

A Priority-Based Model of Routing

Babak Farzad, Neil Olver and Adrian Vetta*

May 13, 2009

Abstract

We consider a priority-based selfish routing model, where agents may have different priorities on a link. An agent with a higher priority on a link can traverse it with a smaller delay or cost than an agent with lower priority. This general framework can be used to model a number of different problems. The structural properties that lead to inefficiencies in routing choices appear different in this priority-based model compared to the classical model. In particular, in parallel link networks with nonatomic agents, the price of anarchy is exactly one in the priority-based model; that is, selfish behaviour leads to optimal routings. In contrast, in the standard model the worst possible price of anarchy can be achieved in a simple two-link network. For multi-commodity networks, selfish routing does lead to inefficiencies in the priority-based model. We present tight bounds on the price of anarchy for such networks. Specifically, in the nonatomic case the worst-case price of anarchy is exactly $(d + 1)^{d+1}$ for polynomial latency functions of degree d (hence 4 for linear cost functions). For atomic games, the worst-case price of anarchy is exactly $3 + 2\sqrt{2}$ in the weighted case, and exactly $17/3$ in the unweighted case. An upper bound of $O(2^d d^d)$ is also shown for polynomial cost functions in the atomic case, although this is not tight. Our framework (and results) also generalise to include models similar to congestion games.

ACM Classification: F.2.0, F.2.2

AMS Classification: 68Q25, 68M10, 90B18

Key words and phrases: selfish routing, price of anarchy

1 Introduction

This work is motivated by the simple observation that, in a transportation network, a car traversing a road can only cause congestion delays to those cars that use the road at a *later* time. Moreover, this is a common feature of most traffic networks and queuing models.

The study of congestion and transportation networks is not new. The ideas were first discussed qualitatively by Pigou [14] in 1920, and later placed on a sound mathematical footing by Wardrop [20]. The book of Beckmann, McGuire and Winsten [4] gives a very thorough treatment. More recently, applications to communication networks such as the

*McGill University. Email: {babak,olver,vetta}@math.mcgill.ca




Figure 1: Classical versus priority-based selfish routing

internet spurred interest from the computer science community. The concept of the *price of anarchy* was introduced by Koutsoupias and Papadimitriou [12]. It is the ratio between the costs of the worst Nash equilibrium and the optimal routing, and is essentially a quantitative measure of the loss of efficiency attributable to the lack of a central coordinating authority.

In this classical selfish routing model, each link e has an associated *cost function* $f_e(x)$; the delay experienced by the users of this link is then $f_e(x_e)$, where x_e is the total traffic on the link. Thus all users of a link experience the same latency. One practical situation in which this may arise is when users are continuously using a network, making the concept of time redundant. There are many situations where this assumption is not valid however. For example, imagine someone driving home during rush hour in a large city. The time that they leave will make a big difference to how long the trip takes; it will be much shorter if they leave early enough to avoid the worst of the traffic.

Here we show that a simple modification to the classical model does allow us to incorporate some elements of time dependence. In the classical model, the total cost associated with a link e is the delay experienced on that link, multiplied by the number of players using it, i.e. $f_e(x_e)x_e$. In our model, the total cost will instead be given by the area under the cost function, i.e. the integral $\int_0^{x_e} f_e(z)dz$. The idea is that the area under the cost function can be partitioned amongst the users so that earlier users are associated with smaller latencies. If a player j has an amount $x_e^{(j)}$ of flow ahead of it, it will experience a delay of $f_e(x_e^{(j)})$. The difference between the models is represented visually in Figure 1; the total cost associated with link e in the classical model is given by the area of the lightly shaded rectangle, and in the priority-based model it is given by the area under the curve.

There are other reasons aside from time-based considerations why different users might experience different delays or costs; for instance, certain users might simply be given priority, and always experience lower latencies. Our model allows the ordering of the players to be defined very generally; various examples will be discussed later.

Both the classical and priority-based models can be broadly divided into two variants; atomic and nonatomic. In the atomic case, there are a finite number of agents, each with a certain amount of flow to route. The flow may be *splittable*, or *unsplittable*, in which case

each agent must pick a single path for its entire flow. In the nonatomic case, there are an infinite number of agents, and each controls only a negligible fraction of the total flow. The results and techniques will be different for these two variants; the atomic case is generally more difficult to analyse.

This simple modification of using an integral rather than a rectangle to measure the cost on a link has more of an effect than one might expect. Roughgarden and Tardos [18, 16] showed that in the classical model with nonatomic agents, the worst-case price of anarchy is $4/3$ for linear cost functions, and $(1 - d(d+1)^{-1-1/d})^{-1} = \Theta(d \log d)$ for polynomial cost functions of degree d . By contrast, we show that in the nonatomic priority-based model, the worst-case price of anarchy is exactly 4 for linear cost functions, and $(d+1)^{d+1}$ for polynomial cost functions of degree d ; these are considerably larger.

In addition, some of the causes of the inefficiencies due to selfish routing appear to be different. In particular, it is known [16] that even in single-commodity networks (in fact, even in simple parallel link networks), the above worst-case bounds in the standard nonatomic model can be achieved. Thus the worst-case price of anarchy is essentially independent of the network topology in the standard model. By contrast, we will show that selfish routing leads to optimal solutions in parallel link networks in the priority-based model, for any choice of priority scheme. For some important special cases of our model (including the time-based model mentioned earlier), this is still true for arbitrary single-commodity networks, where all agents have the same origin and destination.

