Improving the Price of Anarchy via Predictions in Parallel-Link Networks

George Christodoulou
Aristotle University of Thessaloniki
Thessaloniki, Greece
Archimedes, Athena Research Center
Athens, Greece
gichristo@csd.auth.gr

Vasilis Christoforidis
Aristotle University of Thessaloniki
Thessaloniki, Greece
Archimedes, Athena Research Center
Athens, Greece
vgchrist@csd.auth.gr

Alkmini Sgouritsa
Athens University of Economics and Business
Athens, Greece
Archimedes, Athena Research Center
Athens, Greece
alkmini@aueb.gr

Ioannis Vlachos
Athens University of Economics and Business
Athens, Greece
Archimedes, Athena Research Center
Athens, Greece
ioa.vlahos@aueb.gr

Abstract

We study non-atomic congestion games on parallel-link networks with polynomial latencies. We investigate the power of machine-learned predictions in the design of *coordination mechanisms* aimed at minimizing the impact of selfishness. Our main results demonstrate that enhancing coordination mechanisms with simple advice on the input rate can optimize the social cost whenever the advice is accurate (*consistency*), while only incurring minimal losses even when the predictions are arbitrarily inaccurate (*bounded robustness*). Moreover, we provide a full characterization of consistent mechanisms, which holds for all monotone cost functions, and show that our proposed mechanism is optimal with respect to robustness. We further explore the notion of error-tolerance within this context, i.e., we provide an approximation guarantee that degrades smoothly as a function of the prediction error, up to a predetermined threshold, while achieving a bounded robustness.

CCS Concepts

• **Theory of computation → Algorithmic game theory**; **Network games**; **Quality of equilibria**.

Keywords

Congestion Games, Algorithms with Predictions, Price of Anarchy, Coordination Mechanisms

ACM Reference Format:
George Christodoulou, Vasilis Christoforidis, Alkmini Sgouritsa, and Ioannis Vlachos. 2026. Improving the Price of Anarchy via Predictions in Parallel-Link Networks. In *Proceedings of the ACM Web Conference 2026 (WWW '26), April 13–17, 2026, Dubai, United Arab Emirates.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3774904.3792480

ACM ISBN 979-8-4007-2307-0/2026/04
https://doi.org/10.1145/3774904.3792480

1 Introduction

A fundamental issue in large decentralized systems is the impact of selfish behavior, wherein each player selects a strategy to minimize their own cost without regard for the overall system efficiency. Such decentralized decision-making often leads to suboptimal outcomes for the population as a whole. In this context, the celebrated notion of the Price of Anarchy (PoA) [38] is the standard metric for quantifying the inefficiency arising due to players' selfish behavior. Specifically, it measures the ratio of the social cost in the worst Nash equilibrium to the optimal social cost. In this work, we consider congestion games, a fundamental resource allocation problem, from the perspective of anticipating players' strategic behavior in order to improve overall system performance.

Parallel-link networks serve as a fundamental model for studying congestion phenomena, where each link represents an independent resource and players choose among them. Analyzing such networks yields key insights into the performance of mechanisms in more complex networked environments. Moreover, this model captures a wide range of well-studied games and applications, including load balancing and scheduling games (e.g., load balancing in distributed systems, scheduling on cloud computing environments), where tasks must be assigned to machines with potentially different delays. More precisely, there is a set of available machines and a set of users, each associated with an infinitesimally small task that needs to be processed by a machine. Each machine is characterized by a cost function that depends on the congestion that appears on that machine. Each user selects a strategy, i.e., choosing a machine, by seeking to minimize their individual completion time.

To mitigate the impact of selfish behavior, we explore coordination mechanisms (CMs) [16, 17] as interventions to improve overall system performance. Additionally, motivated by the observation that additional information about the input can have a substantial impact on the quality of induced equilibria, we study the effect of exogenous information provided in the form of machine-learned predictions on the design of CMs aimed at steering the players towards more efficient outcomes. This approach of enhancing algorithmic decision-making with exogenous information has had a profound impact in recent algorithmic research [40, 43]. In line with this direction, closely related work by [28] incorporates predictions about anticipated demand in the design of learning-augmented cost-sharing mechanisms for scheduling games.

1.1 Our Results

We study learning-augmented CMs in network non-atomic congestion games over parallel-link networks with polynomial cost functions. Given a prediction of the total demand, we design *consistent* CMs that achieve the optimal performance with respect to the PoA when the prediction is accurate, while remaining near-optimal even under arbitrarily erroneous predictions; that is, they have bounded robustness. More precisely:

- In Section 3 we provide a thorough study of consistent CMs. First, we provide a complete characterization of consistent CMs for any monotone cost function (Lemma 3.1). We then propose a mechanism that applies the minimal necessary modification, which we call MinCharge, and show that it is consistent and also $(d + 1)$-robust for polynomial cost functions of maximum degree d with nonnegative coefficients (Theorem 3.4). We then use our characterization to show that any consistent mechanism must be at least $(d + 1)$-robust (Lemma 3.5), establishing the optimality of MinCharge among all consistent CMs.
- In Section 4 we study the trade-off between robustness and *error-tolerance* in the design of consistent CMs. Error-tolerance is a desirable property of learning-augmented algorithms, as it ensures that the performance of the mechanism degrades gracefully as a function of the prediction error up to a specified error-tolerance threshold. First, we extend the characterization of consistent CMs to capture error-tolerance (Lemma 4.1). We then propose the ErrorTolerant mechanism and analyze its approximation and robustness for multiple links with affine cost functions (Theorem 4.5). For the special case of two links, we also provide a tight approximation guarantee when the prediction error does not exceed the error-tolerance threshold.

1.2 Related Work

We consider the problem of routing traffic to improve network performance under congestion. In the absence of a central authority, users behave selfishly and seek to minimize their individual cost. Since equilibria generally do not minimize the social cost, we adopt the notion of the Price of Anarchy to quantify the performance degradation due to selfish behavior [38]. A vast literature has followed this notion and obtained tight bounds for a wide variety of games, e.g., [45, 46]. Considerable attention has been given specifically to networks of parallel links [1, 31]. [27] studied the problem of routing n users on m parallel machines. [15, 29, 47] considered resource-aware cost-sharing protocols for scheduling games.

Interventions. Coping with the effects of selfish behavior has been a central theme in algorithmic game theory. A prominent line of research explores taxation mechanisms and tolls to incentivize agents toward more efficient outcomes, either by incorporating them into the total disutility experienced by each player or by treating them as internal transfers within the system that can be refunded later. It is well-known that marginal taxes induce the optimal flow for the original cost functions as a Nash flow in the modified network (e.g. see [20, 26, 34]). In this work, we focus on the concept of CMs.[1] They were introduced by [16] to improve the efficiency of equilibria. In scheduling settings, CMs are implemented as local policies on each machine that dictate the order of processing tasks and/or introduce delays, and they have since been extended to a variety of settings [4, 11–13, 17, 19, 33, 36, 37].

We underline that a fundamental element in our analysis is the fact that CMs work for any given demand, i.e., they do not assume foreknowledge of the total demand in the network, as is required for marginal edge pricing. [21] examined whether optimal flows can be induced via demand-independent tolls for any travel demand, albeit without PoA guarantees. [35] showed that for affine cost functions and a fixed demand rate, enforcing the optimal flow with constant edge taxes results in a PoA of 2 when both the tax and the latency contribute to the social cost. Perhaps most closely related to our work is [17], which studies coordination mechanisms for parallel-link networks. They propose a mechanism in which the ratio of the cost of the resulting Nash flow to the original total latency is strictly less than 4/3 for parallel-link networks; they also provide a tight result for the case of two links, demonstrating the existence of a coordination mechanism that is 1.19-competitive. [48] obtained comparable improvements using alternative tax schemes.

In this work, we take a further step by augmenting the mechanism with predictions about the total demand expected to flow through the network. Our work is at the intersection of two main threads: improving the equilibrium performance in congestion games, and the use of exogenous information in the form of machine-learned predictions, commonly referred to as the *learning-augmented framework*.

Algorithms with Predictions. In recent years, machine-learned predictions have been widely adopted in algorithm design [40, 41, 43]. This approach aims to circumvent the limitations of worst-case analysis that has dominated algorithm analysis. More closely related to our work is the line of work on learning-augmented mechanisms in game-theoretic settings, specifically on strategic scheduling [6, 49], network formation games [28], facility location games [2, 5, 8, 49], mechanism design [3, 18, 22, 49], and auctions [7, 14, 30, 39].

