r^*$ and hence $x \leq x_1$. Thus $y = r - x \geq r - x_1 = (r^* - x_1 + y_0)/2 > y_0$ and $r - y_0 > x_1$. Consider the pair $(r - y_0, y_0)$. Its rate is less than r^* and its flow across the first link is $r - y_0$ which is larger than x_1 . Thus it is not an equilibrium by the definition of r^* . Therefore there is either an $x'' \in (r - y_0, r]$ with $\hat{\ell}_1(x'') < \hat{\ell}_2(y_0)$ or a $y'' \in (y_0, r]$ with $\hat{\ell}_2(y'') < \hat{\ell}_1(r - y_0)$. We now distinguish cases.

Assume the former. Since (x, y) is an equilibrium, we have $\hat{\ell}_1(x') \geq \hat{\ell}_2(y)$ for all $x' \in (x, x + y]$ and in particular for x'' ; observe that $r - y_0 \geq x$ since $r - x = y > y_0$. Thus $\hat{\ell}_2(y_0) > \hat{\ell}_2(y)$, a contradiction to the monotonicity of $\hat{\ell}_2$.

Assume the latter. Since (x_0, y_0) is an equilibrium, we have $\hat{\ell}_2(y') \geq \hat{\ell}_1(x_0)$ for all $y' \in (y_0, x_0 + y_0]$ and in particular for y'' . Thus $\hat{\ell}_1(r - y_0) > \hat{\ell}_1(x_0)$, a contradiction to the monotonicity of $\hat{\ell}_2$; observe that $r - y_0 < x_0$ since $r < r^* \leq x_0 + y_0$.

Lemma 11.

$$ePoA \geq \min \left(1.2, \min_{x_1 \leq 1} \max \left(\frac{x_1^2}{C_{opt}(x_1)}, \min_{r^* \geq 1} \left(\frac{\ell_2(r^* - x_1)r^*}{C_{opt}(r^*)} \right) \right) \right). \quad (11)$$

Proof. If $x_1 \geq 1$ or $r^* \leq 1$ or $r^* = \infty$, we have $ePoA \geq 1.2$. So assume $x_1 < 1$ and $1 < r^* < \infty$. The argument preceding this Lemma shows that the hypothesis of either Lemma 9 or 10 is satisfied. In the former case, $ePoA \geq 1.2$. In the latter case, $ePoA \geq \max \left(\frac{x_1^2}{C_{opt}(x_1)}, \frac{\ell_2(r^* - x_1)r^*}{C_{opt}(r^*)} \right)$. This completes the proof. \square

It remains to bound

$$\min_{x_1 \leq 1} \max \left(\frac{x_1^2}{C_{opt}(x_1)}, \min_{r^* \geq 1} \frac{\ell_2(r^* - x_1)r^*}{C_{opt}(r^*)} \right) = \min_{x_1 \leq 1} \max \left(\frac{x_1^2}{C_{opt}(x_1)}, \min_{r^* \geq 1} \frac{4r^*(R+1)(r^* - x_1 + R)}{R(4(r^*)^2 + 4Rr^* - R)} \right) \quad (12)$$

from below. We prove a lower bound of 1.191. The term $x_1^2/C_{opt}(x_1)$ is increasing in x_1 . Thus there is a unique value $\alpha_1 \in [1/2, 1]$ such that the first term is larger than 1.191 for $x_1 > \alpha_1$. If the minimizing x_1 is larger than α_1 we have established the bound. The second term is minimized for $r^* = \max(1, (R + \sqrt{R^2 + 4R^2x_1 - 4Rx_1^2})/(4x_1))$. Since $x_1 \leq 1$ and hence $x_1^2 \leq x_1$, we have $(R + \sqrt{R^2 + 4R^2x_1 - 4Rx_1^2})/(4x_1) \geq 2R/4 \geq 1$ and hence $r^* = (R + \sqrt{R^2 + 4R^2x_1 - 4Rx_1^2})/(4x_1)$. The second term is decreasing in x_1 and hence we may substitute x_1 by α_1 for the purpose of establishing a lower bound. We now specialize R to 21/10. For this value of R and $x_1 = \alpha_1$

$$\frac{\ell_2(r^* - \alpha_1)r^*}{C_{opt}(r^*)} \Big|_{r^*=(R+\sqrt{R^2+4R^2\alpha_1-4R\alpha_1^2})/(4\alpha_1) \text{ and } R=21/10} \geq 1.191.$$

This completes the proof of the lower bound.