The atomic unsplittable case of the classical model was considered by Azar, Awerbuch and Epstein [3]. They show that for linear cost functions, $(3 + \sqrt{5})/2$ is a tight upper bound for the price of anarchy; this is reduced to 2.5 in the unweighted case, where all users route one unit of demand (see Christodoulou and Koutsoupias [5] for an independent proof). We show that in the priority-based model with linear cost functions, the worst-case price of anarchy is $3 + 2\sqrt{2}$, and reduces to $17/3$ in the unweighted case. They also show that for polynomial cost functions of degree d , the worst possible price of anarchy is $d^{\Theta(d)}$; this was later determined exactly by Aland et al. [2] (see also Olver [13]). We show an upper bound of $O(2^d d^d)$ in our model for this case.

Related work Rosenthal [15] introduced atomic selfish routing games, as well as *congestion games*, an important generalisation which removes the network structure. Our model generalises in an analogous way.

Correa, Schulz and Stier-Moses [6] gave shorter proofs of some of the price of anarchy results for nonatomic games, as well as some new results. Some of our proofs are inspired by their technique.

Independently of this work, Harks, Heinz and Pfetsch [9] consider an online version of the multicommodity routing problem. The greedy online algorithm they consider can be interpreted as the Nash equilibrium in an instance of what we call the *global priority model*, a special case of the priority-based model. Thus price of anarchy results in the global priority model are related to online competitiveness results in their model. Harks and Végé [10] generalise [9] to an online selfish routing game; in this game, a sequence of standard selfish routing games are played, and players in a particular game are aware of the choices made by the players in earlier games, but not later ones. This model are in some sense a generalisation of the global priority model where some players have identical

priorities.

Paper outline In Section 2 we present the priority-based selfish routing model, considering atomic and nonatomic agents. We also give a number of motivating applications, and show some sufficient conditions for the existence of Nash equilibria. The bulk of the paper is devoted to deriving the exact value for the price of anarchy under certain restrictions on the cost functions. In Section 3, we consider the price of anarchy of nonatomic agents; single-commodity networks are considered first, followed by the general multicommodity case. Atomic agents are dealt with in Section 4.

2 The Model

We now define the model rigorously. We begin with the atomic unsplittable case, since this is actually easier to define (although more difficult to analyse).

2.1 The unsplittable atomic case

We begin with a network, represented as a directed graph $G = (V, E)$, and a finite number n of players. Each player j has a flow requirement of w_j units, which must be routed from node s_j to node t_j . The players must each pick a single path to route their entire demand. A particular routing is then defined by $\mathcal{P} = \{P_1, \dots, P_n\}$, where P_j is an $s_j - t_j$ path for each j , representing the route taken by that player. We also define the flow vector $\mathbf{x}(\mathcal{P})$ by $x_e(\mathcal{P}) = \sum_{j:e \in P_j} w_j$, the total flow on edge e . We will write simply x_e if the desired routing is clear. Each edge has an associated cost function f_e that is nonnegative and increasing. We will also sometimes refer to these as *latency functions*, since they represent the delay experienced by the user on the edge. So far, nothing we have described differs from the standard network routing model. But now we introduce a priority scheme that will allow us to order the users of a particular edge, prescribing different latencies to the users based on this order. We will allow this to be very general—the priority ordering on an edge can depend arbitrarily on the current routing \mathcal{P} . This can include dependence on routings that do not use that edge. If player i has higher priority than player j on edge e under routing \mathcal{P} , we write $i \succ_{\mathcal{P},e} j$. For a fixed e and \mathcal{P} , the relation $\succ_{\mathcal{P},e}$ must define a total ordering of the players using edge e ; this is the only restriction we impose. If it is clear from the context what edge or routing is being referred to, we will omit it.

For an arbitrarily defined priority scheme, it might not be computationally feasible to calculate a player’s best response, since there are an exponential number of paths to consider and the priority orderings could be different for all of them. In that case, best response dynamics and Nash equilibria would not be of much practical interest. Many natural priority schemes that we consider do allow best responses to be easily calculated. For example, one possibility would be to give an ordering to the players, and assign the priority along all the edges based on this ordering. Another option would be to assign priorities based on the time that the players arrive at the beginning of the link. We will describe in detail a number of possible priority schemes, including these ones, in Section 2.2.

In the classical selfish routing model, the total (or social) cost of a routing is given by $C(\mathcal{P}) = \sum_{e \in E} f_e(x_e)x_e$. As mentioned in the introduction, we will modify this in our

model, and define the total cost as

$$C(\mathcal{P}) = \sum_{e \in E} \int_0^{x_e} f_e(z) dz. \quad (1)$$

Let $x_e^{(j)}(\mathcal{P})$ be the amount of flow on edge e with a higher priority than player j under routing \mathcal{P} , i.e.

$$x_e^{(j)}(\mathcal{P}) = \sum_{i: i \succ_{\mathcal{P}, e} j} w_i.$$

Then we define $C_j(\mathcal{P})$, the cost attributable to player j , as

$$C_j(\mathcal{P}) = \sum_{e \in P_j} \int_{x_e^{(j)}(\mathcal{P})}^{x_e^{(j)}(\mathcal{P}) + w_j} f_e(x) dx. \quad (2)$$

For \mathcal{P} to be a Nash equilibrium, we must have for any player j and any $s_j - t_j$ path P' ,

$$C_j(\mathcal{P}) \leq C_j(\mathcal{P}') \quad (3)$$

where $\mathcal{P}' = \mathcal{P} \setminus P_j \cup P'$. This is simply a restatement of the condition that player j cannot switch to a cheaper route.