2 Preliminaries

Setting. We focus on the well-studied selfish routing problem, which serves as a clear and illustrative example of non-atomic congestion games. We consider single-commodity congestion games on parallel-link networks, defined by a directed graph with only two vertices s, t representing the source and sink, respectively, and a set M of m links directed from s to t. Each link $i \in M$ comes with a non-decreasing, non-negative latency function $\ell_i(\cdot)$ that describes the travel time as a function of the amount of traffic on that particular link. We denote by r the units of traffic (demand) routed from s to t and call this the *input rate*. A flow $f = (f_i)_{i \in M}$ specifies for each link $i \in M$ the amount of flow passed through i, which is denoted by f_i; it holds that $\sum_{i \in M} f_i = r$. Then, $\ell_i(f_i)$ denotes the latency on link $i \in M$ under flow f. Naturally, we define the

[1] CMs can be seen as flow-dependent local taxes on each resource, where the cost functions are allowed to increase in an arbitrary manner.

total latency of f as $C(f) = \sum_{i\in M} f_i \ell_i(fi)$. For any input rate r, we denote by $C_{opt}(r)$ the cost of the optimal flow, i.e., the minimum cost over all feasible flows at rate r.

Price of Anarchy (PoA). A feasible flow f that routes r units of flow from s to t is a Nash (or Wardrop) equilibrium, or simply a Nash flow, if and only if for any two links $i, j \in M$ with $f_i > 0$, $\ell_i(f_i) \le \ell_j(f_j)$. That is, all used links (with positive flow) have equal cost under a Nash flow, which makes the Nash flows to be essentially unique in terms of the edge costs [9]. We use $C_N(r)$ to denote the cost of any Nash flow at input rate r. Then, the Price of Anarchy (PoA) for input rate r and the PoA for *any* rate are defined, respectively, as:

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

The tight 4/3 PoA bound for affine latencies is already attained in a simple network consisting of two parallel links with affine latency functions. This network is commonly referred to as Pigou's network and the corresponding PoA bound as Pigou's bound [46].

Latency functions. We consider latency functions that are polynomials of maximum degree $d \ge 1$ with nonnegative coefficients. Formally, for each link $i \in M$, $\ell_i(x) = \sum_{j=0}^{d} a_{i,j} x^j$, with $a_{i,j} \ge 0$ for all j. In addition, we treat the case of affine latency functions separately in Section 4. An affine latency function is of the form $\ell(x) = ax + b$, for $a, b \ge 0$. We write $\ell_i(x) = a_i x + b_i$ (with $a_i, b_i \ge 0$) for the latency function of link $i \in M$.

We remark that if two links i, j have the same constant term, i.e., $b_i = b_j$ then replacing them with one link of latency function $\ell(x) = \frac{a_i a_j}{a_i + a_j} x + b_i$ gives an equivalent network in the following sense: a flow f is an equilibrium flow (optimal flow, respectively) in the initial latencies if and only if it is an equilibrium (optimal) flow under the merged latencies by adding the respective link flows (see the full version for details). Therefore, without loss of generality, we assume that all constant terms are distinct, and we relabel links in increasing order of their constant term, i.e., $0 \le b_1 < b_2 < \ldots < b_m$. Moreover, we may assume that there is at most one link with $a_i = 0$, as in the case of multiple constant-latency links, only the one with the smallest b_i is used in any Nash or optimal flow.

Coordination Mechanisms (CMs). A coordination mechanism (CM) is defined as a set of modified latencies $\hat{\ell} = (\hat{\ell}_i)_{i\in M}$ such that, for any $i \in M$, $\hat{\ell}_i$ is strictly increasing[2] and $\hat{\ell}_i(x) \ge \ell_i(x)$ for all $x \ge 0$ [16, 17]. Let $\hat{C}_N(r)$ be the maximum cost of any Nash flow at rate r under the modified latency functions, i.e., the cost of the worst Nash flow when $\hat{\ell}$ is used rather than ℓ. Then, the engineered (modified) PoA, termed ePoA, for input rate r, and the ePoA for *any* rate are defined, respectively, as

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

Note that the optimum is taken with respect to the *original* latency functions, whereas the cost of the Nash flow includes any increments applied on the latency functions by the CM.

[2]With strictly increasing latencies, the CM admits a unique Nash flow in link loads.

Learning-augmented framework. We consider a prediction on the input rate, $\bar{r}$, that can be viewed as an estimate of the expected traffic that will flow in the network. Our goal is to design CMs that leverage the prediction to obtain improved ePoA bounds. Since the induced cost of a Nash flow and the ePoA of a CM now depends on the predicted input rate $\bar{r}$, we write:

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

When the prediction is correct, we want to induce the optimal routing with minimum cost for this input rate $\bar{r}$, without sacrificing worst-case guarantees, even if the prediction is arbitrarily wrong. Given $\bar{r}$, we say that the CM is *consistent* if it achieves $ePoA(\bar{r},\bar{r}) = 1$, for any $\bar{r} > 0$, and it is *ρ-robust* if it achieves $ePoA \le \rho$.

Prediction error - Error-tolerance. We use the prediction error to measure the inaccuracy of the prediction. Given a predicted input rate $\bar{r}$ and the actual rate r, the *prediction error* is given by $\eta(r,\bar{r}) = |r - \bar{r}|$.

A desirable property is that performance degrades gracefully as the error increases. Given an error-tolerance threshold $\bar{\eta} \ge 0$, we say that a CM is *error-tolerant* if for any $\bar{r} > 0$, the function $r \mapsto ePoA(r,\bar{r})$ is continuous for all $r \in [\bar{r} - \bar{\eta}, \bar{r} + \bar{\eta}]$, i.e., whenever $\eta(r,\bar{r}) \le \bar{\eta}$.

Equilibrium existence. We use the term Nash flow to refer to the equilibrium concept throughout the paper; for the connection between Nash flows and Wardrop equilibria, see [32]. In our setting, we allow latency functions to be discontinuous, since restricting to continuous modified latencies yields no improvement [17]. However, standard Wardrop equilibria need not always exist under discontinuous latencies. To address this, we adopt the notion of *user equilibrium* [10], which aligns with Wardrop's first principle: no individual commuter can unilaterally improve their travel time by switching routes. Formally, we say that a feasible flow f that routes r units of flow is a user equilibrium, if and only if for any two links $i, j \in M$ with $f_i > 0$,

$$\ell_i(f_i) \le \liminf_{\epsilon \downarrow 0} \ell_j(f + \epsilon \mathbf{1}_j - \epsilon \mathbf{1}_i),$$

where $\mathbf{1}_i$ denotes a unit of flow passing through link i. Under lower-semicontinuous, non-decreasing latencies, the existence of user equilibria is guaranteed [10]; henceforth, we refer to them simply as Nash flows. For a more thorough treatment, we refer the reader to [25, 42, 44].

3 A Learning-Augmented Coordination Mechanism for Parallel-Link Networks

In this section, we first derive a characterization of all consistent mechanisms (Lemma 3.1). We then define a natural class of parametrized mechanisms that satisfy the necessary and sufficient conditions of the characterization. We propose a mechanism within this class that applies the minimal required modification, and show that it is $(d+1)$-robust (Theorem 3.4). Finally, we establish a matching lower bound, proving that no consistent mechanism can be $(d+1-\epsilon)$-robust, for any $\epsilon > 0$ (Lemma 3.5).

3.1 Characterization of Consistent Mechanisms

We begin with the characterization of all consistent mechanisms. Notably, this result applies to all monotone latency functions, and not merely polynomial ones.

