Improving the Price of Anarchy for Selfish Routing via Coordination Mechanisms

George Christodoulou[1], Kurt Mehlhorn[2], and Evangelia Pyrga[3]

[1] University of Liverpool, United Kingdom
gchristo@liv.ac.uk
[2] Max-Planck-Institut für Informatik, Saarbrücken, Germany
mehlhorn@mpi-inf.mpg.de
[3] Technische Universität München, Germany
pyrga@in.tum.de

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 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 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

C. Demetrescu and M.M. Halldórsson (Eds.): ESA 2011, LNCS 6942, pp. 119–130, 2011.
© Springer-Verlag Berlin Heidelberg 2011

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 [28]) Equilibrium*[1] 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)* [22] (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 [27,12]. The worst-case is already assumed for a simple network of two parallel links, known as Pigou's network; see Figure 1.

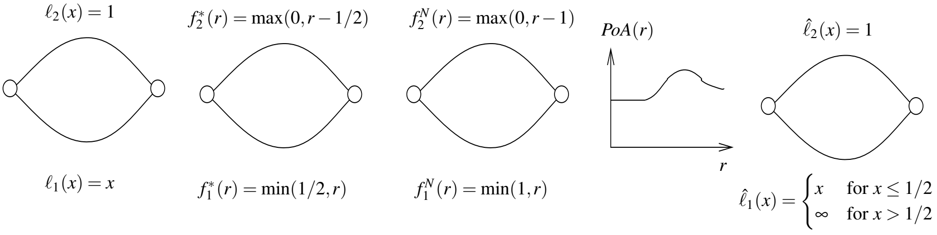
$\ell_2(x) = 1$
$\ell_1(x) = x$
$f_2^*(r) = \max(0, r - 1/2)$
$f_1^*(r) = \min(1/2, r)$
$f_2^N(r) = \max(0, r - 1)$
$f_1^N(r) = \min(1, r)$
$PoA(r)$
r
$\hat{\ell}_2(x) = 1$
$\hat{\ell}_1(x) = \begin{cases} x & \text{for } x \leq 1/2 \\ \infty & \text{for } x > 1/2 \end{cases}$

Fig. 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

A *Coordination Mechanism*[2] replaces the cost functions $(\ell_e)_{e \in E}$ by functions[3] $\hat{\ell} = (\hat{\ell}_e)_{e \in E}$ 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.

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

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

[2] Technically, we consider *symmetric* coordination mechanisms in this work, as defined in [8] i.e., the latency modifications affect the users in a symmetric fashion.

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

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 literature[4]. The ones that are closer in spirit to Nash equilibria, are those proposed by Dafermos[5] [14] and Berstein and Smith [4]. According to the Dafermos' [14] 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 path[6]. 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* of the users, that was 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 $f_{p_1} > 0$,

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

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 guaranteed[7]. It is important to emphasize, that all the mechanisms that we suggest in this paper use both lower semicontinuous and regular[8] 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. Finally, the lower bounds that we provide, do not rely on any monotonicity assumptions, and hold for general coordination mechanisms as defined above. 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.

[4] See [25,23] for an excellent exposure of the relevant concepts, the relation among them, as well as for conditions that guarantee their existence.

[5] In [14], Dafermos weakened the orginal definition by [13]to make it closer to the concept of Nash Equilibrium.

[6] See Section 5 for a formal definition.

[7] See for example [15,4] for examples where equilibria do not exist even for the simplest case of two parallel links and non-decreasing functions.

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

- 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

$$\hat{\ell}_e(x) = \begin{cases} \ell_e(x) & \text{for } x \leq r_e \\ \infty & \text{for } x > r_e. \end{cases} \tag{1}$$

  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*[9] mechanism; its engineered Price of Anarchy lies between 1.191 and 1.192 (see Sections 4 and 5). It uses modified cost functions of the form

$$\hat{\ell}_e(x) = \begin{cases} \ell_e(x) & \text{for } x \leq r_e \text{ and } x \geq u_e \\ \ell_e(u_e) & \text{for } r_e < x < u_e. \end{cases} \tag{2}$$

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 [24]). 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.

Due to space limitations, some of the proofs are omitted and deferred to the full version of this paper.

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 (2) 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 (1) 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) [8]. Coordination Mechanisms have been used to improve the Price of Anarchy in scheduling problems for parallel and related machines[8,18,21] as well as for unrelated machines [3,6]; the objective is makespan minimization. Very recently,

[9] The lower bound that we provide in Section 5 holds for all deterministic coordination mechanisms with respect to the Dafermos-Sparrow[13] definition of equilibria. However, the arguments of our proof work for all deterministic coordination mechanisms that use *non-decreasing modified latencies* even for the weaker definition of User Equilibrium. Therefore the mechanism of Section 4 is optimal for these two classes of mechanisms.