There is a slight subtlety to the interpretation of our modified cost. The definition is chosen so that the total cost on an edge is given by the integral $\int_0^{x_e} f_e(x) dx$, and so the total cost of a routing \mathcal{P} is given by (1). The analogous definition for the classical model was that player j contributed $f_e(x)w_j$ to the total cost; note that this is the delay experienced by the player multiplied by his weight. Our definition should be interpreted similarly; so the delay experienced by player j is $C_j(\mathcal{P})/w_j$. Depending on the application, this might not always be the “correct” choice; for instance, another natural option would be a delay of $f_e(x_e^{(r)})$. The difference will often not be significant, and our choice is more amenable to analysis. All of these difficulties disappear in the nonatomic version of the model, discussed in Section 2.3

Analogously to the classical case [15], we can also consider the congestion game generalisation of this priority-based model. Let I be a set of items; these will take the place of the edges in the network model. A cost function and priority ordering is associated with each item; again, the priority ordering can depend on the strategies chosen by the players. But now, each player has a set of possible strategies \mathcal{S}_j , where each strategy is some subset of the items. There is no restriction on what subsets can be specified as a player’s allowed strategies, or how many strategies a player may have. Notice that a network game is a special case of a congestion game where the strategies of player j are exactly the subsets corresponding to $s_j - t_j$ paths.

We will call a routing \mathcal{P} *optimal* if it has the minimum cost $C(\mathcal{P})$ over all feasible routings. We will often use the notation \mathcal{P}^* to refer to an optimal solution. Note that the optimality of a routing does not depend in any way on the priority scheme used.

2.2 Some possible priority schemes

Here we list some natural and interesting games that fall within the framework of our model, by picking the priority scheme appropriately.

The global priority game This is the simplest possible case; the ordering is independent of the routing, and is also the same for all edges. In other words, there is a fixed priority ordering of the players.

This model has an interesting interpretation as an online routing problem, as discussed by Harks et al. [9]. Suppose players are charged for their use of a network, and would like to choose the cheapest route for their demand. The players arrive one at a time however, and the amount charged to a player for using an edge depends on the *current* congestion on the edge; later players may be charged more for some edges than earlier players. The social optimum is taken to be the total cost incurred by all the players. In this setting, a Nash equilibrium can be interpreted as a greedy online algorithm, and the price of anarchy corresponds to the competitiveness of this algorithm. Most of the results in Harks et al. correspond to the atomic model, but with splittable flow; for instance, they show that with n unweighted splittable players, and linear cost functions, the price of anarchy cannot exceed $4n/(2+n)$.

The fixed priority game A more general model than the global priority one, here we still insist that the priorities are independent of the routing, but we allow different orderings on different edges. One application within this framework is Quality of Service in telecommunication networks such as the internet. In the absence of network neutrality, telecommunication companies could charge for faster access to the portion of the internet that they own. The players in this case would be companies with a large internet presence; these companies want to serve content to their users as quickly as possible. The priorities would then be determined by contracts between the companies and the service providers.

The timestamp game The priorities of agents are determined by their arrival times at the start of the edge. Associate with each agent j an additional value τ_j that represents the starting time of that agent. Now take a specific routing $\mathcal{P} = \{P_1, \dots, P_n\}$. The time agent j arrives at a vertex $u \in P_j$ is then τ_j plus the time taken to traverse all the edges on the subpath of P_j from s_j to u , denoted $P_j[s_j, u]$.

Of course, the latency of player j along an edge in $P_j[s_j, u]$ depends on the priority of j on that edge, which in turn depends on the start times of other agents. To see that we have enough information to uniquely determine the priorities, imagine simulating the game. Take the player with the smallest starting time, and move her along the first edge of her path. Her timestamp is then adjusted to include the time taken to traverse this edge. We then repeat, taking the player with the smallest timestamp after this update (this could be the same player). When the simulation terminates and all players have reached their destinations, we can read off the priority ordering on any edge; it is simply the order in which the players traversed that edge in the simulation.

The technical issue of ties—two agents taking the same edge with the same timestamp—can easily be resolved, either by prescribing a tie-breaking order for the players, or by perturbing the starting times by sufficiently small values to break the ties without modifying the ordering.

One problem with this model is that a player using a link delays *all* players that traverse the link afterwards. This is not very realistic—rather, the delaying effect of a player should only last for a short time. Congestion effects, however, complicate the situation; the duration of the delaying effect can vary quite dramatically depending on the situation. If a link is heavily congested, a player might be delayed by traffic from quite some time ago. A better model could be obtained using a proper model of flow over time, where packets have a well defined position at each moment in time, and so the amount of traffic on an edge can vary. Such a “dynamic” model was introduced by Ford and Fulkerson [7], and has received considerable attention. Köhler and Skutella [11] have proposed a dynamic model with load-dependent transit times which can incorporate congestion effects. Investigating the behaviour of such a dynamic model with selfish behaviour would be an interesting avenue of research.

For the congestion game variant of our model, the following is one motivating example:

The subcontractor game Suppose there are a number of construction companies, each involved with one or more large project. In order to complete various parts of these projects, the construction companies need to enlist the services of subcontractors. There are many subcontractors to choose from, and different subcontractors provide different subsets of services (more than one subcontractor may offer the same service). Also, there might be more than one way of completing a project, and so the construction company might have a choice of which services are required. However, if two companies decide to use the same subcontractor, and require some of the same services, the subcontractor will not be able to complete both requests simultaneously. A delay in construction will negatively affect the profit of the construction company, so essentially the company that gets delayed is paying more for the service.¹ The subcontractor’s choice as to which company to delay could depend on many factors—the time that the contracts were made, the total value of the contracts, previous business relationships with the companies, etc.