Lemma 3.1 (Characterization lemma). *Given any set of arbitrary monotone latency functions $(\ell_i)_{i\in M}$, a predicted input rate $\bar{r}$, and any CM, $(\hat{\ell}_i)_{i\in M}$, we denote by $\bar{f}^* = (\bar{f}_i^*)_{i\in M}$ the optimal flow for input rate $\bar{r}$ and we define the maximum latency on any used link under $\bar{f}^*$ as $L = L(\bar{r}) = \max_{i:\bar{f}_i^*>0} \ell_i(\bar{f}_i^*)$. Then, the CM is consistent, if and only if for every link i with $\bar{f}_i^* > 0$:*

- $\hat{\ell}_i(\bar{f}_i^*) = \ell_i(\bar{f}_i^*)$ *and*
- $\liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon) \geq L$.[3]

Proof. We first show that any CM that satisfies the two conditions of the statement is consistent. For this we consider the input rate $r = \bar{r}$ and show that the only Nash flow is $\bar{f}^*$. Suppose, for contradiction, that there exists a Nash flow $f = (f_i)_{i\in M} \neq \bar{f}^*$, meaning that there exists some link i with $0 \leq f_i < \bar{f}_i^*$. Then, there exists some link $j \neq i$ with $f_j > \bar{f}_j^* \geq 0$. If $\bar{f}_j^* > 0$, by the second condition we have $\hat{\ell}_j(f_j) \geq L$. If $\bar{f}_j^* = 0$, $\ell_j(0) \geq L$, otherwise the optimal flow would also route positive flow on j; therefore, it is also $\hat{\ell}_j(f_j) \geq \ell_j(0) \geq L$ in this case. Overall, it holds $\hat{\ell}_j(f_j) \geq L \geq \hat{\ell}_i(\bar{f}_i^*) \geq \hat{\ell}_i(f_i)$, where the second inequality comes from the definition of L. It cannot be $\hat{\ell}_j(f_j) > \hat{\ell}_i(f_i)$, since $f_j > 0$, and this would violate the fact that f is a Nash flow. Therefore, $\hat{\ell}_j(f_j) = \hat{\ell}_i(f_i) = L$, and by Footnote 3 it is $\hat{\ell}_j(f_j) > L$, which leads to a contradiction since it should be $\hat{\ell}_j(f_j) = \hat{\ell}_i(f_i)$.

Now we show that for any consistent CM the two lemma's conditions hold. Suppose again that the input rate is $r = \bar{r}$. Since all latency functions $(\ell_i)_{i\in M}$ are strictly increasing apart maybe from one, it is easy to check that the optimal flow $\bar{f}^*$ is unique. Therefore, in order for any CM to be consistent it needs to enforce $\bar{f}^*$ as the Nash flow and not to increase the latencies on that flow. This proves the necessity of the first condition in the lemma. Regarding the second condition, suppose for the sake of contradiction that for some link i with $\bar{f}_i^* > 0$, $\liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon) < L$. We show that $\bar{f}^*$ is not a Nash flow. Let $q = \arg\max_{j:\bar{f}_j^*>0} \ell_j(\bar{f}_j^*)$ be the link with the maximum latency under $\bar{f}^*$. Then, it holds that $\bar{f}_q^* > 0$, and $\ell_q(\bar{f}_q^*) = L > \liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon)$, which means that $\bar{f}^*$ is not a Nash flow and the CM is not consistent. Thus, the second condition in the lemma is also necessary. □

We define $q(\bar{r})$ for some predicted input rate $\bar{r}$ to be the link attaining the maximum latency under the optimal flow $\bar{f}^*$ for $\bar{r}$, i.e., $q(\bar{r}) = \arg\max_{j:\bar{f}_j^*>0} \ell_j(\bar{f}_j^*)$. We further define $k(\bar{r})$ as the number of links used in the optimal flow for input rate $\bar{r}$.

3.2 Constant Mechanisms

We define a class of parametrized CMs that we call Constant mechanisms. Given a parameter c and a predicted input rate $\bar{r}$, such a mechanism increases a link's latency to c for flow greater than the optimal flow under $\bar{r}$, provided that c is greater than the original latency. Every Constant CM is consistent when $c \geq L$, where L is as defined in Lemma 3.1. We then specify as MinCharge mechanism the Constant mechanism with parameter L, and show that it is optimal among all consistent CMs.

Definition 3.2 (Constant Mechanisms). Given any latency functions $(\ell_i)_{i\in M}$ and a predicted input rate $\bar{r}$, let $\bar{f}^* = (\bar{f}_i^*)_{i\in M}$ be the optimal flow for $\bar{r}$. For any parameter $c \geq \ell_i(\bar{f}_i^*)$, for all $i \in M$, a Constant mechanism is defined for $i \geq k(\bar{r})$ as $\hat{\ell}_i(x) = \ell_i(x), \forall x$, and for $i < k(\bar{r})$ as

$$\hat{\ell}_i(x) = \begin{cases} c + \delta_i(x), & \bar{f}_i^* < x < f_i^c \\ \ell_i(x), & \text{otherwise} \end{cases}$$

where $\delta_i(x)$ is an arbitrarily small value to ensure that the modified latency functions are strictly increasing, and f_i^c denotes the amount of flow such that $c + \delta_i(f_i^c) = \ell_i(f_i^c)$.[4]

Focusing on a specific instantiation of the Constant mechanisms, we define the following mechanism as the MinCharge mechanism.

Definition 3.3 (MinCharge Mechanism). Given any latency functions $(\ell_i)_{i\in M}$ and a predicted input rate $\bar{r}$, let $\bar{f}^* = (\bar{f}_i^*)_{i\in M}$ be the optimal flow for $\bar{r}$ and $L = \max_{i:\bar{f}_i^*>0} \ell_i(\bar{f}_i^*)$. The MinCharge mechanism is defined as the Constant Mechanisms with parameter L, or more analytically for $i \geq k(\bar{r})$ as $\hat{\ell}_i(x) = \ell_i(x), \forall x$, and for $i < k(\bar{r})$ as

$$\hat{\ell}_i(x) = \begin{cases} L + \delta_i(x), & \bar{f}_i^* < x < f_i^L \\ \ell_i(x), & \text{otherwise} \end{cases}$$

where $\delta_i(x)$ is an arbitrarily small value, and f_i^L is such that $L + \delta_i(f_i^L) = \ell_i(f_i^L)$.

Next, we provide a robustness guarantee for the MinCharge mechanism.

Theorem 3.4. *The MinCharge mechanism is consistent and $(d+1)$-robust on parallel-link networks with polynomial cost functions of maximum degree d with nonnegative coefficients.*

Proof. The MinCharge mechanism trivially satisfies the requirements of Lemma 3.1, hence it is consistent. We now analyze the behavior of the $\mathrm{ePoA}(r, \bar{r})$ as a function of the input rate r and the predicted input rate $\bar{r}$. We show that the worst-case inefficiency is realized for $r = \bar{r} + \epsilon$ for $\epsilon > 0$ that goes to zero. We distinguish between the cases of $r \leq \bar{r}$ and $r > \bar{r}$. If $r \leq \bar{r}$, we show that $ePoA(r, \bar{r}) \leq PoA(r) \in \Theta(d/\log d)$.

Case 1: Overprediction ($r \leq \bar{r}$). Suppose that f^* is the optimal flow under r with respect to the original latency functions and f denotes the Nash flow under r with respect to the modified latency functions. In the overprediction case, for each link i it holds that $f_i^* \leq \bar{f}_i^*$; thus, $\hat{\ell}_i(f_i^*) = \ell_i(f_i^*)$ and $\hat{\ell}_i(f_i) = \ell_i(f_i)$. Then, the variational inequality [24, 25] holds for f and f^*: $\sum_i \ell_i(f_i)f_i \leq \sum_i \ell_i(f_i)f_i^*$. Following the proof of the celebrated result on the PoA

[3]If the maximum latency link is a constant, then this inequality should be strict for all other links.

[4]Specifically, $\delta_i(x) = \frac{\varepsilon}{f_i^c - \bar{f}_i^*}(x - \bar{f}_i^*)$, where $\ell_i(f_i^c) = c + \varepsilon$. In the special case that $\ell_i(x)$ is constant, we set $\varepsilon = 0$. This technicality guarantees some nice properties: the CM admits a unique Nash flow and the link flows in the Nash flow are nondecreasing in the input rate [23].

of selfish routing due to [46], we obtain $ePoA(r, \bar{r}) \leq PoA(r) \in \Theta(d/\log d)$.

Case 2: Underprediction ($r > \bar{r}$). We now consider the second case, where $r > \bar{r}$. Suppose again that f^* and f are the optimal flow with respect to the original latencies and Nash flow w.r.t. the CM, respectively, both under r.

