Nash Equilibrium and the Price of Anarchy in Priority Based Network Routing

Benjamin Grimmer, Sanjiv Kapoor

Abstract—We consider distributed network routing for networks that support differentiated services, where services are prioritized by a proportional weighting system. We use the classical Generalized Processor Sharing (GPS) scheme for scheduling traffic on network links. In such a scheme, each type of traffic is guaranteed a minimum capacity rate based on its priority.
To model the performance of this scheme and to account for autonomous routing we consider scheduling games on networks. We consider both networks with a set of parallel links (which also applies to processor scheduling) and more general scenarios where the network is a multi-graph. In each of these settings we consider two different routing schemes: Atomic and Non-Atomic. Atomic routing requires all traffic of one type to follow a single path. Non-Atomic routing splits traffic into a flow over multiple paths.
For each type of game, we prove either the existence of Nash Equilibrium or give a counterexample. We consider the inefficiency of equilibrium (termed as the price of anarchy) and provide price of anarchy upper bounds under reasonable assumptions. In general, this inefficiency in queuing systems is unbounded. We also provide complexity results on computing optimal solutions and the existence of equilibrium in these games.

Index Terms—Net Neutrality, Game Theory, Generalized Processor Sharing, Price of Anarchy

1 Introduction

THIS Differentiated services and prioritized traffic have been suggested as a means to offer certain Quality of Service (QoS) requirements. These requirements may vary based on the type of traffic. Certain traffic may be very dependent on bandwidth, delay, jitter, or loss rate. This is the case with many types of multimedia traffic (e.g. video or audio traffic requires low latency). Differentiated services may also be offered for economic reasons and the implications of using these services are also relevant to the debate on Net Neutrality.

One of the problems that arises from this is to devise methods of guaranteeing these requirements for each class of traffic. Different scheduling schemes can be used to guarantee certain properties for each type of traffic. We are particularly interested in one such mechanism, Generalized Processor Sharing (GPS), which was introduced by [1] and has attracted much interest [2]. GPS scheduling allows us to guarantee each type of traffic will be provided a minimum portion of a link's bandwidth.

Large communication networks (such as the Internet) are impractical to manage from a centralized source. Realistically, routing decisions must be distributed across the network. Then it is reasonable to view each data source as an independent rational agent in a non-cooperative routing game [3]. All of the routing agents attempt to optimally route their traffic based on some performance metric. For the games we consider, this metric will be to minimize the delay experienced by their data.

We will investigate two different types of routing: Atomic and Non-Atomic. These are formally defined in Section 1.3, but in essence Atomic routing requires all of a player's data to follow a single path whereas Non-Atomic routing allows traffic to be split into a flow over multiple paths. We assume our players will minimize the delay of their worst path. This behavior occurs in networks when traffic is allowed to take multiple paths and follows ECMP (Equal Cost Multi-Path routing). We also consider the special case where our network is a set of parallel links in addition to the general case of a multigraph. When the game's network is limited to parallel links, it can be used to model processor scheduling. By combining these two distinctions, we have four different games of interest: *GPS Atomic Game on Parallel Links*, *GPS Atomic Game on a Graph*, *GPS Non-Atomic Game on Parallel Links*, and *GPS Non-Atomic Game on a Graph*.

Of particular interest in this type of system are equilibrium points where all players are satisfied with their current routing decisions. These stable points are called Nash Equilibrium (NE). Stable points have been studied thoroughly in the context of many different economic and network routing models in [4], [5], and [6]. Different properties of Nash Equilibrium are useful when evaluating the quality of a system. One standard quantity is the ratio between the total delay of the worst NE to the total delay of the optimal assignment. This is known as the *Price of Anarchy* (PoA). Having a low upper bound on this value helps justify that the system is efficient when players act autonomously. It has already been shown that similar routing games can have arbitrarily large PoA [7]. One previous work has considered PoA in the context of GPS-Scheduling [8]. They only consider Non-Atomic Routing on Parallel Links and prove upper bounds on the PoA in the specific case where each type of traffic has high demand. Our results comprehensively improve this previous work, generalizing it to arbitrary networks.

We also consider the complexity of computing optimal

- *Department of Computer Science,*
  Illinois Institute of Technology, Chicago, IL, 60616.
  E-mail: bgrimmer@hawk.iit.edu
  E-mail: kapoor@iit.edu

solutions and equilibrium assignments in our games. Previous works have been able to classify computing NE for different games as being polynomial time, PPAD-complete, PLS-complete, and NP-complete (for example, [5] and [9]).

1.1 Our Contributions

We investigate the effectiveness of GPS scheduling on our models of network routing games. In each type of game, we consider the existence of equilibrium, the price of anarchy, the computational complexity of related problems, and simulations of average behavior. Our results are summarized as follows:

- In Section 2, we prove that a GPS Atomic Game on Parallel Links will always have a Nash Equilibrium. Further, we give counterexamples in our three remaining games where no NE exists. Thus in these games there in no guarantee that players will eventually stop rerouting their traffic.
- In Section 3, we prove upper bounds on the rate of growth of the PoA in our atomic and non-atomic games. All four of the games we study could have arbitrarily large PoA. Further, we give examples of atomic games where the price of anarchy exactly matches our bound in Parallel Link networks and is asymptotically tight for multigraphs.
- In Section 4, we prove that it is NP-Complete to find an optimal assignment minimizing total delay. We give a quick algorithm that computes a NE for a GPS Atomic Game on Parallel Links. Further, we prove that computing whether a NE exists in any of the three remaining games is NP-Complete.
- In Section 5, we give experimental results from randomly generated parallel links networks. We find average price of anarchy values to contrast our tight worst case bounds from Section 3. We also give results showing the average number times traffic will reroute before the network reaches equilibrium.

1.2 Related Works

The efficiency of NE in scheduling and routing has been considered actively by many researchers. Papadimitriou considered the special case when the network has two nodes and many parallel links [10]. Improving on this, Roughgarden et al. in [4] studied this problem for a general network with continuous and monotone cost functions on each link. They were able to bound the PoA tightly in such a system using a "Pigou" value, which only depends on the link cost functions (i.e. it is independent of network topology). In this special case of linear cost functions, the Pigou bound is equal to 4/3. These results are further studied in [11], [12], [13], and [14].