To model this as a priority-based congestion game, we will consider each company as an agent, and each service offered by a subcontractor as an item (the same service offered by multiple contractors will be considered as multiple items). The possible strategies for an agent will then be any combination of services from the subcontractors that together provide all the needs of the agent. Since our model is so general, the priority ordering could be as complicated as needed to take the various factors noted above into account.

2.3 The nonatomic case

If we let the fraction of the total flow controlled by any single player diminish to zero, the game becomes nonatomic. We have to be quite careful in defining things formally however—there are some subtleties that do not occur in the standard model. Our approach to nonatomic games follows Schmeidler [19], in that the game is represented by an atomless space of players, and each player has an associated payoff function.

Label the players by elements in the interval $R = [0, 1]$. We also have two measurable

¹Alternatively, the subcontractor may experience increasing marginal costs; for example, these may be due to overtime payments, increased costs arising from the need for additional production, etc. These additional costs are then passed on to the construction companies.

functions $s, t : R \rightarrow V$, which specify the origin and destination of each player respectively.

For each $r \in [0, 1]$, the strategy of player r is given by a unit $s_r - t_r$ flow \mathbf{y}_r . We are allowing splittable flow—the reason for this is discussed later in this section. A feasible solution is given by $\mathcal{P} = \{\mathbf{y}_r : r \in R\}$. A priority ordering is defined as before; for any edge e and feasible solution \mathcal{P} , $\succ_{\mathcal{P}, e}$ is a total ordering of the players.

The total flow on edge e will then be $x_e = \int_R y_r(e) d\mu$. The amount of flow on edge e ahead of player r is given by

$$x_e^{(r)}(\mathcal{P}) = \int_{L_e^{(r)}(\mathcal{P})} y_s(e) d\mu,$$

where $L_e^{(r)}(\mathcal{P}) = \{s : s \succ_{\mathcal{P}, e} r\}$. We require that $L_e^{(r)}(\mathcal{P})$ be Lebesgue measurable for all $e \in E$, $r \in R$ and feasible \mathcal{P} ; any reasonable ordering will satisfy this technical requirement.

The total latency experienced by player r , i.e. the time taken for the player to traverse from the source to the sink, is

$$\ell_r(\mathcal{P}) = \sum_{e \in P_r} f_e(x_e^{(r)}(\mathcal{P})). \quad (4)$$

Analogously to the atomic case, the requirement for \mathcal{P} to be a Nash equilibrium is that for any $r \in R$ and any $s_r - t_r$ path P' ,

$$\ell_r(\mathcal{P}) \leq \ell_r(\mathcal{P}') \quad (5)$$

where $\mathcal{P}' = \mathcal{P} \setminus P_j \cup P'$.

The model is most easily thought of as a nonatomic version with *splittable* flow. There is a reason for this; in general, we cannot assign an unsplittable flow to each nonatomic player. For consider a simple two-link network with arcs e and e' , and define the cost functions $f_e(x) = f_{e'}(x) = x$. Assign the global priority ordering $r \succ s$ iff $r < s$. Now suppose we demand that each player routes an unsplittable flow; consider any such solution \mathcal{P} where the flows \mathbf{y}_r are all 0 – 1 vectors. Any Nash equilibrium must satisfy $x_e^{(r)} = x_{e'}^{(r)} = r/2$ for all $r \in R$. Now define $F = \{s \in R : y_r(e) = 1\}$. Then

$$x_e^{(r)} = \int_R y_s(e) \mathbf{1}_{s \succ r} d\mu = \mu(F \cap [0, r]).$$

This implies that $\mu(F \cap [r, s]) = (s - r)/2$ for all $[r, s] \subset [0, 1]$. This however contradicts Lebesgue's density theorem, which states that a measurable set has density 1 almost everywhere in the set. Thus even in this very simple example, no Nash equilibrium exists with unsplittable flow. On the other hand, setting $y_r(e) = y_r(e') = 1/2$ for all r is a Nash equilibrium.

The nonatomic case in the classical model is comparatively much easier to define. In the classical model, a solution is defined completely by the flow vector \mathbf{x} —only the total flow on an edge is important. Such games, where each player's payoff depends only on the aggregate action of the other players, have some useful simplifying properties (see e.g. Schmeidler [19]).

It might not be clear then why this nonatomic game can be considered a suitable limit of the atomic unsplittable game when $w_j \rightarrow 0$. For some intuition, take again this simple

two-link example, but with atomic agents; suppose we have n agents (labelled from 1 to n), all of size $1/n$, going from the origin to the destination; the priority ordering is $i \succ j$ iff $i < j$. Then one possible Nash equilibrium is for odd-numbered players to take arc e , and all even-numbered players to take edge e' ; call this solution \mathcal{P}_n . Considering the set of strategies, this does not lead to any sensible limit as $n \rightarrow \infty$. For a fixed $r \in [0, 1]$, we do however have that

$$\lim_{n \rightarrow \infty} x_e^{(\lfloor rn \rfloor)}(\mathcal{P}_n) = r/2.$$

In that sense, we converge to the Nash equilibrium of the nonatomic game as we have defined it.

Generalising to congestion games is done analogously to the atomic case.