In the special case that $\ell_{q(\bar{r})}$ is a constant, it should be $\ell_{q(\bar{r})} = L$, and all the flow $r - \bar{r}$, will be routed through link $k(\bar{r})$. Then,

$$ePoA(r, \bar{r}) \leq \frac{rL}{C_{opt}(r)} .$$

On the other hand, if $\ell_{q(\bar{r})}$ is a strictly increasing function, then it follows that every ℓ_i is strictly increasing for $i \leq k(\bar{r})$ (see the full version for more details). In this case, every link incurs a latency higher than L under f, where L is the parameter of the MinCharge mechanism. To see this, note that for $r > \bar{r}$ there exists a link i with $f_i > \bar{f}_i^*$ and therefore, $\hat{\ell}_i(f_i) > L$; then, if there exists a link j with $\hat{\ell}_j(f_j) \leq L$, then there exists a sufficiently small $\epsilon > 0$ such that $\hat{\ell}_j(f_j + \epsilon) < \hat{\ell}_i(f_i)$, which violates the fact that f is a Nash flow.

Having established that the minimum cost in any link under r is more than L, we define r_{max} to be the maximum rate for which the latency on each link in the Nash flow does not exceed $L + \varepsilon$ (as defined in Definition 3.3). If $r > r_{max}$, all link flows in the Nash flow are beyond the modification point, i.e., $f_i > f_i^L$ for all i such that $\bar{f}_i^* > 0$, and so $\hat{\ell}_i(f_i) = \ell_i(f_i)$, for all $i \in M$. It follows that $\hat{C}_N(r, \bar{r}) = C_N(r)$. Hence, $ePoA(r, \bar{r}) = PoA(r)$. Therefore, we focus on the interval $r \in [\bar{r}, r_{max}]$.

In this regime, by omitting ε as it is arbitrarily small, the ePoA is given by

$$ePoA(r, \bar{r}) = \frac{rL}{C_{opt}(r)} .$$

which we analyze below; note that this is the same bound as in the case of $\ell_{q(\bar{r})} = L$. We first give the following claim to show that the ePoA is maximized when r approaches $\bar{r}$.

Claim 1. *$\frac{r}{C_{opt}(r)}$ is a decreasing function of r.*

The claim follows by convexity of the optimum social cost function $r \mapsto C_{opt}(r)$ (see e.g. [23]). The details are deferred to the full version. Under the MinCharge mechanism, the ePoA$(r, \bar{r})$ is maximized for $r = \lim_{\epsilon \to 0^+}(\bar{r} + \epsilon)$. The $ePoA(r, \bar{r})$ is defined as:

$$ePoA(r, \bar{r}) = \frac{\hat{C}_N(r, \bar{r})}{C_{opt}(r)} = \frac{rL}{C_{opt}(r)},$$

where L is the constant latency applied by the MinCharge CM. By Claim 1, $ePoA(r, \bar{r})$ of MinCharge is strictly decreasing for $r > \bar{r}$, and the maximum is attained in the limit as $r \to \bar{r}$. Therefore for $r \in [\bar{r}, r_{max}]$ the ePoA$(r, \bar{r})$ is as follows:

$$ePoA(r, \bar{r}) = \frac{\bar{r}L}{C_{opt}(\bar{r})} = \frac{\bar{r}L}{\sum_{i:\bar{f}_i^*>0} \bar{f}_i^* \ell_i(\bar{f}_i^*)} \qquad (1)$$

We next show that for any $i \leq k(\bar{r})$, which coincides with all links i with $\bar{f}_i^* > 0$, it holds $\ell_i(\bar{f}_i^*) \geq L/(d+1)$. Since $\bar{f}^*$ is an optimal flow, it holds that the marginal costs are equal in all links up to $k(\bar{r})$, and it is equal to $\ell_i(\bar{f}_i^*) + \bar{f}_i^* \ell_i'(\bar{f}_i^*) = \sum_{j=0}^{d} (j+1) a_{i,j} (\bar{f}_i^*)^j$. Then,

$$\ell_i(\bar{f}_i^*) = \sum_{j=0}^{d} a_{i,j} (\bar{f}_i^*)^j \geq \frac{1}{d+1} \sum_{j=0}^{d} (j+1) a_{i,j} (\bar{f}_i^*)^j \geq \frac{L}{d+1} ,$$

where for the last inequality, we used the definition of L. We plug in that $\ell_i(\bar{f}_i^*) \geq \frac{L}{d+1}$ for all i such that $\bar{f}_i^* > 0$, to bound the optimal flow under $\bar{r}$:

$$C_{opt}(\bar{r}) = \sum_{i:\bar{f}_i^*>0} \bar{f}_i^* \ell_i(\bar{f}_i^*) \geq \frac{L}{d+1} \sum_{i:\bar{f}_i^*>0} \bar{f}_i^* = \frac{L}{d+1} \bar{r} .$$

Substituting this into equation (1) we get:

$$ePoA(r, \bar{r}) \leq \frac{\bar{r}L}{\frac{L}{d+1}\bar{r}} = d + 1 . \qquad \square$$

We then show that the MinCharge CM is optimal with respect to robustness among all consistent CMs. For this we present an instance that precludes the existence of a consistent CM that is $(d + 1 - \epsilon)$-robust for any $\epsilon > 0$.

Lemma 3.5 (Lower bound). *Any consistent CM is at least $(d+1)$-robust for polynomial cost functions of maximum degree $d \geq 1$ with nonnegative coefficients.*

Proof. Fix the predicted input rate $\bar{r} = 1$. Consider two parallel links with latencies $\ell_1(x) = x^d$ and $\ell_2(x) = (d+1)(x+1)^d - \delta$, for a sufficiently small parameter $\delta \in (0, 1)$. We will eventually let $\delta \to 0$. Let $\bar{f}^*$ be the optimal flow for $\bar{r}$. The marginal costs are $m_1(x) = (d+1)x^d$ and $m_2(x) = (d+1)(x+1)^d - \delta + x \cdot d(d+1)(x+1)^{d-1} = (d+1)(x+1)^{d-1}((d+1)x+1) - \delta$. If both links are used, optimality implies $m_1(x) = m_2(1-x)$. Define $F(x) = m_1(x) - m_2(1-x)$. Since $m_1(x)$ is increasing in x and $m_2(1-x)$ is decreasing in x, it follows that F is strictly increasing in $[0, 1]$ and can have at most one root. Moreover, $F(0) = \Big((d+1)2^{d-1}(d+2) - \delta\Big) < 0$ and $F(1) = d + 1 - ((d+1) - \delta) = \delta > 0$. So by continuity F has a unique root in $(0, 1)$, and thus $\bar{f}_2^* > 0$. Moreover, by the Mean Value Theorem in $[\bar{f}_1^*, 1]$ for F, we get that there exists some $\xi \in (\bar{f}_1^*, 1)$ such that $F(1) - F(\bar{f}_1^*) = F'(\xi)(1 - \bar{f}_1^*)$, hence, $\delta = F'(\xi)\bar{f}_2^*$. Lastly, it holds that $F'(x) = m_1'(x) + m_2'(1-x) \geq d(d+1)$. Plugging this into the equation above, yields $\bar{f}_2^* \leq \delta/(d^2 + d) < \delta$, which goes to 0 as δ goes to 0.

Note that $\max_{i \in \{1,2\}} \ell_i(\bar{f}_i^*) = (d+1)(\bar{f}_2^* + 1)^d - \delta \geq (d+1)$ after letting $\delta \to 0$. By the characterization of consistent CMs (Lemma 3.1), the mechanism's modified latencies $\hat{\ell}_i$ must satisfy:

$$\hat{\ell}_i(\bar{f}_i^*) = \ell_i(\bar{f}_i^*), \text{ for } i \in \{1, 2\} ; \quad \liminf_{\epsilon \to 0^+} \hat{\ell}_1(\bar{f}_1^* + \epsilon) \geq (d+1) ,$$

Let f be a Nash flow of the CM, and let $r > \bar{r}$ denote the minimum input rate such that $\hat{\ell}_1(f_1) = \hat{\ell}_2(f_2)$. This happens precisely when link 1 in the Nash flow carries more than $\bar{f}_1^*$. Hence, $r = \bar{f}_1^* + f_2 = 1 - \bar{f}_2^* + f_2$, where $\bar{f}_2^* < \delta$. Let f^* be the optimal flow under r with respect to the original latency functions. Then:

$$ePoA(r, \bar{r}) = \frac{\hat{C}_N(r)}{C_{opt}(r)} = \frac{(1 + f_2)(d+1)(1 + f_2)^d}{C_{opt}(r)} ,$$

We can upper bound the optimal social cost at rate r by the cost of any feasible flow f, and hence, from the flow sending all flow through the first link, i.e., $C_{opt}(r) \le r\ell_1(r) = (1 + f_2)^{d+1}$. We obtain:

$$ePoA \ge d + 1 . \quad \square$$

The value for the modified latencies that minimizes the ePoA is precisely the one selected by the MinCharge mechanism. Combining Theorem 3.4 and Lemma 3.5, we obtain the following corollary.