In this paper, we are particularly interested in routing games where the delay is modeled via queuing delay functions. The first work to study this type of delay function in resource allocation games was [15]. Although this work analyzed both socially optimal and NE assignments, it did not investigate the PoA. Later, in [16], it was shown that the PoA is unbounded in a general network. Further, it has been shown that in networks where the maximum amount of traffic is less than the capacity of the smallest link, the PoA is independent of network topology [4]. In the case of M/M/1 servers, the PoA in Non-Atomic Games network is tightly bounded by the number of backend servers as shown in [17] extending [15]. Existence of Nash equilibria has been proven for general product form network where the metric is the average throughput [18]. The queuing models used in the paper include PS (processor sharing). However the inefficiency of the game considered has not been considered.

A more realistic model would allow differentiation between types of traffic. In [19] extending the results of [17], Altman et al. consider traffic with different packet sizes. They show that even for a graph with two nodes and two links the PoA is unbounded. However, this model of differentiated traffic does not guarantee any Quality of Service requirements to different classes of data. Another differentiated routing game is proposed in [20], which considers an atomic dynamic routing system. In their model, servers follow discrete queuing behavior and process traffic in a first-come first-server manner, where ties are broken by the traffic's priority. They prove PoA is bounded by the number of users and show computing social optimal and NE assignments are both NP-complete.

1.3 Network Model

In general, we represent our network as a graph $G = (V, E)$. We will consider two different network topologies, where G is a two node graph consisting of a set of parallel links, and where G can be any multigraph. Each link in our network has a capacity μ_e. Each link is treated as two directed edges that operate independently and have identical capacities. The traffic in one direction on e and its capacity determine the delay $d_{i,e}$ experienced by traffic of type i using the link. In Section 1.4, we will formally define the delay function $d_{i,e}$ based each link's scheduling scheme.

Now we can define our routing game over the network. Let K represent the set of types of traffic. Each type of traffic $i \in K$ has a source node s_i and sink node t_i, and an amount of traffic λ_i that must be routed from s_i to t_i. Each type of traffic is also associated with a priority ϕ_i that will be used in our scheduling scheme (without loss of generality, we normalize these values: $\sum \phi_i = 1$). Then a game is defined by the following tuple $(G, \mu, K, \lambda, \phi)$, where $G = (V, E)$ is the network, $\mu : E \to \mathbb{R}$ is the capacity function, K is the set of traffic types (each type is defined by a pair of vertices s_i and t_i), $\lambda : K \to \mathbb{R}$ is the demand function, and $\phi : K \to \mathbb{R}$ is the priority function.

We consider two different models of assigning traffic to paths in G: Atomic Routing and Non-Atomic Routing. In Atomic Routing, each type of traffic is a player in our game. Each player must send all their traffic on a single s_i, t_i-path. The player will choose to route their traffic on the path with minimum total delay. In Non-Atomic Routing, each type of traffic is split into an infinitesimal flow. Then each unit of flow can be treated as a player and greedily routes itself from s_i to t_i. We say that traffic of type i is *stable* in an assignment S if each path being used by traffic of type i in S has minimum total delay. This corresponds to ECMP routing. Let $\lambda_{i,e}$ be the amount of traffic of type i on link e. Then the social objective (total delay) in both types of games is $\sum_{i\in K}\sum_{e\in E}\lambda_{i,e}d_{i,e}$. These two variations of routing rules along with the two types of allowed networks generate our four different games.

1.4 GPS-Scheduling Model

Each link in our network implements a GPS scheduler giving a capacity guarantee to each type of traffic. Regardless of other traffic present on link e, traffic of type i will have at least $\phi_i\mu_e$ capacity available to it. For an assignment of traffic to paths S, we let $S(e)$ denote the set of all types of traffic using link e. For ease of notation, we will define $C_{S(e)} = \sum_{j\in S(e)}\phi_j$. Then according to the GPS scheduling algorithm, the capacity perceived by traffic of type i on link e is the following [1]:

$$\tilde{\mu}_{i,e} = \frac{\phi_i}{C_{S(e)}}\mu_e$$

We assume that traffic will arrive at links according to a Poisson process, and packets can be handled according to an exponential distribution of service times. So we have a M/M/1 system. Note that the output of such a server is also Poisson, so indeed all traffic in the network will be Poissonian. For a standard First-Come-First-Serve (FCFS) server with capacity μ and traffic $\lambda < \mu$, the expected delay is $d = 1/(\mu - \lambda)$.

Our M/M/1-GPS queue handles each type of traffic separately. Thus we can split one such server into a set of parallel M/M/1-FCFS queues, as done in [21]. Then we can view our entire system of priority based queues as FCFS queues using the modified capacity $\tilde{\mu}_{i,e}$. When $\lambda_{i,e}$ traffic of type i is routed on link e, the delay experienced by traffic of type i on link e is the following:

$$d_{i,e} = \begin{cases} \frac{1}{\tilde{\mu}_{i,e}-\lambda_{i,e}} & \text{if } \tilde{\mu}_{i,e} > \lambda_{i,e} \\ +\infty & \text{otherwise} \end{cases}$$

In the case that the traffic exceeds the available capacity, the player is essentially unable to send that amount of traffic on the link. We say that player i *starves* on link e if $\lambda_{i,e} \geq \tilde{\mu}_{i,e}$.

2 Existence of Nash Equilibrium

2.1 Atomic Games

First, we will show that a GPS Atomic Game on Parallel Links will always have a NE. We prove this by giving a potential function that decreases whenever a player makes an improving move. Suppose player i improved its total delay by moving from link e to link e'. Let S be the player assignments before this move and S' be the player assignments afterwards. Then the following inequality must hold:

$$\frac{1}{\frac{\mu_e\phi_i}{C_{S(e)}} - \lambda_i} > \frac{1}{\frac{\mu_{e'}\phi_i}{C_{S'(e')}} - \lambda_i} \tag{1}$$

$$\implies \frac{\mu_e\phi_i}{C_{S(e)}} < \frac{\mu_{e'}\phi_i}{C_{S'(e')}} \tag{2}$$

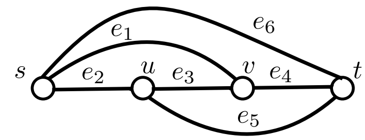

Fig. 1. The generalized link network given in Theorem 2.