2.4 Existence of Nash equilibria

We mention a few existence and nonexistence results regarding pure Nash equilibria. First, an unsurprising negative result; in the atomic unsplittable case, allowing general priority schemes, there need not be a pure Nash equilibrium. In particular, consider the fixed priority game depicted in Figure 2. The edges in this network are undirected, and flow in either direction contributes to the congestion on an edge (we will return to this point shortly). There are two users, each of size 1, with source-destination pairs (s_1, t_1) and (s_2, t_2) respectively. All edges have cost function $f_e(x) = x$. The priorities on each edge are shown in the figure. It is easy to see that no matter which direction each of the two players choose to route their flow, the player with lower priority on the single edge these routes have in common will have an incentive to change to the other route. Thus the game has no pure Nash equilibria. Since this example is unweighted, this is in contrast with the classical mode; unweighted atomic congestion games always have a pure Nash equilibrium (but weighted ones need not) [15].




Figure 2: A fixed-priority game with no pure Nash equilibria.

Of course, we have not explicitly allowed undirected edges in our model. But we can replace each of the undirected edges in the construction with the widget shown in Figure 3; v and w represent the endpoints of the replaced edge (this is a standard technique; see e.g. [1]).

Now for a positive result: in the global priority model, even in the atomic unsplittable case, there is always a pure Nash equilibrium (as long as the cost functions are at least nonnegative and increasing). This can be seen in the atomic case by an explicit algorithm to construct the Nash: simply go through the agents in priority order, and route each along




Figure 3: Widget to imitate an undirected edge $e = (v, w)$ with directed edges (dashed arcs have zero cost).

a shortest path given the congestion effects of the higher priority agents that have already been routed. In the single-commodity case, the timestamp game and the fixed priority game are equivalent—the players can be ordered by their starting time. Our conjecture is that existence is guaranteed even in the multicommodity case for the timestamp game.

In the nonatomic version of our game, existence of Nash equilibria is guaranteed under some continuity conditions.

Definition 1. A nonatomic game is called **continuous** if all the cost functions are continuous, $x_e^{(r)}(\mathcal{P})$ depends continuously on \mathcal{P} , and

$$\{s \in R : \ell_s(\mathcal{P}) < \ell_r(\mathcal{P})\}$$

is measurable for all feasible solutions \mathcal{P} and $r \in R$.

Theorem 2.1. A continuous nonatomic game always has a pure Nash equilibrium.

Proof. Consider the strategy set of player r ; it is the set of all unit $s_r - t_r$ flows, and so forms a convex and compact subset of \mathbb{R}^E . The latency experienced by player r depends continuously on \mathcal{P} , because the game is continuous. This, along with the measurability requirement, guarantees that the conditions for Theorem 1 from Schmeidler [19] are satisfied, and so an equilibrium exists. Since each player’s strategy set is already convex, there is no need to consider mixed equilibria. \square

2.5 A correspondence with the classical model

In this section and in following ones, we will use the standard term *Wardrop equilibrium* when referring to Nash equilibria in a classical nonatomic congestion game. This is simply to aid in distinguishing between the classical and priority-based models.

The optimal flows in the priority-based model can be related to the classical model:

Lemma 2.2. Given an instance $G = (V, E)$ of the (atomic or nonatomic) priority-based network game with cost functions f_e , optimal flows are exactly the same as the optimal flows in the classical game on the same network, but with cost functions

$$\hat{f}_e(x) = \frac{1}{x} \int_0^x f_e(z) dz. \quad (6)$$

Proof. This follows by noting that the cost of a flow \mathbf{x} in the priority-based model,

$$C(\mathbf{x}) = \sum_{e \in E} \int_0^{x_e} f_e(x) dx,$$

is exactly the same as the cost induced in the classical model with cost functions \hat{f}_e :

$$\hat{C}(\mathbf{x}) = \sum_{e \in E} \hat{f}_e(x_e) x_e = \sum_{e \in E} \int_0^{x_e} f_e(x) dx.$$

□

Corollary 2.3. *In a nonatomic priority-based game, the optimal flows are exactly the Wardrop equilibria of the classical network game on the same network, with the same cost functions.*

Proof. The result follows directly from the following characterisation of optimal flows in the classical model [4, 14, 17]:

A flow \mathbf{x} is optimal for a classical nonatomic game with continuously differentiable, semiconvex² cost functions \hat{f}_e iff it is a Wardrop equilibrium for a game on the same network, where the cost functions are replaced by

$$f_e^*(y) = \frac{d}{dy} (y \cdot \hat{f}_e(y)).$$

But if the \hat{f}_e 's are defined as in Equation (6), then $f_e^*(y) = f_e(y)$, and the result follows. □

3 The price of anarchy of nonatomic agents

The central topic of this paper is the question: how bad can the cost of a Nash equilibrium be compared to the cost of an optimal solution? The price of anarchy is a quantitative answer to this question:

Definition 2. *The (pure) price of anarchy of an instance is the ratio between the cost of the worst possible pure Nash equilibrium, and the cost of the optimal solution.*

An analogous definition can be made for mixed Nash equilibria; however, we will only consider pure Nash equilibria in this paper. As such, we will usually omit “pure”. We will also use the term *worst-case price of anarchy* in relation to a class of possible instances (for example, all instances with linear cost functions) to refer to the supremum of the price of anarchy over all these instances.

Analysing the price of anarchy is easier in the nonatomic case, where there are an infinite number of players, each controlling a negligible amount of flow. The special case of single-commodity networks (particularly parallel link networks) give very different results to general networks, and we discuss these separately.

²A function $f(y)$ is semiconvex iff $yf(y)$ is convex.

3.1 Single-commodity networks

Single-commodity networks refer to the case where all agents have the same source s and destination t . We only require the cost functions be continuous, nonnegative and increasing for the following results.