Corollary 3.6. *The MinCharge CM is optimal with respect to robustness among all consistent CMs.*

4 Error-Tolerant Design

In this section, we extend our CMs to be *error-tolerant*, i.e., mechanisms whose approximation guarantee degrades gracefully as a function of the prediction error, up to a specified threshold. We derive a characterization of all consistent error-tolerant mechanisms (Lemma 4.1). We then focus on affine latencies; we derive an approximation guarantee as a function of the prediction error whenever it is upper bounded by the error-tolerance threshold, and show that this is tight for two links. Additionally, we establish a robustness bound in terms of $\bar{\eta}$ that the mechanism satisfies for arbitrary input rates. Combining the above, we derive our main result on error-tolerance in Theorem 4.5.

4.1 Characterization of Error-Tolerant Mechanisms

We provide a full characterization of consistent, error-tolerant CMs, extending Lemma 3.1, by requiring that, within the tolerance window around $\bar{r}$, equilibrium flows avoid discontinuities. This goes in addition to the consistency of the CM, meaning that the two conditions of Lemma 3.1 are still necessary. Therefore, if the optimal flow at $\bar{r}$ uses the first k links,[5] then $k-1$ links must have discontinuities beyond their optimal loads. Consequently, any additional demand must be routed along link $q(\bar{r})$ and links from $k+1$ to m.

Lemma 4.1 (Characterization lemma). *Given any latency functions $(\ell_i)_{i\in M}$, a predicted input rate $\bar{r}$, an error-tolerance threshold $\bar{\eta}$, and any CM, $(\hat{\ell}_i)_{i\in M}$, let $\bar{f}^* = (\bar{f}_i^*)_{i\in M}$ be the optimal flow for input rate $\bar{r}$. Suppose that $\bar{f}^{\bar{\eta}}$ is the Nash flow at input rate $\bar{f}^*_{q(\bar{r})} + \bar{\eta}$ by considering the network with links $k(\bar{r})+1$ through m with the modified latency functions $(\hat{\ell}_i)_{i\in\{k(\bar{r}),\dots m\}}$. We define as $\hat{L}_{\bar{\eta}}$ the maximum latency of any used link in $\bar{f}^{\bar{\eta}}$, and we further define $\bar{f}^{-\bar{\eta}}$ to be the Nash flow of the CM for input rate $\max\{\bar{r}-\bar{\eta}, 0\}$. Then, the CM is consistent and error-tolerant, if and only if:*

- *for every $i \le k(\bar{r})$, $\hat{\ell}_i(\bar{f}_i^*) = \ell_i(\bar{f}_i^*)$, and*
- *for every $i < k(\bar{r})$, $\liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon) \ge \hat{L}_{\bar{\eta}}$, and*
- *for every $i \le k(\bar{r})$, $\hat{\ell}_i(x)$ is continuous for $x \in [\bar{f}_i^{-\bar{\eta}}, \bar{f}_i^*]$, and for every $i \ge k$, $\hat{\ell}_i(x)$ is continuous for $x \in [\bar{f}_i^*, \bar{f}_i^{\bar{\eta}}]$.*

Proof. We first show that any CM satisfying the conditions of the statement is consistent. Since links $k+1$ to m have higher baseline latencies (constant terms $\ell_i(0)$) than $\ell_k(\bar{f}_k^*)$, otherwise they would be assigned a positive flow, it follows that $\hat{L}_{\bar{\eta}} \ge L$ (L as defined in Lemma 3.1). Therefore, any CM meeting the conditions of the lemma also satisfies the conditions of Lemma 3.1 and is consequently consistent. Next, we prove that any CM satisfying the conditions of the lemma is error-tolerant. We show that for any $r \in [\max\{\bar{r}-\bar{\eta}, 0\}, \bar{r}+\bar{\eta}]$, for any Nash flow f, no discontinuity appears, i.e., $f_i \le \bar{f}_i^*$ for $i < k$. Suppose on the contrary that there exists some $i < k$ with $f_i > \bar{f}_i^*$, meaning that $\hat{\ell}_i(f_i) \ge \hat{L}_{\bar{\eta}}$. Then, there exists some link $j < k$ with $f_j < \bar{f}_j^*$ or some link $j \ge k$ with $f_j < \bar{f}_j^{\bar{\eta}}$. In any case, $\hat{\ell}_j(f_j) < \hat{L}_{\bar{\eta}}$, meaning that f is not a Nash flow.

We now proceed to show that any CM that is consistent and error-tolerant satisfies the conditions of the lemma. The first condition is necessary for consistency similarly to the proof of Lemma 3.1. Regarding the second condition, suppose for the sake of contradiction that there exists a link $i < k$ such that $\liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon) < \hat{L}_{\bar{\eta}}$. We show that for the input rate $r = \bar{r} + \bar{\eta}$, any Nash flow f has $f_i > \bar{f}_i^*$, which would violate continuity within the tolerance window. Suppose for the sake of contradiction that f is a Nash flow that satisfies continuity, which means that the aggregated link flow for the first $k-1$ links is at most $\bar{r} - \bar{f}_k^*$. Therefore, the rest of the flow, which is at least $\bar{f}_k^* + \bar{\eta}$ is routed via links k through m under f. Based on the definition of $\hat{L}_{\bar{\eta}}$, there exists at least one link among them with positive flow and latency at least $\hat{L}_{\bar{\eta}}$. Since we assumed $\liminf_{\epsilon\to 0^+} \hat{\ell}_i(\bar{f}_i^* + \epsilon) < \hat{L}_{\bar{\eta}}$ for some $i < k$, f cannot be a Nash flow. Therefore, the second condition of the lemma's statement is also necessary. The third condition is trivially necessary for error-tolerance. $\square$

4.2 ErrorTolerant Mechanism

We extend MinCharge to achieve error-tolerance for affine latencies by introducing a the ErrorTolerant mechanism, which can be viewed as the direct analog of MinCharge, adapted to account for a bounded prediction error. Given an error-tolerance parameter $\bar{\eta}$, the CM sets a high latency for flow greater than the optimal flow under $\bar{r}$ to steer any additional flow towards links $k(\bar{r})$ through m, ensuring that no link incurs a steep increase in latencies. We instantiate the Constant mechanism by setting $c = L_{\bar{\eta}}$ which is the minimum latency on links $k(\bar{r})$ through m when the additional flow $\bar{\eta}$ is routed through them.

Definition 4.2 (ErrorTolerant Mechanism). Given any latency functions $(\ell_i)_{i\in M}$, a predicted input rate $\bar{r}$, and an error-tolerance threshold $\bar{\eta}$, let $\bar{f}^* = (\bar{f}_i^*)_{i\in M}$ be the optimal flow for $\bar{r}$. Suppose that $\bar{f}^{\bar{\eta}}$ is the Nash flow at input rate $\bar{f}^*_{k(\bar{r})} + \bar{\eta}$ by considering the network with links $k(\bar{r})$ through m, with $(\ell_i)_{i\in\{k(\bar{r}),\dots,m\}}$. We define as $L_{\bar{\eta}}$ the latency of any used link in $\bar{f}^{\bar{\eta}}$. The ErrorTolerant mechanism is defined for any $i \ge k(\bar{r})$ as $\hat{\ell}_i(x) = \ell_i(x), \forall x \ge 0$, and for any $i < k(\bar{r})$ as:

$$\hat{\ell}_i(x) = \begin{cases} L_{\bar{\eta}} + \frac{\varepsilon}{f_i^{L_{\bar{\eta}}} - \bar{f}_i^*}(x - \bar{f}_i^*), & \bar{f}_i^* < x < f_i^{L_{\bar{\eta}}} \\ \ell_i(x), & \text{otherwise} \end{cases}$$

where $\varepsilon > 0$ is an arbitrarily small value, and $f_i^{L_{\bar{\eta}}}$ is such that $\ell_i(f_i^{L_{\bar{\eta}}}) = L_{\bar{\eta}} + \varepsilon$.