Thus, such a move can be characterized by the following *variational inequality*.

$$\frac{C_{S(e)}}{\mu_e} > \frac{C_{S'(e')}}{\mu_{e'}} \tag{3}$$

Theorem 1. *Every GPS Atomic Game on Parallel Links instance (G,μ,K,λ,ϕ) has a Nash Equilibrium.*

Proof. We consider the following potential function:

$$\Phi(S) = \sum_{e\in E}\frac{C^2_{S(e)}}{\mu_e} - \sum_{i\in K}\frac{\phi_i^2}{\mu_{e_i}}$$

Now we consider the change in this function during the move from S to S': $\Phi(S') - \Phi(S)$

$$= \frac{C^2_{S'(e')}}{\mu_{e'}} - \frac{C^2_{S(e')}}{\mu_{e'}} + \frac{C^2_{S'(e)}}{\mu_e} - \frac{C^2_{S(e)}}{\mu_e} - \frac{\phi_i^2}{\mu_{e'}} + \frac{\phi_i^2}{\mu_e}$$
$$= \frac{2C_{S'(e')}\phi_i + \phi_i^2}{\mu_{e'}} - \frac{2C_{S(e)}\phi_i + \phi_i^2}{\mu_e} - \frac{\phi_i^2}{\mu_{e'}} + \frac{\phi_i^2}{\mu_e}$$
$$= \frac{2C_{S'(e')}\phi_i}{\mu_{e'}} - \frac{2C_{S(e)}\phi_i}{\mu_e} < 0$$

Therefore the potential function decreases when a player makes an improving move. As a result, the minimum value of this function will correspond to a NE. ☐

Next we consider the more general problem that allows the network to be any multigraph. In this case, we are able to give an example network where no NE assignment exists.

We introduce a new notation to simplify the description of our counter-example. A *Generalized Link* is a series of links with identical capacities forming a path. We can combine such a path of length α into a single generalized link with delay:

$$d_{i,e} = \begin{cases} \frac{\alpha}{\tilde{\mu}_{i,e}-\lambda_{i,e}} & \text{if } \tilde{\mu}_{i,e} > \lambda_{i,e} \\ +\infty & \text{otherwise} \end{cases}$$

Theorem 2. *A GPS Atomic Game on a Graph instance (G,μ,K,λ,ϕ) with no Nash Equilibrium exists.*

Proof. We consider a game with two players. Let $\phi_1 = \phi_2 = 1/2$, and $\lambda_1 = 5$ and $\lambda_2 = 0.0001$. Let G be the graph on four vertices $\{s,u,v,t\}$ with the following six generalized edges (shown in Figure 1):

$$\begin{aligned}
e_1 &= sv, & \mu_{e_1} &= 0.25, & \alpha_{e_1} &= 0.01 \\
e_2 &= su, & \mu_{e_2} &= 900, & \alpha_{e_2} &= 100 \\
e_3 &= uv, & \mu_{e_3} &= 10, & \alpha_{e_3} &= 1 \\
e_4 &= vt, & \mu_{e_4} &= 11, & \alpha_{e_4} &= 1 \\
e_5 &= ut, & \mu_{e_5} &= 1, & \alpha_{e_5} &= 0.01 \\
e_6 &= st, & \mu_{e_6} &= 20, & \alpha_{e_6} &= 10
\end{aligned}$$

Player 1 has traffic larger than the capacity on the links sv and ut. Thus player 1 can only use the two paths: $p_1 =$

s, u, v, t and $p_2 = s, t$. Player 2 can use any of the five paths from s to t. By enumerating these ten cases, one can verify that this game has no equilibrium assignment. □

2.2 Non-Atomic Games

Theorem 3. *A GPS Non-Atomic Game on Parallel Links instance with three players and three links that has no Nash Equilibrium exists.*

Proof. In a GPS Non-Atomic Game, the delay function for player i depends on whether $\lambda_{j,e}$ is non-zero, but doesn't depend on its value for each other player j. This allows us to view the other player's assignments as boolean information when computing the best responses. Using a lemma equivalent to the one proven by [17], Monsef, Anjali and Kapoor have shown how to compute the NE of traffic of a single type if all other traffic is fixed as follows [8]:

Lemma 1 ([8]). *Given all traffic types other than i are fixed. Sort the links in decreasing order of $\tilde{\mu}_{i,e}$. Then there exists a unique NE for the traffic of type i that exactly uses the first m_i links, where*

$$m_i = \min \left\{ m \geq 1 : \tilde{\mu}_{i,m+1} \leq \frac{\sum_{e=1}^{m} \tilde{\mu}_{i,e} - \lambda_i}{m} \right\}$$

Now we can view the GPS Non-Atomic Game as a discrete game where each type of traffic is a player with best response defined by Lemma 1. Despite the game being discrete, the strategy sets of each player are still large. If there are n links, then there are $2^n - 1$ non-empty subsets of edges a player can chose to use. Even for a game with three players and three links, there are $7^3 = 343$ possible assignments. This issue makes it increasingly difficult to construct a counter example with no NE by hand.

To remedy this, we had a program randomly sample the nine game parameters (ϕ_i, λ_i, and μ_e) and check if any of the 343 assignments are a NE. We sample each ϕ_i and λ_i uniform randomly from $[0, 1]$ and μ_e uniform randomly from $[0, 2]$. Then we normalize the value of ϕ. Our program was able to find a counter example, but we found that such instances are extremely rare. After running 10^8 random samples, we found with 95% confidence that counter examples make up $8.24 \times 10^{-6} \pm 5.626 \times 10^{-7}$ proportion of our sample space. One counter example found by our program is given below, completing our proof of Theorem 3:

$$\begin{array}{lll} \phi_1 = 0.1, & \phi_2 = 0.48, & \phi_3 = 0.42 \\ \lambda_1 = 0.2, & \lambda_2 = 0.5, & \lambda_3 = 0.28 \\ \mu_1 = 0.6, & \mu_2 = 1.27, & \mu_3 = 0.3 \end{array}$$

□

In the case of a GPS Non-Atomic Game on Parallel Links with exactly two players, we can prove that a NE will always exist. We sketch the proof of this result next.