[9] considered as 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 [10,17,19,20,16,5]. 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 [10,17,19,20]) 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.[10] Roughgarden [26] seeks a subnetwork of a given network that has optimal Price of Anarchy for a given demand. [11] 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.

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 $ePoA(r) \geq PoA(r)$.*

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

[10] It 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.

assume $b_1 < b_2 < \ldots < b_k$ 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.[11]

Theorem 1. *Let $\Lambda_j = \sum_{i \le j} \lambda_i$ and $\Gamma_j = \sum_{i \le j} \gamma_j$. Consider a fixed rate r and let f_i^* and f_i^N, $1 \le i \le k$ be the optimal flow and the Nash flow for rate r respectively. Let*

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

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

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

and

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

If Nash uses exactly j links at rate r then

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

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

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

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

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

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, $R_2, \ldots, R_{k-1}$; $R_i \ge 2$ 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

[11] In 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 $a_1 f_1^N + b_1 = \ldots = a_j f_j^N + b_j \le b_{j+1}$ and $r = f_1^N + \ldots + f_j^N$. 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.

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 $\hat{\ell}(x) = \infty$ 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 $j_1 + 1, \ldots, j_t + 1$ 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 $C_{MN}(r) = \hat{C}_N^{R_1,\ldots,R_{k-1}}(r)$ 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 $B(R_1,\ldots,R_{k-1})$ which is smaller than $4/3$ for all choices of the R_i's and will prove $C_{MN}(r) \leq B(R_1,\ldots,R_{k-1})C_{opt}(r)$. We then investigate the non-benign case. We will derive a recurrence relation for

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

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* $R = 4$.

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* $\leq B(R_1, \ldots, R_{k-1}) < 4/3$; 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 (Sketch). 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

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

A tedious calculation shows that $F(r)$ is bounded by $2\Lambda_h/(\Lambda_i + \Lambda_h)$.

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

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 $\mu \leq 2$. □

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*

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

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 $i \le j$ and $L \le b_i = a_i f_i^N + b_i$ for $i > j$. Thus

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

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

The General Case: We come to the case where $\lambda_{i+1} \ge R_i \Lambda_i$ for some i. Let j be the smallest such i. For $r \le r_{j+1}/2$, MN and Opt use only links 1 to j and we are in the benign case. Hence *ePoA* is bounded by $B(R_1, \ldots, R_{j-1}) < 4/3$. 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 $s_h = r_{j+1}/2 + \sum_{j+1 \le i < h} (b_{i+1} - b_i)(\Lambda_i - \Lambda_j)$ and Opt starts to use it at $r_h/2 = r_{j+1}/2 + \sum_{j+1 \le i < h} (b_{i+1} - b_i)\Lambda_i/2$. We have $s_h \ge r_h/2$ since $\Lambda_i - \Lambda_j \ge \Lambda_i/2$ for $i > j$. □

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 $r_s = \sum_{i \le j} f_i^* \ge r_{j+1}/2$ 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 $t \le 1 + 1/R_j$ 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, $j+1 \le i \le k$, 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

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

The cost of MN on the high links is at most $ePoA(R_{j+1},\ldots,R_{k-1})$ times this cost by the induction hypothesis. We can now bound the cost of MN as follows:

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

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

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*

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

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 defined R_{k-1} down to R_{i+1} so that $ePoA(R_{i+1},\ldots,R_{k-1}) < 4/3$ and $(1+1/R_j)^2 ePoA(R_{j+1},\ldots,R_{k-1}) < 4/3$ for $j \geq i+1$. We next define R_i. We have

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

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

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* $= 1.192$.

5 A Lower Bound for the Case of Two Links

We prove that the construction of the previous section is optimal among the following class of deterministic mechanisms; mechanisms that consider Dafermos-Sparrow equilibria, and mechanisms that use non-decreasing[12] latency functions *even for the weaker notion of User Equilibria*. For the above mechanisms we show that *ePoA* ≥ 1.191. The proof of the following theorem does not make any assumptions on the monotonicity of the modified latencies.

Theorem 5. *The construction of Section 4 is optimal and ePoA* ≥ 1.191.

6 Open Problems

Clearly the ultimate goal is to design coordination mechanisms for general networks. Our mechanism approaches $4/3$, as k grows. Can we improve the upper bound for the case of k parallel links? 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. What can be said about atomic (weighted or unweighted) scheduling games and for games with polynomial latencies? What is the exact value of *ePoA* for the case of two parallel links? We conjecture that a reinspection of Sections 4 and 5 settles this question.