We then proceed to bound the approximation guarantee of the ErrorTolerant CM as a function of the prediction error. In fact,

[5] We refer the reader to the full version for clarifications on which links are used in the optimal flow. We set $k = k(\bar{r})$ for simplicity.

we show that the analysis of the ErrorTolerant approximation guarantee is tight for any $\eta \leq \bar{\eta}$ in the case of two links. We generalize our results to networks with any number of links, deriving upper bounds on both the approximation and robustness guarantees.

4.3 Two Links

We begin with the case of two links. We establish a tight bound on the approximation guarantee when $\eta \leq \bar{\eta}$, and a tight bound on the robustness guarantee when $\eta > \bar{\eta}$ that the mechanism satisfies at all times for the case of two links.

We first establish that, when considering the Nash and the optimal flow over two links, it is without loss of generality to set the latency of the first link as $\ell_1(x) = x$, and that of the second link as $\ell_2(x) = ax + b$, for some constants $a, b \geq 0$. We also provide an auxiliary lemma describing the PoA of a two-link network as a function of r. We defer the details to the full version. Then, we proceed with the analysis of the approximation guarantee when $\eta \leq \bar{\eta}$, and the bound on the robustness guarantee when $\eta > \bar{\eta}$ that the mechanism satisfies at all times.

Lemma 4.3. *Consider any network G of two links with $\ell_1(x) = x$, and $\ell_2(x) = ax+b$, for some constants $a, b \geq 0$. Given a predicted input rate $\bar{r}$ and an optimal flow $\bar{f}^*$ under $\bar{r}$, let r' be the largest input rate such that the Nash flow on link 1 is less than $\bar{f}_1^*$, and $r'' = f_1^{L_{\bar{\eta}}} + f_2^{L_{\bar{\eta}}}$ (where $f_i^{L_{\bar{\eta}}}$ is as defined in the ErrorTolerant mechanism). Given an error-tolerance threshold $\bar{\eta}$, for any input rate r, let the actual error $\eta = r - \bar{r}$ (which we allow to be negative), with $|\eta| \leq \bar{\eta}$. Then, the ErrorTolerant mechanism achieves an approximation guarantee:*

- *for $r \leq r'$ and $r \geq r''$, $ePoA(r, \bar{r}) = PoA(r)$,*
- *for $r \in [r', \bar{r} + \bar{\eta}]$,*

$$ePoA(r, \bar{r}) = 1 + a - \frac{ab(a+1)(\frac{b}{4} + \eta)}{a\eta^2 + (a+1)b\eta + \frac{(a+1)b^2}{4}},$$

- *for $r \in [\bar{r} + \bar{\eta}, r'']$,*

$$ePoA(r, \bar{r}) = \frac{r\left(a\left(\frac{\left(\bar{r}-\frac{b}{2}\right)}{a+1} + \bar{\eta}\right) + b\right)}{r^2 - \frac{\left(r-\frac{b}{2}\right)^2}{a+1}}.$$

Proof. **Case (a)**: For $r \leq r'$ and $r \geq r''$ the modifications of the cost functions do not affect the ePoA, therefore $ePoA(r, \bar{r}) = PoA(r)$.

Case (b): It follows by definition of r' that there is enough capacity in both links such that any additional flow will not cause the social cost to increase abruptly. Define the difference δ as $\delta = f_1^* - \bar{f}_1^*$. It holds that $2f_1^* = 2a(r - f_1^*) + b = 2a(\bar{r} + \eta - f_1^*) + b$ and $2\bar{f}_1^* = 2a(\bar{r} - \bar{f}_1^*) + b$, therefore, $\delta = \frac{a\eta}{a+1}$.

Since $r \in [r', \bar{r} + \bar{\eta}]$, in the Nash flow $f_1 = \bar{f}_1^*$, therefore,

$$\begin{aligned}
ePoA(r, \bar{r}) &= \frac{\hat{C}_N(r)}{C_{opt}(r)} \\
&= \frac{\bar{f}_1^* \hat{\ell}_1(\bar{f}_1^*) + (r - \bar{f}_1^*)\hat{\ell}_2(r - \bar{f}_1^*)}{C_{opt}(r)} \\
&= \frac{(f_1^* - \delta)\ell_1(f_1^* - \delta) + (r - f_1^* + \delta)\ell_2(r - f_1^* + \delta)}{C_{opt}(r)} \\
&= \frac{f_1^*\ell_1(f_1^* - \delta) - \delta\ell_1(f_1^* - \delta)}{C_{opt}(r)} \\
&+ \frac{f_2^*\ell_2(r - f_1^* + \delta) + \delta\ell_2(r - f_1^* + \delta)}{C_{opt}(r)}.
\end{aligned}$$

Using the fact that the latencies satisfy $\ell_1(\bar{f}_1^*) = \ell_1(f_1^* - \delta) = \ell_1(f_1^*) - \delta$ and $\ell_2(r - \bar{f}_1^*) = \ell_2(r - f_1^* + \delta) = \ell_2(f_2^*) + a\delta$, we substitute these into the previous expression to obtain:

$$\begin{aligned}
ePoA(r, \bar{r}) &= \frac{f_1^*(\ell_1(f_1^*) - \delta) - \delta(\ell_1(f_1^*) - \delta)}{C_{opt}(r)} \\
&+ \frac{f_2^*(\ell_2(f_2^*) + a\delta) + \delta(\ell_2(f_2^*) + a\delta)}{C_{opt}(r)} \\
&= \frac{C_{opt}(r) + \delta(\ell_2(f_2^*) - \ell_1(f_1^*)) + \delta(af_2^* - f_1^*) + \delta^2(a+1)}{C_{opt}(r)}.
\end{aligned}$$

Since the optimal flows f^* satisfy $\ell_2(f_2^*) - \ell_1(f_1^*) = b/2$, we get:

$$\begin{aligned}
&= 1 + \frac{\delta^2(a+1)}{C_{opt}(r)} \\
&= 1 + \frac{a^2\eta^2}{(a+1)C_{opt}(r)}.
\end{aligned}$$

We have already established that $f_1^* - \bar{f}_1^* = a\eta/(a+1)$. Moreover, we have noted that that the optimal and Nash flows begin to diverge once $\bar{f}_1^* \geq b/2$ (meaning that if $\bar{f}_1^* < b/2$ the mechanism would not make any changes). Combining these observations, we have that $f_1^* = a\eta/(a+1) + \bar{f}_1^* \geq a\eta/(a+1) + b/2$. Similarly, for the second link, we have $f_2^* - \bar{f}_2^* = \eta/(a+1)$, and hence, $f_2^* \geq \eta/(a+1)$. We proceed to derive a lower bound on the optimal cost $C_{opt}(r)$ using the aforementioned inequalities.

$$\begin{aligned}
C_{opt}(r) &= f_1^{*2} + af_2^{*2} + bf_2^* \\
&\geq \frac{a^2\eta^2}{(a+1)^2} + \frac{a\eta b}{a+1} + \frac{b^2}{4} + \frac{a\eta^2}{(a+1)^2} + \frac{b\eta}{a+1} \\
&= \frac{a\eta^2 + b\eta + a\eta b}{a+1} + \frac{b^2}{4}.
\end{aligned}$$

Substituting, we have:

$$\begin{aligned}
ePoA(r, \bar{r}) &\leq 1 + \frac{a^2\eta^2}{(a+1)\left(\frac{a\eta^2 + b\eta + a\eta b}{a+1} + \frac{b^2}{4}\right)} \\
&= 1 + \frac{a^2\eta^2}{a\eta^2 + b\eta + a\eta b + (a+1)\frac{b^2}{4}} \\
&= 1 + \frac{a^2\eta^2 + a\eta b + a^2\eta b - a\eta b - a^2\eta b + a(a+1)\frac{b^2}{4} - a(a+1)\frac{b^2}{4}}{a\eta^2 + b\eta + a\eta b + (a+1)\frac{b^2}{4}} \\
&= 1 + a - \frac{ab(a+1)(\frac{b}{4} + \eta)}{a\eta^2 + (a+1)b\eta + \frac{(a+1)b^2}{4}}
\end{aligned}$$

For $\eta = 0$, we get that $ePoA(r, \bar{r}) = 1$, while for $\eta \rightarrow \infty$, $ePoA(r, \bar{r}) = 1 + a$.