Theorem 4. *Every GPS Non-Atomic Game on Parallel Links instance $(G, \mu, K, \lambda, \phi)$ with exactly two players has a Nash Equilibrium.*

Proof. We show that a NE must exist by giving a direct construction: Assign player 1 to use all links. Then assign player 2 to its best response given player 1's current strategy. Then we alternate each player making their best response move until we reach a NE. We claim that each iteration of this algorithm the total delay of both types of traffic is strictly decreasing. This can be shown through an analysis of cases. Since the total delay is bounded below, the procedure must eventually terminate with an equilibrium. □

3 Price of Anarchy

In this Section, we give an upper bound on the Price of Anarchy for our routing games. We prove our bounds for the routing problems on multi-graph topologies, and then bounds on the PoA for Parallel Links will follow. We give examples showing that our bound is exactly tight for GPS Atomic Games on Parallel Links, and that the bound for GPS Atomic Games on a Graph is asymptotically tight. Our upper bounds are based on the following quantity:

$$\beta = \sup \left\{ \frac{\sum_{e \in E} \lambda_{i,e}(S) d_{i,e}(S)}{\sum_{e \in E} \lambda_{i,e}(S^*) d_{i,e}(S^*)} : S \text{ is a NE, any } S^*, i \in K \right\}$$

Then simple algebra can show the following relationship between β and the PoA.

Lemma 2. *The Price of Anarchy for a GPS Non-Atomic Game on a Graph is at most β.*

Proof. Let λ^* and d^* correspond to the optimal assignment for our game. Similarly, let λ and d correspond to the worst case NE in our game.

$$|OPT| = \sum_{i \in K} \sum_{e \in E} \lambda^*_{i,e} d^*_{i,e} \geq \frac{1}{\beta} \sum_{i \in K} \sum_{e \in E} \lambda_{i,e} d_{i,e}$$

□

3.1 Price of Anarchy in Atomic Games

For an atomic game, we use the shorthand $d_i(S)$ as the sum of the delay on each link used by player i in assignment S. Then we can simplify the value of β as follows in the case of atomic games:

$$\beta = \sup \left\{ \frac{d_i(S)}{d_i(S^*)} : S \text{ is a NE, any } S^*, i \in K \right\}$$

If there is a NE in the game where a player cannot have their entire demand met, then the price of anarchy will be infinite. Then for all our PoA bounds, we need to ensure the network has enough surplus capacity to prevent such starvation. We represent this using a surplus factor C. For each player i, we consider the s_i, t_i-path p_i with the largest minimum capacity among its links. Let μ^i_{min} be the minimum capacity along p_i. Then to prevent starvation, we need $\phi_i \mu^i_{min} > \lambda_i$. This ensures that player i can always route their traffic along p_i. Then our surplus factor C is the following:

$$C = \min_{i \in K} \left\{ 1 - \frac{\lambda_i}{\phi_i \mu^i_{min}} \right\}$$

Theorem 5. *Assuming no player starves (i.e. $C > 0$), in a GPS Atomic Game on Graph topology, the Price of Anarchy has the following upperbound:*

$$PoA \leq \Big(|V| - 1\Big)\Big(1 + \frac{1 - \phi_{min}}{C\phi_{min}}\Big)$$

Proof. Let S^* be any assignment and let S the worst case NE. We can guarantee that each $d_i(S)$ is at most $|p_i|/(\phi_i \mu^i_{min} - \lambda_i)$ since player i is making a best response move and has a strategy with at most this delay available. Further, we can bound $d_i(S^*)$ below by $1/(\mu^i_{min} - \lambda_i)$ since the path used in S^* must have a link with capacity at most μ^i_{min}. Applying these two bounds on the definition of β (and that $|p_i| \leq |V| - 1$) yields the following PoA bound:

$$\beta \leq \sup_{i \in K} \frac{|p_i|(\mu^i_{min} - \lambda_i)}{\phi_i \mu^i_{min} - \lambda_i} \tag{4}$$

$$= \sup_{i \in K} |p_i| \Big(1 + \frac{\mu^i_{min} - \phi_i \mu^i_{min}}{\phi_i \mu^i_{min} - \lambda_i}\Big) \tag{5}$$

$$\leq \sup_{i \in K} \Big(|V| - 1\Big)\Big(1 + \frac{1 - \phi_i}{C\phi_i}\Big) \tag{6}$$

Then the Theorem follows from Lemma 2. □

Corollary 1. *Assuming no player starves (i.e. $C > 0$), in a GPS Atomic Game on Parallel Links, the Price of Anarchy has the following upperbound:*

$$PoA \leq 1 + \frac{1 - \phi_{min}}{C\phi_{min}}$$

3.2 Price of Anarchy in Non-Atomic Games

We have a more relaxed surplus condition for non-atomic games than its Atomic counterpart. We now need that each player can flow λ_i traffic from s_i to t_i given $\phi_i \mu_e$ capacity on each link. From the standard max flow-min cut theorem, we can express our surplus fraction in terms of s_i, t_i-cuts. We let E_W be the set of all edges crossing a cut $\emptyset \subset W \subset V$. For each $i \in K$, we are interested in the s_i, t_i-cut that will bottleneck our flow. This is defined as $\{s_i\} \subseteq W_i \subseteq V \setminus \{t_i\}$ based on the cost function from Lemma 1 as follows:

$$W_i = \operatorname*{argmax}_{s_i, t_i\text{-cut } W} \Big\{ \frac{|E_W|}{\phi_i \sum\limits_{e \in E_W} \mu_e - \lambda_i} \Big\}$$

Then our surplus factor for a given game is the following:

$$C = \min_{i \in K} \Big\{ 1 - \frac{\lambda_i}{\phi_i \sum_{e \in E_{W_i}} \mu_e} \Big\}$$

Theorem 6. *Assuming no player starves (i.e. $C > 0$), in a GPS Non-Atomic Game on a Graph, the Price of Anarchy has the following upperbound:*

$$PoA \leq (|V| - 1)|E| \Big(1 + \frac{1 - \phi_{min}}{C\phi_{min}}\Big)$$

Proof. First, we will give an upperbound for each type of traffic i on its total delay $\sum_{e \in E} \lambda_{i,e} d_{i,e}$. At equilibrium, all the traffic of type i must be using paths with minimum total delay (otherwise some traffic would change to a faster path). Then our upperbound on the total delay will follow from finding a s_i, t_i-path with bounded delay. We let d^p_i denote the sum of delays $d_{i,e}$ on the links of p for player i.