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

References

1. Angel, E., Bampis, E., Pascual, F.: Truthful algorithms for scheduling selfish tasks on parallel machines. Theor. Comput. Sci. 369(1-3), 157–168 (2006)
2. Angel, E., Bampis, E., Pascual, F., Tchetgnia, A.-A.: On truthfulness and approximation for scheduling selfish tasks. J. Scheduling 12(5), 437–445 (2009)
3. Azar, Y., Jain, K., Mirrokni, V.S.: (Almost) optimal coordination mechanisms for unrelated machine scheduling. In: SODA, pp. 323–332 (2008)
4. Bernstein, D., Smith, T.E.: Equilibria for networks with lower semicontinuous costs: With an application to congestion pricing. Transportation Science 28(3), 221–235 (1994)

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

5. Bonifaci, V., Salek, M., Schäfer, G.: On the efficiency of restricted tolls in network routing games. In: SAGT (2011) (to appear)
6. Caragiannis, I.: Efficient coordination mechanisms for unrelated machines scheduling. In: SODA, pp. 815–824 (2009)
7. Christodoulou, G., Gourvès, L., Pascual, F.: Scheduling selfish tasks: About the performance of truthful algorithms. In: Lin, G. (ed.) COCOON 2007. LNCS, vol. 4598, pp. 187–197. Springer, Heidelberg (2007)
8. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. Theor. Comput. Sci. 410(36), 3327–3336 (2009)
9. Cole, R., Correa, J.R., Gkatzelis, V., Mirrokni, V., Olver, N.: Inner product spaces for minsum coordination mechanisms. In: STOC (2011)
10. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: STOC, pp. 521–530 (2003)
11. Cole, R., Dodis, Y., Roughgarden, T.: How much can taxes help selfish routing? J. Comput. Syst. Sci. 72(3), 444–467 (2006)
12. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: A geometric approach to the Price of Anarchy in nonatomic congestion games. Games and Economic Behavior 64, 457–469 (2008)
13. Dafermos, S.C., Sparrow, F.T.: The traffic assignment problem for a general network. Journal of Research of the National Bureau of Standards, Series B 73B(2), 91–118 (1969)
14. Dafermos, S.: An extended traffic assignment model with applications to two-way traffic. Transportation Science 5, 366–389 (1971)
15. de Palma, A., Nesterov, Y.: Optimization formulations and static equilibrium in congested transportation networks. In: CORE Discussion Paper 9861, pp. 12–17. Université Catholique de Louvain, Louvain-la-Neuve (1998)
16. Fleischer, L.: Linear tolls suffice: New bounds and algorithms for tolls in single source networks. Theor. Comput. Sci. 348(2-3), 217–225 (2005)
17. Fleischer, L., Jain, K., Mahdian, M.: Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In: FOCS, pp. 277–285 (2004)
18. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.: Coordination mechanisms for selfish scheduling. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
19. Karakostas, G., Kolliopoulos, S.G.: Edge pricing of multicommodity networks for heterogeneous selfish users. In: FOCS 2004, pp. 268–276 (2004)
20. Karakostas, G., Kolliopoulos, S.G.: The efficiency of optimal taxes. In: López-Ortiz, A., Hamel, A.M. (eds.) CAAN 2004. LNCS, vol. 3405, pp. 3–12. Springer, Heidelberg (2005)
21. Kollias, K.: Non-preemptive coordination mechanisms for identical machine scheduling games. In: Shvartsman, A.A., Felber, P. (eds.) SIROCCO 2008. LNCS, vol. 5058, pp. 197–208. Springer, Heidelberg (2008)
22. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. Computer Science Review 3(2), 65–69 (2009)
23. Marcotte, P., Patriksson, M.: Traffic equilibrium. In: Transportation, Handbooks in Operations Research and Management Science. ch.10, vol. 14, pp. 623–713. North-Holland, Amsterdam (2007)
24. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
25. Patriksson, M.: The Traffic Assignment Problem: Models and Methods. V.S.P. Intl. Science (1994)
26. Roughgarden, T.: Designing networks for selfish users is hard. In: FOCS (2001)
27. Roughgarden, T., Tardos, É.: How bad is selfish routing? J.ACM 49, 236–259 (2002)
28. Wardrop, J.G.: Some theoretical aspects of road traffic research. In: Proceedings of the Institute of Civil Engineers, Part II, vol. 1, pp. 325–378 (1952)