This is tight because there exists an input rate $\bar{r}$ for which the inequalities are equalities. For $\bar{r} = b/2 + \epsilon$ for ϵ arbitrarily small, $\bar{f}_2^* = 0$, and therefore, the analysis is tight.

Case (c): In this case the cost of the Nash flow is approximately $\hat{C}_N(r, \bar{r}) = r\hat{L}_{\bar{\eta}} = r(a(\bar{f}_2^* + \bar{\eta}) + b)$ based on the definition of the ErrorTolerant mechanism. Since $\bar{f}^*$ is the optimal flow for $\bar{r}$, it holds that $2(\bar{r} - \bar{f}_2^*) = 2a\bar{f}_2^* + b$ and so

$$\hat{C}_N(r, \bar{r}) = r\left(a\left(\frac{\left(\bar{r} - \frac{b}{2}\right)}{a+1} + \bar{\eta}\right) + b\right).$$

The cost of the optimum flow is $C_{opt}(r) = r^2 - \frac{\left(r - \frac{b}{2}\right)^2}{a+1}$. Then, we derive,

$$ePoA(r, \bar{r}) = \frac{\hat{C}_N(r)}{C_{opt}(r)} = \frac{r\left(a\left(\frac{\left(\bar{r} - \frac{b}{2}\right)}{a+1} + \bar{\eta}\right) + b\right)}{r^2 - \frac{\left(r - \frac{b}{2}\right)^2}{a+1}}. \qquad \square$$

Lemma 4.4. *Consider any network G of two links with $\ell_1(x) = x$, and $\ell_2(x) = ax + b$, for some constants $a, b \geq 0$. Given any predicted input rate $\bar{r}$ and any error-tolerance threshold $\bar{\eta}$, the ErrorTolerant mechanism achieves the following robustness guarantee:*

$$ePoA = \max\{2, 1 + a\}.$$

Proof. **Upper bound.** Consider the third case of Lemma 4.3. The $ePoA(r, \bar{r})$ is given by:

$$ePoA(r, \bar{r}) = \frac{\hat{C}_N(r)}{C_{opt}(r)} = \frac{r\hat{L}_{\bar{\eta}}}{C_{opt}(r)}$$

Since $\hat{L}_{\bar{\eta}} = \left(a\left(\frac{\left(\bar{r} - \frac{b}{2}\right)}{a+1} + \bar{\eta}\right) + b\right)$ does not depend on r, by Claim 1, ePoA$(r, \bar{r})$ is a decreasing function of r and since $r \geq \bar{r} + \bar{\eta}$, we get:

$$ePoA(r, \bar{r}) \leq \frac{(\bar{r} + \bar{\eta})\left(a\left(\frac{\left(\bar{r} - \frac{b}{2}\right)}{a+1} + \bar{\eta}\right) + b\right)}{(\bar{r} + \bar{\eta})^2 - \frac{\left(\bar{r} + \bar{\eta} - \frac{b}{2}\right)^2}{a+1}}.$$

It is known that the optimal flow begins using link 2 at rate $b/2$ [17, 48]. The derivative of the above expression is non-positive, implying that the expression is maximized when $\bar{r} = b/2$. Therefore, we set $\bar{r}$ with $b/2$ and get:

$$\begin{aligned} ePoA(r, \bar{r}) &\leq \frac{(b + a\bar{\eta})(\frac{b}{2} + \bar{\eta})}{\left(\frac{b(a+1)+2a\bar{\eta}}{2(a+1)}\right)^2 + a\left(\frac{\bar{\eta}}{2(a+1)}\right)^2 + \frac{b\bar{\eta}}{2(a+1)}} \\ &= \frac{(b + a\bar{\eta})(\frac{b}{2} + \bar{\eta})}{\left(\frac{b}{2} + \frac{a\bar{\eta}}{a+1}\right)^2 + a\left(\frac{\bar{\eta}}{a+1}\right)^2 + \frac{b\bar{\eta}}{2(a+1)}} \\ &= \left(1 + \frac{a}{b}\bar{\eta}\right) \cdot \frac{2b + 4\bar{\eta}}{b + 4\bar{\eta} + \frac{4a}{b(a+1)}\bar{\eta}^2} \end{aligned}$$

To analyze the bound more effectively, we divide both the numerator and denominator of the right-hand side fraction by b, expressing the bound as a function of $u = \bar{\eta}/b \geq 0$:

$$ePoA(r, \bar{r}) \leq \frac{4au^2 + (4 + 2a)u + 2}{\frac{4a}{a+1}u^2 + 4u + 1}.$$

We inspect this function for $u \geq 0$. The expression evaluates to 2 at $u = 0$ and approaches $1 + a$ as $u \to \infty$. When $a \geq 2$, the function is increasing. When $a < 2$, the function is decreasing on the interval $[0, \frac{1-a+\sqrt{3-a}}{2a}]$ and increasing thereafter. Therefore, the worst-case $ePoA(r, \bar{r})$ is $\max\{2, 1 + a\}$, which is strictly worse than the 4/3 bound of the PoA. It is also strictly worse than the guarantee of case (b), since for $r \in [r', \bar{r} + \bar{\eta}]$, $ePoA(r, \bar{r})$ is less than $1 + a$, as the fraction in the expression of Lemma 4.3 is positive.

Hence, case (c) gives the worst case robustness guarantee.

Lower bound. The analysis of the ErrorTolerant robustness guarantee is tight in the case of two links. It suffices to set $\bar{r} = b/2$; then, we have $\bar{f}_1^* = \frac{2a\bar{r}+b}{2(a+1)} = \bar{r}$, and $\bar{f}_2^* = 0$. The ErrorTolerant CM dictates that:

$$\hat{\ell}_1(x) = \begin{cases} a\bar{\eta} + b, & b/2 \leq x < a\bar{\eta} + b \\ \ell_1(x), & \text{otherwise} \end{cases}$$

For $r = \bar{r} + \bar{\eta} = b/2 + \bar{\eta}$, we have $f_1^* = (2ar + b)/(2(a + 1)) = b/2 + a\bar{\eta}/(a+1)$ and correspondingly, $f_2^* = r - f_1^* = \bar{\eta}/(a+1)$. Then:

$$\begin{aligned} ePoA(r, \bar{r}) &= \frac{\hat{C}_N(r)}{C_{opt}(r)} \\ &= \frac{r\hat{L}_{\bar{\eta}}}{f_1^{*2} + af_2^{*2} + bf_2^*} \\ &= \frac{(b + a\bar{\eta})(\frac{b}{2} + \bar{\eta})}{\left(\frac{b}{2} + \frac{a\bar{\eta}}{a+1}\right)^2 + a\left(\frac{\bar{\eta}}{a+1}\right)^2 + \frac{b\bar{\eta}}{a+1}} \\ &= \left(1 + \frac{a}{b}\bar{\eta}\right) \cdot \frac{2b + 4\bar{\eta}}{b + 4\bar{\eta} + \frac{4a}{b(a+1)}\bar{\eta}^2} \qquad \square \end{aligned}$$

4.4 Many Links

We extend our results to networks with an arbitrary number of links, establishing upper bounds on the approximation and robustness guarantees. We include the statements here and defer the proofs of this section to the full version due to space constraints.

For simplicity of the presentation we set $k = k(\bar{r})$. Moreover, we define Λ_k and Γ_k as $\Lambda_k = \sum_{i \leq k} \frac{1}{a_i}$ and $\Gamma_k = \sum_{i \leq k} \frac{b_i}{a_i}$, respectively. Let $\bar{f}^*$ be the optimal flow under some predicted input rate $\bar{r}$.

Theorem 4.5. *For any error-tolerance threshold $\bar{\eta} > 0$, the ErrorTolerant coordination mechanism achieves an approximation of $ePoA(\eta) \leq 1 + \Lambda_k(a_{max} + a_k)\frac{\eta^2}{(\bar{r}+\eta)^2 + \Gamma_k(\bar{r}+\eta) - \Lambda_k C_k}$ if $\eta \leq \bar{\eta}$, and $2 + 2\bar{\eta}\frac{a_k}{b_k}$ otherwise, where η is the prediction error.*

Acknowledgments

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

References

[1] Daron Acemoglu and Asuman E. Ozdaglar. 2007. Competition and Efficiency in Congested Markets. *Math. Oper. Res.* 32, 1 (2007), 1–31.