Claim 1. *Let S be any Nash Equilibrium in a GPS Non-Atomic Game on a Graph and $i \in K$ be any traffic of type. There exists an s_i, t_i-path p satisfying the following:*

$$d^p_i \leq \frac{(|V| - 1)|E|}{\phi_i \sum_{e \in E_W} \mu_e - \lambda_i}$$

Proof. Let $E(S)$ be the set of directed edges being used by traffic of type i in the assignment S. Each path being used must have minimum delay among the possible s_i, t_i-paths. Then a standard result from network flows gives us that the digraph $(V, E(S))$ must be acyclic (a cycle would imply one of the paths can be shortened). Then we use this directed acyclic graph to define a partial ordering $\succ_S$ over V.

Let W be any cut such that no vertex in $V \setminus W$ precedes a vertex of W in our $\succ_S$. We call such a set *orderly*. Since all edges of $E(S)$ are directed out of an orderly set W, we know there is exactly λ_i traffic of type i crossing the cut. Next we consider the average value of $d^{-1}_{i,e} = \tilde{\mu}_{i,e} - \lambda_{i,e}$ of an edge $e \in E_W$ to show an edge with bounded delay exists:

$$\operatorname*{average}_{e \in E_W} \big(d^{-1}_{i,e}\big) \geq \frac{\phi_i \sum\limits_{e \in E_W} \mu_e - \lambda_i}{|E_W|} \geq \frac{\phi_i \sum\limits_{e \in E_{W_i}} \mu_e - \lambda_i}{|E_{W_i}|}$$

$$\implies \exists e \in E_W, \text{ s.t. } d_{i,e} \leq \frac{|E_{W_i}|}{\phi_i \sum\limits_{e \in E_{W_i}} \mu_e - \lambda_i}$$

Now we will use this to show there exists a path from s_i to t_i with delay at most $|V| - 1$ times this delay bound. We construct this path by using Dijkstra's Shortest Path Algorithm. At each iteration j let R_j be the set of vertices in our shortest path tree. We maintain two invariants during this procedure.

First, we claim the set R_j is orderly. If this was not the case, then one of the paths in equilibrium could be shortened by using the shortest path tree.

Second, we claim that all vertices in R_j can be reached by a path of delay at most $j \times |E_{W_i}|/(\phi_i \sum_{e \in E_W} \mu_e - \lambda_i)$. We prove this by induction on j. It is trivially true for $R_0 = \{s_i\}$. Since R_j is orderly, there must be an edge $uv \in E_{R_j}$ with delay at most $|E_{W_i}|/(\phi_i \sum_{e \in E_W} \mu_e - \lambda_i)$. By induction, there is a path to v with delay at most $(j + 1) \times |E_{W_i}|/(\phi_i \sum_{e \in E_W} \mu_e - \lambda_i)$. Thus, all vertices in R_{j+1} must also have a path from s_i under this bound.

Since the shortest path algorithm will add t_i in at most $|V| - 1$ iterations, our claim follows from the second invariant. □

We use this path to give an upperbound on the total delay of traffic of type i at NE. We lower bound the total delay of an optimal assignment using the fact that it must cross the cut W_i. Combing these two bounds, we can manipulate the definition of β, as done in Theorem 5, to get our PoA bound.

$$\beta = \sup_{i \in K} \frac{\sum_{e \in E} \lambda_{i,e}(S) d_{i,e}(S)}{\sum_{e \in E} \lambda_{i,e}(S^*) d_{i,e}(S^*)} \tag{7}$$

$$\leq \sup_{i \in K} (|V| - 1)|E| \Big(1 + \frac{1 - \phi_i}{C\phi_i}\Big) \tag{8}$$

Then the Theorem follows from Lemma 2. □

Corollary 2. *Assuming no player starves (i.e. $C > 0$), in a GPS Non-Atomic Game on Parallel Links, the Price of Anarchy has the following upperbound:*

$$PoA \leq |E| \Big(1 + \frac{1 - \phi_{min}}{C\phi_{min}}\Big)$$

3.3 Quality of Price of Anarchy Bounds

Theorem 7. *For any $C > 0$ and $0 < \phi_{min} \leq 1/2$, there exists a GPS Atomic Game on Parallel Links instance with surplus factor C and minimum priority ϕ_{min}, with Price of Anarchy arbitrarily close to the upperbound in Corollary 1.*

Proof. Let $\epsilon > 0$ be a small number. We consider a GPS Atomic Game on Parallel Links with two players and two links. Our players are defined by $\lambda_1 = \epsilon$, $\lambda_2 = 1 - C$, $\phi_1 = 1 - \phi_{min}$, and $\phi_2 = \phi_{min}$. Our links have capacities $\mu_1 = 1$ and $\mu_2 = 1/\phi_{min}$. Note that the surplus fraction of this game is in fact C. Now we give two assignments that will give us a lowerbound on the PoA in this game. First, consider the assignment of player 1 to link 1 and player 2 to link 2. The total delay of this assignment must be larger or equal to the total delay of the optimal assignment. Then we can conclude the following:

$$\lim_{\epsilon \to 0} |OPT| \leq \frac{1 - C}{1/\phi_{min} - 1 + C}$$

Second consider the assignment of player 1 to link 2 and player 2 to link 1. Note that this assignment is an equilibrium. Thus the total delay of this assignment is less than or equal to the total delay of the worst NE. Combining this with our bound on $|OPT|$, we get the following lower bound on PoA:

$$\lim_{\epsilon \to 0} \frac{|NE|}{|OPT|} \geq \frac{1/\phi_{min} - 1 + C}{1 - C} \Big(\frac{1 - C}{1 - 1 + C}\Big)$$
$$= \frac{1 - \phi_{min}}{C\phi_{min}} + 1$$

This value is exactly tight with our bound for games on a Parallel Link topology. □

Theorem 8. *For any $|V| \geq 4$ and $0 < \phi_{min} \leq 1/2$, there exists a GPS Atomic Game on a Graph instance with a $|V|$ node graph, minimum priority ϕ_{min} and surplus factor $C = \phi_{min}$, with Price of Anarchy within a small factor of the upperbound in Theorem 5.*

Proof. Let $\epsilon > 0$ be a small number. Let G be a cycle on $|V|$ vertices labeled in order $v_1, v_2, ..., v_n$. We assign the link $v_1 v_n$ capacity $\mu_{v_1 v_n} = (1 - \phi_{min})/\phi_{min}$ and all other links capacity $\mu_e = 1/\phi_{min}$. Consider the game on this network with two players defined as follows. Player 1 routes $\lambda_1 = 1 - C$ traffic from v_1 to v_n with priority $\phi_1 = \phi_{min}$. Player 2 routes $\lambda_2 = \epsilon$ traffic from v_1 to $v_{\lceil n/2 \rceil + 1}$ with priority $\phi_2 = 1 - \phi_{min}$. Note that the surplus fraction of this game is in fact $C = \phi_{min}$.