We next argue that the construction of Section 4 is optimal. Equations (5) of Section 4 for $b_1 = 0$ and Equation (12) agree. Hence our refined solution is optimal. \square

6 Open Problems

Clearly the ultimate goal is to design coordination mechanisms that work for general networks. In the case of parallel links that we studied, we showed that our mechanism approaches $4/3$, as the number of links k grows. It is still an open problem to show a bound of the form $4/5 - \alpha$, for some strictly positive α . A possible approach could be to use the ideas of Section 4. Another approach would be to define the benign case more restrictively. Assuming $R_i = 8$ for all i , we would call the following latencies benign: $\ell_1(x) = x$, and $\ell_i(x) = 1 + \varepsilon \cdot i + x/8^i$ for $i > 1$ and small positive ε . However, Opt starts using the k -th link shortly after $1/2$ and hence uses an extremely efficient link for small rates.

Also, our results hold only for affine original latency functions. What can be said for the case of more general latencies, for instance polynomials? On the more technical side, it would be interesting to study whether our lower bound construction of Section 5 can be extended to modified latency functions $\hat{\ell}$ that do not need to satisfy monotonicity.

Acknowledgments: We would like to thank Elias Koutsoupias and Spyros Angelopoulos for many fruitful discussions.

References

- [1] E. Angel, E. Bampis, and F. Pascual. Truthful algorithms for scheduling selfish tasks on parallel machines. *Theor. Comput. Sci.*, 369(1-3):157–168, 2006.
- [2] E. Angel, E. Bampis, F. Pascual, and A.-A. Tchetgna. On truthfulness and approximation for scheduling selfish tasks. *J. Scheduling*, 12(5):437–445, 2009.
- [3] Y. Azar, K. Jain, and V. S. Mirrokni. (Almost) optimal coordination mechanisms for unrelated machine scheduling. In *SODA*, pages 323–332, 2008.
- [4] D. Bernstein and T. E. Smith. Equilibria for networks with lower semicontinuous costs: With an application to congestion pricing. *Transportation Science*, 28(3):221–235, 1994.
- [5] V. Bonifaci, M. Salek, and G. Schäfer. On the efficiency of restricted tolls in network routing games. In *SAGT*, 2011, to appear.
- [6] I. Caragiannis. Efficient coordination mechanisms for unrelated machines scheduling. In *SODA*, pages 815–824, 2009.
- [7] G. Christodoulou, L. Gourvès, and F. Pascual. Scheduling selfish tasks: about the performance of truthful algorithms. In *COCOON*, pages 187–197, 2007.
- [8] G. Christodoulou, K. Mehlhorn, and E. Pyrga. Improving the Price of Anarchy for Selfish Routing via Coordination Mechanisms. In *ESA*, pages 119–130, 2011.
- [9] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
- [10] R. Cole, J. R. Correa, V. Gkatzelis, V. Mirrokni, and N. Olver. Inner product spaces for minsum coordination mechanisms. In *STOC*, 2011.
- [11] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *STOC*, pages 521–530, 2003.
- [12] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.
- [13] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. A geometric approach to the Price of Anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64:457–469, 2008.
- [14] S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, 73B(2):91–118, 1969.
- [15] S. Dafermos. An extended traffic assignment model with applications to two-way traffic. *Transportation Science*, 5:366–389, 1971.
- [16] A. de Palma and Y. Nesterov. Optimization formulations and static equilibrium in congested transportation networks. In *CORE Discussion Paper 9861, Université Catholique de Louvain, Louvain-la-Neuve*, pages 12–17, 1998.

- [17] L. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theor. Comput. Sci.*, 348(2-3):217–225, 2005.
- [18] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *FOCS*, pages 277–285, 2004.
- [19] N. Immorlica, L. Li, V. S. Mirrokni, and A. Schulz. Coordination mechanisms for selfish scheduling. In *WINE*, pages 55–69, 2005.
- [20] G. Karakostas and S. G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *FOCS*, pages 268–276, 2004.
- [21] G. Karakostas and S. G. Kolliopoulos. The efficiency of optimal taxes. In *CAAN*, pages 3–12, 2004.
- [22] K. Kollias. Non-preemptive coordination mechanisms for identical machine scheduling games. In *SIROCCO*, pages 197–208, 2008.
- [23] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [24] P. Marcotte and M. Patriksson. Traffic equilibrium. In *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, chapter 10, pages 623–713. North-Holland, 2007.
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [26] M. Patriksson. *The Traffic Assignment Problem: Models and Methods*. V.S.P. Intl Science, 1994.
- [27] T. Roughgarden. Designing networks for selfish users is hard. In *FOCS*, 2001.
- [28] T. Roughgarden and É. Tardos. How bad is selfish routing? *JACM*, 49:236–259, 2002.
- [29] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Part II*, volume 1, pages 325–378, 1952.