Observation 1. *For a single-commodity game with nonatomic agents, any flow \mathbf{x} where all of the $s - t$ paths with non-zero flow are shortest paths (where length is determined by the metric $l_e = f_e(x_e)$) is a Wardrop equilibrium in the classical model, and hence by Corollary 2.3, an optimal flow in the priority-based model. In other words, a flow satisfying*

$$\sum_{e \in P} f_e(x_e) \leq \sum_{e \in P'} f_e(x_e). \quad (7)$$

for any $s - t$ path P with $x_P > 0$, and all $s - t$ paths P' , is optimal.

A particularly simple class of networks of this type are parallel link networks, which consist only of a source node, a sink node, and some number of links between them. We show the following:

Theorem 3.1. *For parallel link networks with nonatomic agents and any choice of priority scheme, the price of anarchy is one.*

Proof. Let \mathcal{P} be an arbitrary Nash equilibrium. Consider Equation (5). In our case, it can be written

$$l_r(\mathcal{P}) \leq f_{e'}(x_{e'}^{(r)}) \quad \text{for all } e' \in E, \quad (8)$$

for all players r . Now for each link e , either $x_e = 0$ or there is a player r such that $P_r = e$ and $x_e^{(r)} = x_e$. Equation (8) then yields

$$f_e(x_e) \leq f_{e'}(x_{e'}) \quad \text{for all } e' \in E.$$

Hence the result follows by Observation 1. \square

These results are in contrast to the classical model, where Pigou’s two-link network yields the largest possible price of anarchy in most cases [16]. If we consider general single-commodity instances however, the price of anarchy can be larger than one, even for the global priority scheme. Consider the example shown in Figure 4, which is recognisable as a standard example of “Braess’s paradox”. There is one unit of demand between the source s and the sink t . The optimal solution is to send half the flow along the top path, and half along the bottom, with no flow using the zero cost “shortcut” edge. This solution has cost $C(OPT) = 1 + 1/(d + 1)$.

With the global priority model, the unique Nash equilibrium uses the shortcut edge for the higher priority half of the flow. The lower priority half of the flow then splits between the upper and lower paths, but must now pay a larger cost; the non-constant edges now have a total load of $3/4$. The Nash has cost $C(Nash) = 1/2 + (3/2)^{d+1}/(d + 1)$. So the price of anarchy is

$$\rho = \frac{1/2 + (3/2)^{d+1}/(d + 1)}{1 + 1/(d + 1)};$$




Figure 4: A single-commodity game with price of anarchy larger than one.

this grows exponentially with d . For linear cost functions, this gives $\rho = 13/12$.

On the positive side, we can extend the parallel link result slightly for the global priority model:

Theorem 3.2. *Consider an instance where the underlying network is a series-parallel graph, with all demand between the source and sink, with cost functions that are nondecreasing and nonnegative. Then in the global priority model, the price of anarchy is one.*

The proof is a simple induction based on Theorem 3.1, and we omit it. (*TODO: Is that OK?*)

3.2 Multicommodity networks

We now investigate the price of anarchy of the priority-based model for general networks, where the behaviour is very different from the single-commodity case. We will obtain tight bounds for linear and polynomial cost functions.

First a useful inequality:

Theorem 3.3. *For any Nash flow \mathcal{P} , under any priority scheme,*

$$C(\mathcal{P}) \leq \sum_{e \in E} f_e(x_e(\mathcal{P}))x_e^*. \quad (9)$$

where \mathbf{x}^* is an arbitrary flow (in particular, it may be an optimum flow).

Proof. We will use just x_e to denote $x_e(\mathcal{P})$. Let $\mathcal{P}^* = \{P_r^* : r \in R\}$ be some assignment of paths to players that obtains the flow \mathbf{x}^* (so formally, it is a valid routing such that $\int_{r: e \in P_r^*} 1 \, d\mu = x_e^*$). Apply (5) with $P_r' = P_r^*$:

$$\begin{aligned} \ell_r(\mathcal{P}) &\leq \ell_r(\mathcal{P}^{(r)}) \\ &= \sum_{e \in P_r^*} f_e(x_e^{(r)}(\mathcal{P}^{(r)})), \end{aligned}$$

where $\mathcal{P}^{(r)} = \mathcal{P} \setminus P_r \cup P_r^*$; we have used Equation (4). Now clearly $x_e^{(r)}(\mathcal{P}^{(r)}) \leq x_e$, so

$$\ell_r(\mathcal{P}) \leq \sum_{e \in P_r^*} f_e(x_e).$$

Thus

$$\begin{aligned}
C(\mathcal{P}) &= \int_R \ell_r(\mathcal{P}) dr \\
&\leq \int_R \sum_{e \in P_r^*} f_e(x_e) dr \\
&= \sum_{e \in E} f_e(x_e) \int_R (\mathbf{1}_{e \in P_r^*}) dr \\
&= \sum_{e \in E} f_e(x_e) x_e^*.
\end{aligned}$$

□

It is interesting to compare this to the classical model, where the variational inequality

$$\sum_{e \in E} f_e(x_e) x_e \leq \sum_{e \in E} f_e(x_e) x'_e$$

holds if and only if \mathbf{x} is a Wardrop equilibrium; note that the left hand side is the cost of the flow \mathbf{x} in the classical model. In our model, the inequality is necessary, but not sufficient.

Let us now find an upper bound in the case of linear cost functions. The result is superseded by the more general polynomial case considered next, but the proof in the linear case is simpler and more transparent.