Now we give two assignments that will give us a lowerbound on the PoA in this game. First, consider the assignment of player 1 to the path $v_1 v_n$ and player 2 to the path $v_1, v_2, ...v_{\lceil n/2 \rceil + 1}$. The total delay of this assignment must be larger or equal to the total delay of the optimal assignment. Then we can conclude the following as ϵ approaches zero (note we can disregard the objective of player 2 since its multiplied by ϵ):

$$\lim_{\epsilon \to 0} |OPT| \leq \frac{1 - C}{(1 - \phi_{min})/\phi_{min} - 1 + C}$$

Second consider the assignment of player 1 to the path $v_1, v_2, ...v_n$ and player 2 to the path $v_1, v_n, v_{n-1}, ...v_{\lceil n/2 \rceil + 1}$. Note that the capacity guaranteed to player 2 on every link is $(1 - \phi_{min})/\phi_{min}$. Thus player 2 is stable. Further since player 1 has demand equal to its guaranteed capacity on $v_1 v_n$, it will not change its move. Then this is an equilibrium and has delay at most equal to the worst NE. Combining this with our bound on $|OPT|$, we get the following lower bound on PoA:

$$\lim_{\epsilon \to 0} \frac{|NE|}{|OPT|} \geq \frac{(1 - \phi_{min})/\phi_{min} - 1 + C}{1 - C} \frac{(|V|/2)(1 - C)}{1 - 1 + C}$$
$$\geq \frac{|V|}{2} \Big(\frac{1 - 2\phi_{min}}{C\phi_{min}} + 1\Big)$$

□

Now we compare our Non-Atomic bound to the bounds for systems without priorities. The PoA when all traffic is of a single type is tightly bounded by $|E|$ [17]. Our bounds on the price of anarchy show that adding priorities and capacity guarantees to the system increases the PoA by at most a factor of $(|V| - 1)(1 + \frac{1 - \phi_{min}}{C\phi_{min}})$. This is a reasonable increase since we have already shown that GPS scheduling has this behavior in the Atomic case.

4 Computational Complexities

In this Section, we consider the complexity of computing various assignments for an instance of one of our games. Specifically, we focus on the assignment with minimum total delay and NE assignments. We formulate these into two decision problems, SOCIAL-OPTIMUM and NE-EXISTENCE. First we define the decision problem, SOCIAL-OPTIMUM: The input is a game instance and a positive number R. It answers yes if and only if the optimal assignment in the game has total delay at most R. Secondly, we define the decision problem, NE-EXISTENCE: The input is a game instance. It answers yes if and only if the game has a Nash Equilibrium assignment.

Using a reduction from the PARTITION problem, we can show SOCIAL-OPTIMUM is NP-Complete. We omit this proof due to length constraints, but it follows from a simplified version of our proof for Theorem 11.

Theorem 9. SOCIAL-OPTIMUM *for a GPS Atomic Game on Parallel Links (and thus on a Graph topology) is NP-Complete.*

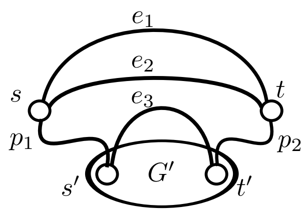

Fig. 2. The graph used to show computing the existence of Nash Equilibrium in atomic games is NP-Complete.

4.1 Nash Equilibrium in Atomic Games

Next we consider the problem of computing a NE in a given game, or determining that none exist. We have already proven that an equilibrium will exist in any GPS Atomic Game on Parallel Links. In Algorithm 1, we give a fast algorithm for computing such an assignment.

Algorithm 1 NE in GPS Atomic Game on Parallel Links

```
1: Order φ_1 ≥ φ_2 ≥ ... ≥ φ_k
2: for p = 1...k do
3:     S_p := Best response for player p given S_1, ..., S_{p-1}
4: end for
```

Theorem 10. *For any GPS Atomic Game on Parallel Links instance with k players and m links, Algorithm 1 outputs a NE in time $O(k\dot{l}og(k) + km)$.*

Proof. We prove this by showing that at the end of iteration i, players 1 through i are stable. Then after all k iterations, we will have a NE.

Our invariant trivially holds for $i = 1$, since player 1 will make his best move (ie choosing the highest capacity link). Similarly, player i will be stable at the end of iteration i. Then we only need to show that after player i's assignment, each player j for $j < i$ is still stable. Any player $j < i$ made unstable by player i must be using the same link e as player i. Let e' be any other link. Since player i chose e, we know the following:

$$\frac{\phi_i}{C_{S(e)} + \phi_i}\mu_e > \frac{\phi_i}{C_{S(e')} + \phi_i}\mu_{e'} \tag{9}$$

$$\implies \frac{\phi_j}{C_{S(e)} + \phi_i}\mu_e > \frac{\phi_j}{C_{S(e')} + \phi_j}\mu_{e'} \tag{10}$$

The second inequality follows since $\phi_j > \phi_i$. Then we can conclude that player j does not get a higher capacity on any other link after player i moves in. Therefore each player $j < i$ is still stable. The Theorem follows. □

Now, we focus on the more difficult problem of whether or not an equilibrium exists in the general graph topology. In our next Theorem, we show that this problem is NP-Complete.