[2] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. 2024. Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location. *Math. Oper. Res.* 49, 4 (2024), 2626–2651.

[3] Georgios Amanatidis, Evangelos Markakis, Christodoulos Santorinaios, Guido Schäfer, Panagiotis Tsamopoulos, and Artem Tsikiridis. 2025. Online Budget-Feasible Mechanism Design with Predictions. In *SAGT (Lecture Notes in Computer Science, Vol. 15953)*. Springer, 402–421.

[4] Yossi Azar, Lisa Fleischer, Kamal Jain, Vahab S. Mirrokni, and Zoya Svitkina. 2015. Optimal Coordination Mechanisms for Unrelated Machine Scheduling. *Oper. Res.* 63, 3 (2015), 489–500.

[5] Eric Balkanski, Vasilis Gkatzelis, and Golnoosh Shahkarami. 2024. Randomized Strategic Facility Location with Predictions. In *NeurIPS*.

[6] Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. 2023. Strategyproof Scheduling with Predictions. In *ITCS (LIPIcs, Vol. 251)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:22.

[7] Eric Balkanski, Vasilis Gkatzelis, Xizhi Tan, and Cherlin Zhu. 2024. Online Mechanism Design with Predictions. In *EC*. ACM, 1184.

[8] Zohar Barak, Anupam Gupta, and Inbal Talgam-Cohen. 2024. MAC Advice for facility location mechanism design. In *NeurIPS*.

[9] Martin Beckmann, Bartlett McGuire, and Christopher B. Winsten. 1956. *Studies in the economics of transportation.*

[10] David Bernstein and Tony E. Smith. 1994. Equilibria for Networks with Lower Semicontinuous Costs: With an Application to Congestion Pricing. *Transp. Sci.* 28, 3 (1994), 221–235.

[11] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. 2014. Coordination mechanisms from (almost) all scheduling policies. In *ITCS*. ACM, 121–134.

[12] Ioannis Caragiannis. 2013. Efficient Coordination Mechanisms for Unrelated Machine Scheduling. *Algorithmica* 66, 3 (2013), 512–540.

[13] Ioannis Caragiannis and Angelo Fanelli. 2019. An Almost Ideal Coordination Mechanism for Unrelated Machine Scheduling. *Theory Comput. Syst.* 63, 1 (2019), 114–127.

[14] Ioannis Caragiannis and Georgios Kalantzis. 2024. Randomized Learning-Augmented Auctions with Revenue Guarantees. In *IJCAI*. ijcai.org, 2687–2694.

[15] George Christodoulou, Vasilis Gkatzelis, and Alkmini Sgouritsa. 2024. Resource-Aware Cost-Sharing Methods for Scheduling Games. *Oper. Res.* 72, 1 (2024), 167–184.

[16] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. 2009. Coordination mechanisms. *Theor. Comput. Sci.* 410, 36 (2009), 3327–3336.

[17] Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. 2014. Improving the Price of Anarchy for Selfish Routing via Coordination Mechanisms. *Algorithmica* 69, 3 (2014), 619–640.

[18] George Christodoulou, Alkmini Sgouritsa, and Ioannis Vlachos. 2024. Mechanism design augmented with output advice. In *NeurIPS*.

[19] Richard Cole, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, and Neil Olver. 2015. Decentralized utilitarian mechanisms for scheduling games. *Games Econ. Behav.* 92 (2015), 306–326.

[20] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. 2003. Pricing network edges for heterogeneous selfish users. In *STOC*. ACM, 521–530.

[21] Riccardo Colini-Baldeschi, Max Klimm, and Marco Scarsini. 2018. Demand-Independent Optimal Tolls. In *ICALP (LIPIcs, Vol. 107)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 151:1–151:14.

[22] Riccardo Colini-Baldeschi, Sophie Klumper, Guido Schäfer, and Artem Tsikiridis. 2024. To Trust or Not to Trust: Assignment Mechanisms with Predictions in the Private Graph Model. In *EC*. ACM, 1134–1154.

[23] Roberto Cominetti, Valerio Dose, and Marco Scarsini. 2024. The price of anarchy in routing games as a function of the demand. *Math. Program.* 203, 1 (2024), 531–558.

[24] Stella Dafermos. 1980. Traffic Equilibrium and Variational Inequalities. *Transportation Science* 14, 1 (1980), 42–54.

[25] Andre De Palma, Yurii Nesterov, et al. 1998. Optimization formulations and static equilibrium in congested transportation networks. (1998).

[26] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. 2004. Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. In *FOCS*. IEEE Computer Society, 277–285.

[27] Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. 2010. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. *Theory Comput. Syst.* 47, 2 (2010), 405–432.

[28] Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. 2022. Improved Price of Anarchy via Predictions. In *EC*. ACM, 529–557.

[29] Vasilis Gkatzelis, Emmanouil Pountourakis, and Alkmini Sgouritsa. 2021. Resource-Aware Cost-Sharing Mechanisms with Priors. In *EC*. ACM, 541–559.

[30] Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. 2025. Clock Auctions Augmented with Unreliable Advice. In *SODA*. SIAM, 2629–2655.

[31] Tobias Harks, Marc Schröder, and Dries Vermeulen. 2019. Toll caps in privatized road networks. *Eur. J. Oper. Res.* 276, 3 (2019), 947–956.

[32] Alain Haurie and Patrice Marcotte. 1983. On the relationship between Nash - Cournot and Wardrop equilibria. *Networks* 15 (1983), 295–308.

[33] Nicole Immorlica, Li (Erran) Li, Vahab S. Mirrokni, and Andreas S. Schulz. 2009. Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.* 410, 17 (2009), 1589–1598.

[34] George Karakostas and Stavros G. Kolliopoulos. 2004. Edge Pricing of Multicommodity Networks for Heterogeneous Selfish Users. In *FOCS*. IEEE Computer Society, 268–276.

[35] George Karakostas and Stavros G. Kolliopoulos. 2004. The Efficiency of Optimal Taxes. In *CAAN (Lecture Notes in Computer Science, Vol. 3405)*. Springer, 3–12.

[36] Danish Kashaev. 2026. Selfish, Local and Online Scheduling via Vector Fitting. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 3841–3871.

[37] Konstantinos Kollias. 2013. Nonpreemptive Coordination Mechanisms for Identical Machines. *Theory Comput. Syst.* 53, 3 (2013), 424–440.

[38] Elias Koutsoupias and Christos H. Papadimitriou. 1999. Worst-case Equilibria. In *STACS (Lecture Notes in Computer Science, Vol. 1563)*. Springer, 404–413.

[39] Pinyan Lu, Zongqi Wan, and Jialin Zhang. 2024. Competitive Auctions with Imperfect Predictions. In *EC*. ACM, 1155–1183.

[40] Thodoris Lykouris and Sergei Vassilvitskii. 2021. Competitive Caching with Machine Learned Advice. *J. ACM* 68, 4 (2021), 24:1–24:25.

[41] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. 2007. Allocating online advertisement space with unreliable estimates. In *EC*. ACM, 288–294.

[42] Patrice Marcotte and Michael Patriksson. 2007. *Chapter 10 Traffic Equilibrium.* Vol. 14. 623–713. doi:10.1016/S0927-0507(06)14010-4

[43] Michael Mitzenmacher and Sergei Vassilvitskii. 2022. Algorithms with predictions. *Commun. ACM* 65, 7 (2022), 33–35.

[44] Michael Patriksson. 2015. *The traffic assignment problem: models and methods.* Courier Dover Publications.

[45] Tim Roughgarden. 2003. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.* 67, 2 (2003), 341–364.

[46] Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259.

[47] Philipp von Falkenhausen and Tobias Harks. 2013. Optimal Cost Sharing for Resource Selection Games. *Math. Oper. Res.* 38, 1 (2013), 184–208.

[48] Te-Li Wang, Chih-Kuan Yeh, and Ho-Lin Chen. 2016. An Improved Tax Scheme for Selfish Routing. In *ISAAC (LIPIcs, Vol. 64)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 61:1–61:12.

[49] Chenyang Xu and Pinyan Lu. 2022. Mechanism Design with Predictions. In *IJCAI*. ijcai.org, 571–577.