Theorem 3.4. *In the nonatomic case with linear cost functions, 4 is an upper bound on the price of anarchy.*

Proof. Note the following, for any flow vector \mathbf{x}' :

$$\begin{aligned}
\sum_{e \in E} f_e(x'_e) x'_e &= 2 \sum_{e \in E} \frac{1}{2} a_e x_e'^2 + \frac{1}{2} b_e x'_e \\
&\leq 2 \sum_{e \in E} \int_0^{x'_e} a_e x + b_e dx \\
&= 2C(\mathbf{x}').
\end{aligned} \tag{10}$$

Beginning with the result of Theorem 3.3, we use a technique from [6], which they used to give a short proof of the classical price of anarchy result.

$$\begin{aligned}
C(\mathcal{P}) &\leq \sum_{e \in E} f_e(x_e) x_e^* \\
&= \sum_{e \in E} f_e(x_e^*) x_e^* + \sum_{e \in E} (f_e(x_e) - f_e(x_e^*)) x_e^* \\
&\leq 2C(\mathcal{P}^*) + \sum_{e \in E: x_e \geq x_e^*} (f_e(x_e) - f_e(x_e^*)) x_e^* \quad \text{from (10)}.
\end{aligned}$$

Now consider Figure 5.




Figure 5: A visual proof that $(f_e(x_e) - f_e(x_e^*))x_e^* \leq \frac{1}{4}f_e(x_e)x_e$.

Clearly the area of the grey rectangle is at most $\frac{1}{4}$ of the area of the large rectangle. Thus we obtain

$$\begin{aligned} C(\mathcal{P}) &\leq 2C(\mathcal{P}^*) + \frac{1}{4} \sum_{e \in E} f_e(x_e)x_e \\ &\leq 2C(\mathcal{P}^*) + \frac{1}{2}C(\mathcal{P}), \end{aligned}$$

again using (10) in the final step. Thus $C(\mathcal{P})/C(\mathcal{P}^*) \leq 4$, as required. \square

We now extend this result to polynomial cost functions.

Theorem 3.5. *For the nonatomic case with polynomial cost functions of maximum degree d , there is an upper bound of $(d+1)^{d+1}$ for the price of anarchy.*

Proof. The proof uses a generalisation of the technique used to prove Theorem 3.4. Let $\alpha \geq 1$ be a constant to be chosen later. Let $f_e(x) = \sum_{i=0}^d a_{e,i}x^i$. We have

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{e \in E} f_e(x_e)x_e^* \\ &= \alpha \sum_{e \in E} f_e(x_e^*)x_e^* + \sum_{e \in E} (f_e(x_e) - \alpha f_e(x_e^*))x_e^* \\ &\leq \alpha(d+1)C(\mathcal{P}^*) + \sum_{e \in E: f_e(x_e) \geq \alpha f_e(x_e^*)} (f_e(x_e) - \alpha f_e(x_e^*))x_e^* \end{aligned} \quad (11)$$

Now consider:

$$\begin{aligned}
\frac{(f_e(x_e) - \alpha f_e(x_e^*))x_e^*}{f_e(x_e)x_e} &= \frac{x_e^*}{x_e} - \alpha \cdot \frac{f_e(x_e^*)x_e^*}{f_e(x_e)x_e} \\
&\leq \frac{x_e^*}{x_e} - \alpha \cdot \min_{0 \leq i \leq d} \frac{a_{e,i}x_e^{*i+1}}{a_{e,i}x_e^{i+1}} \\
&= \frac{x_e^*}{x_e} - \alpha \left(\frac{x_e^*}{x_e}\right)^{d+1} \quad (\text{since } x_e^* \leq x_e).
\end{aligned}$$

Since the maximum value of the function $\phi - \alpha\phi^{d+1}$ occurs at $\phi_m = (\alpha(d+1))^{-1/d}$, an easy calculation yields

$$(f_e(x_e) - \alpha f_e(x_e^*))x_e^* \leq \frac{d}{d+1} \cdot \frac{1}{(\alpha(d+1))^{1/d}} \cdot f_e(x_e)x_e.$$

Substituting into (11), and using $\sum_{e \in E} f_e(x_e)x_e \leq (d+1)C(\mathcal{P})$, it follows that

$$\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} \leq \frac{\alpha(d+1)}{1 - d(\alpha(d+1))^{-1/d}}.$$

Now set $\alpha = (d+1)^{d-1}$; this gives

$$\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} \leq (d+1)^{d+1}.$$

□

Having obtained an upper bound, we now show that it cannot be improved, by demonstrating how to construct a game with price of anarchy arbitrarily close to this upper bound. We will need (and again later) the following useful lemma:

Lemma 3.6. *For $a, b \geq 0$, $r \geq 1$ and $0 < \gamma < 1$,*

$$(a+b)^r \leq \gamma^{1-r}a^r + (1-\gamma)^{1-r}b^r. \quad (12)$$

Proof.

$$\begin{aligned}
(a+b)^r &= \left(\gamma \left(\frac{a}{\gamma}\right) + (1-\gamma) \left(\frac{b}{1-\gamma}\right) \right)^r \\
&\leq \gamma \left(\frac{a}{\gamma}\right)^r + (1-\gamma) \left(\frac{b}{1-\gamma}\right)^r \quad (\text{by convexity}) \\
&= \gamma^{1-r}a^r + (1-\gamma)^{1-r}b^r.
\end{aligned}$$

□

Theorem 3.7. *For the nonatomic case with polynomial cost functions of maximum degree d , there is a lower bound of $(d+1)^{d+1}$ for the worst-case price of anarchy in the global priority model.*

Proof sketch. Some calculations have been omitted; the full proof is given in the appendix.