Theorem 11. NE-EXISTENCE *for a GPS Atomic Game in a Graph topology is NP-Complete.*

Our proof is based on a reduction from the PARTITION problem. An instance of PARTITION has input of a multiset of positive integers $X = \{x_1, x_2, ...x_n\}$. It accepts if it can find a partition of X into A and B such that

$$\sum_{x_i \in A} x_i = \sum_{x_i \in B} x_i = \frac{1}{2}\sum_{x_i \in X} x_i$$

Let $M = \sum x_i$. Now we define two constants, ϵ and p, to guarantee certain properties, which we will use later. First, let ϵ satisfy the following:

$$0 < \epsilon < min\Big\{\frac{2}{M/2+1}, \frac{2\lambda_{2'}}{M + \lambda_{2'}}\Big\}$$

This choice guarantees the following two properties, which will be used later in our proof:

$$\frac{M}{2} + 1 > \frac{M}{2-\epsilon} \tag{11}$$

$$\frac{\epsilon}{1+\epsilon}\frac{M}{2-\epsilon} < \lambda_{2'} \tag{12}$$

Now let p be an integer satisfy the following inequality to ensure the subsequent property:

$$p > \frac{M}{\epsilon} - \frac{2-\epsilon}{\epsilon} \implies \forall x_i, \frac{p}{\frac{M}{2-\epsilon} - 1} > \frac{1}{x_i\frac{2}{2-\epsilon} - x_i} \tag{13}$$

Our reduction utilizes the example game from Theorem 2 that has no NE. Let G' be a copy of this game's graph with all demands and capacities reduced by a factor of 1000. Note that all capacities are now less than 1. We will refer to the two players from G' as subgame players. Let $\{s', u', v', t'\}$ be the labels for the nodes of G', and let $\lambda_{1'}$ and $\lambda_{2'}$ denote the demands of the subgame players.

To construct our instance of NE-EXISTENCE, we add two nodes, s and t, to G'. We connect s and t with two parallel links of capacity $\mu_1 = \mu_2 = M/(2-\epsilon)$. Now we add two paths of length p connecting s to s' and connecting t to t' where all edges have capacity $M/(2-\epsilon)$. Finally, we add a third link from s' to t' with capacity $M/(2-\epsilon)$. The resulting graph is shown in Figure 2. The two subgame players are both given priority $\phi_{1'} = \phi_{2'} = \epsilon$. Our game will have n other players sending their traffic from s to t. We will refer to these as regular players, and they each have priority $\phi_i = x_i$ and demand $\lambda_i = x_i$. Note that our priorities are not normalized for ease of notation.

Claim 2. *If a partition into A and B of equal size exists, then our GPS Atomic Game instance has a Nash Equilibrium.*

Proof. We prove this directly by giving a NE assignment. For each $x_i \in A$, assign player i to link 1 and for each $x_j \in B$, assign player j to link 2. Then assign both of the players from G' to link 3. Equation (13) guarantees no regular player will prefer a path using p_1 or p_2 over links 1 and 2 with total priority of $M/2$ on each. Thus our n player corresponding to x_i are stable. Players 1′ and 2′ will not want to change to a path in G' since all its capacities are less than 1. Equation (12) ensures the two subgame players cannot share a link with our regular players (because $\frac{\epsilon}{1-\epsilon}\frac{M}{2-\epsilon} < \lambda_{2'}$). Thus they will not want to change to a path through s and t. Then we can conclude all of our players are making best response moves. □

Claim 3. *If our GPS Atomic Game instance has a Nash Equilibrium, then a partition into A and B of equal size exists.*

Proof. First, we will prove by contradiction that all n regular players must be using links 1 and 2 at equilibrium.

Assume a player i is not using links 1 or 2. Then player i must be using the path p_1 then link 3 then p_2, since all links in G' have capacity less than 1. Equation (11) states that our subgame players cannot share links with our n regular players. Since there is a regular player on link 3, our subgame players can only use the links in G'. However, we know that no NE exists for those two players on those two links. This contradicts the fact that the assignment is a NE.

Since all n regular players are using links 1 and 2 in our assignment, we can construct the corresponding partition into A and B. Each player must have a finite delay, since they could switch to using p_1, then link 3, then p_2, which would give them finite total delay. Thus, the total priority on links 1 and 2 do not exceed $M/2$. This directly implies our partition is valid. □

4.2 Complexities in Non-Atomic Games

Now that we have fully categorized the complexity of our problems in atomic games, we focus on non-atomic games. Using the same method as Theorem 11, we can prove that computing socially optimal assignments is also NP-Compete for non-atomic games. We omit the proof of this Theorem since it is very similar to our previous reduction.

Theorem 12. SOCIAL-OPTIMUM *for a GPS Non-Atomic Game on Parallel Links (and thus on a Graph topology) is NP-Complete.*

We have already show that determining the existence of NE is NP-Complete for Atomic Graph games. A slight modification of our previous proof extends this result to GPS Non-Atomic Games on Parallel Links. The construction will embed the example game from Theorem 3 into the network.

Theorem 13. *The problem of* NE-EXISTENCE *for a GPS Non-Atomic Game in a Parallel Links topology (and thus a Graph topology) is NP-Complete.*

5 SIMULATIONS OF AVERAGE CASE BEHAVIOR

In this Section, we investigate the average behavior of these games on parallel links experimentally. All of our simulations are implemented in JAVA. First we compare the actual Price of Anarchy to our bound in a simple Atomic Parallel Links game. Our sample game will have three links and six players. The links have capacities $\mu_1 = 10$, $\mu_2 = 10$, and $\mu_3 = 40$. We will vary the ϕ_1 and λ_1 while keeping all other players i fixed $\phi_i = (1 - \phi_1)/5$ and $\lambda_i = 4$. In Figure 3, we show our PoA bound and the actual PoA in this network for different values of ϕ_1 and λ_1. Repeating this experiment with Non-Atomic Routing yields very similar results.

Now we consider the price of anarchy in a randomly generated Parallel Links network to understand its average cost. For our experiment, we randomly generate a game with 5 players and 5 links as follows: Each player's demand and priority are selected uniformly at random from $[0, 1]$. The capacity of each link is selected uniformly at random from $[0, 3]$. In Figure 4, we show the results of running this experiment for a non-atomic game. For each pair of surplus

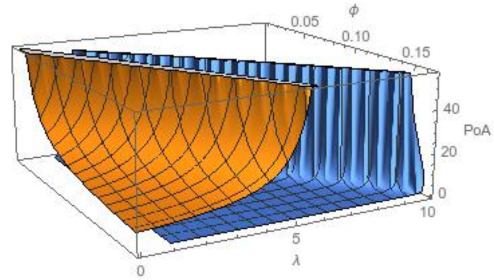

Fig. 3. Shows the Price of Anarchy bound from Corollary 1 and the actual Price of Anarchy in our sample GPS Atomic Game on Parallel Links as functions of ϕ_1 and λ_1.

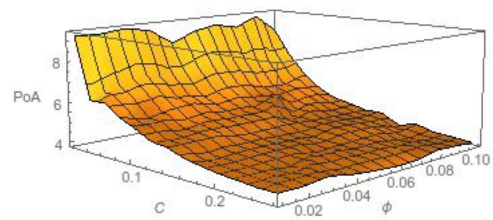

Fig. 4. Average Price of Anarchy for a GPS Non-Atomic Game on a random Parallel Links network with $C \in [0.01, 0.4]$ and $\phi_{min} \in [0.01, 0.1]$.

factor C and minimum priority ϕ values, we average the price of anarchy of 500 random games with the corresponding surplus and minimum priority. All of our average PoAs are at most 10. The result of this procedure for atomic games is a nearly identical plot with maximum values around 1.6.

Next, we examine the average number of best response moves until the network reaches a Nash Equilibrium. In a GPS Atomic Game on Parallel Links, we are guaranteed that this exists. For a non-atomic game, we have no such guarantee, but we know that parallel link networks with no NE are very uncommon. To measure the average time to converge to an equilibrium, we first generate a random graph with n players and n parallel links as follows: Each player's demand and priority are selected uniformly at random from $[0, 1]$. The capacity of each link is selected uniformly at random from $[0, \mu]$ for some fixed μ. Then we compute a random initial assignment of traffic. Then we iteratively let a random type of traffic change to its best response assignment until we reach a NE. In Figures 5 and 6, we show the average number of iterations in this procedure as n and μ vary for atomic and non-atomic games respectively. Each data point comes from averaging the number of steps to converge in 2000 different random games. The results in Figure 5 show that GPS Atomic Games on Parralle Links typically converge to an equilibrium after a linear number of updates. However, this quick convergence is not the case for non-atomic games. We observe an exponential rate of growth in the number of iterations in Figure 6 which might be expected since we have proven that computing

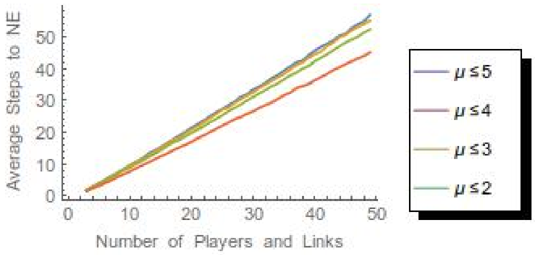

Fig. 5. The average number of best response moves before reaching a Nash Equilibrium in a random atomic game.

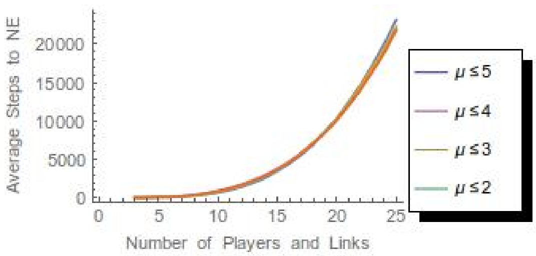

Fig. 6. The average number of best response moves before reaching a Nash Equilibrium in a random non-atomic game.

equilibrium in GPS Non-Atomic Games is NP-Complete.

6 Conclusion and Acknowledgment

We studied the behavior of routing systems that provide priority-based capacity guarantees to users. Our focus was on networks where each link behaves as a M/M/1-GPS queue. We considered four different games under this model and for each characterized the existence of Nash Equilibrium, the Price of Anarchy, its computational complexity, and gave simulation results. These results provide interesting insights into the behaviors of this type of system and the inefficiencies that would come with implementing it. Further, these results may have implications on the debate over Net Neutrality.

Acknowledgments

The research was supported in part by NSF grant: CCF-1451574.

References

[1] Abhay K. Parekh and Robert G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Trans. Netw.*, 1(3):344–357, June 1993.

[2] Y.A. Korilis, A.A. Lazar, and A. Orda. Architecting noncooperative networks. In *Electrical and Electronics Engineers in Israel, 1995., Eighteenth Convention of*, pages 5.3.1/1–5.3.1/5, March 1995.

[3] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[4] Tim Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.

[5] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 604–612, New York, NY, USA, 2004. ACM.

[6] Roberto Cominetti, Jos R. Correa, and Nicols E. Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.

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

[8] E. Monsef, T. Anjali, and S. Kapoor. Price of anarchy in network routing with class based capacity guarantees. In *INFOCOM, 2014 Proceedings IEEE*, pages 2409–2417, April 2014.

[9] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *Commun. ACM*, 52(2):89–97, February 2009.

[10] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, STACS'99, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[11] Andreas S. Schulz and Nicolás Stier Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, pages 86–87, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

[12] Georgia Perakis and Guillaume Roels. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Manage. Sci.*, 53(8):1249–1268, August 2007.

[13] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier-Moses. Selfish routing in capacitated networks. *Math. Oper. Res.*, 29(4):961–976, November 2004.

[14] Sanjiv Kapoor and Junghwan Shin. Price of anarchy with heterogeneous latency functions. *CoRR*, abs/1407.2991, 2014.

[15] Colin E. Bell and Jr. Stidham, Shaler. Individual versus social optimization in the allocation of customers to alternative servers. *Management Science*, 29(7):pp. 831–839, 1983.

[16] E. J. Friedman. A generic analysis of selfish rouing. Technical report, Cornell University Operations Research and Industrial Engineering, 2002.

[17] Moshe Haviv and Tim Roughgarden. The price of anarchy in an exponential multi-server. *Oper. Res. Lett.*, 35(4):421–426, July 2007.

[18] Yannis A. Korilis and Aurel A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *J. ACM*, 42(3):584–613, May 1995.

[19] E. Altman, U. Ayesta, and B.J. Prabhu. Load balancing in processor sharing systems. *Telecommunication Systems*, 47(1-2):35–48, 2011.

[20] T.L. Werth, M. Holzhauser, and S.O. Krumke. Atomic routing in a deterministic queuing model. *Operations Research Perspectives*, 1(1):18 – 41, 2014.

[21] Zhi-Li Zhang, D. Towsley, and J. Kurose. Statistical analysis of the generalized processor sharing scheduling discipline. *Selected Areas in Communications, IEEE Journal on*, 13(6):1071–1080, Aug 1995.