Consider a network of the form shown in Figure 6. There are two types of latency function in the network. Each link of the form (s_i, s_{i+1}) has latency zero, and each link $e_i = (s_i, t)$ has latency $l_e(x) = x^d/i$. We have a large number of infinitesimally small agents, all trying to get to t from one of the s_i 's. The total amount of traffic originating at each s_i is unity. In addition, for all $j < i$ all agents originating at s_i have higher priority than agents originating at s_j . Agents originating at the same vertex are indistinguishable, except for some fixed priority ordering among them.




Figure 6: Lower bound construction for polynomial cost functions.

Let \mathcal{P}_n and \mathcal{P}_n^* be respectively the Nash equilibria and optimum solutions of this construction for each n .

Any agent is unaffected by the choices of lower priority agents, so we can calculate the Nash by working from the highest priority agents (i.e those starting from s_n) to the lowest (starting at s_1). Define $x_{i,j}$ to be the flow on the edge (s_i, t) after all the players with origins in $\{s_j, s_{j+1}, \dots, s_n\}$ have played; in addition, define $x_{i,n+1} = 0$. Let $y_j = f_{e_j}(x_{j,j})$. It is easy to see that the Nash condition implies that

$$f_{e_i}(x_{i,j}) = f_{e_j}(x_{j,j}) = y_j \quad \text{for all } i \leq j.$$

From this, it can be shown that

$$C(\mathcal{P}_n) \geq (d+1)^d \sum_{j=1}^n j^{1/d} \left((j+1)^{-1/d} - (n+2)^{-1/d} \right)^{d+1}. \quad (13)$$

Applying Lemma 3.6 to (13) with $a = (j+1)^{-1/d} - (n+2)^{-1/d}$, $b = (n+2)^{-1/d}$ and $r = d+1$, we obtain that for any constant $0 < \gamma < 1$,

$$C(\mathcal{P}_n) \geq (d+1)^d \left(\gamma^d \sum_{j=1}^n j^{-1} \left(1 + \frac{1}{j}\right)^{-1-1/d} - \left(\frac{\gamma}{1-\gamma}\right)^d (n+2)^{-1-1/d} \sum_{j=1}^n j^{1/d} \right).$$

Some calculations then show that

$$C(\mathcal{P}_n) \geq \gamma^d (d+1)^d H_n - D_\gamma,$$

where D_γ is a constant that depends on γ , but not n .

Let \mathcal{P}'_n be the solution obtained by sending all flow from s_i through arc e_i for each i . This yields a cost of

$$C(\mathcal{P}'_n) = \frac{1}{d+1} \sum_{i=1}^n \frac{1}{i} = \frac{H_n}{d+1},$$

which is an upper bound on the cost of the optimal solution \mathcal{P}_n^* .

We thus get a bound for the price of anarchy for any given n :

$$\frac{C(\mathcal{P}_n)}{C(\mathcal{P}_n^*)} \geq \frac{\gamma^d(d+1)^d H_n - D\gamma}{(d+1)^{-1} H_n} = \gamma^d(d+1)^{d+1} - \frac{D\gamma(d+1)}{H_n}.$$

Consequently, letting $n \rightarrow \infty$, we find that $\gamma^d(d+1)^{d+1}$ is a lower bound for the price of anarchy. Finally, since γ was an arbitrary constant strictly less than 1, we send $\gamma \rightarrow 1$ to obtain $(d+1)^{d+1}$ as a lower bound. \square

Note that the priority ordering used in the above construction can also easily be produced in the timestamp case. Let any agent originating at s_i have an earlier start-time τ_i than any agent originating at s_j , for all $j < i$. The relative ordering of timestamps for agents originating at the same vertex is unimportant. We may assume that that start-times are measured to an arbitrary precision so that ties do not arise.

Combining the previous two theorems, we have an exact value of $(d+1)^{d+1}$ for the worst-case price of anarchy of our model with polynomial latency functions.

4 The price of anarchy for unsplittable atomic agents

In this section we consider the case of unsplittable agents. We will present a tight upper bound for the linear case, as well as a number of matching lower bound constructions for different priority schemes. For polynomial cost functions, we will only provide an upper bound.

Denote the set of players by J . As usual let \mathcal{P} be a Nash flow, \mathcal{P}^* be an unsplittable optimal flow, and define $\mathcal{P}^{(j)} = \mathcal{P} \setminus P_j \cup P_j^*$, where everyone follows \mathcal{P} except player j . We begin with a useful inequality that holds for any Nash flow \mathcal{P} . Using equations (2) and (3),

$$\begin{aligned} C_j(\mathcal{P}) &\leq C_j(\mathcal{P}^{(j)}) \\ &= \sum_{e \in P_j^*} \int_{x_e^{(j)}(\mathcal{P}^{(j)})}^{x_e^{(j)}(\mathcal{P}^{(j)})+w_j} f_e(x) dx. \end{aligned}$$

But $x_e^{(j)}(\mathcal{P}^{(j)}) \leq x_e$, so

$$C_j(\mathcal{P}) \leq \sum_{e \in P_j^*} \int_{x_e}^{x_e+w_j} f_e(x) dx.$$

Summing over all j yields

$$\begin{aligned}
C(\mathcal{P}) &\leq \sum_{j \in J} \sum_{e \in P_j^*} \int_{x_e}^{x_e + w_j} f_e(x) dx \\
&= \sum_{e \in E} \sum_{j: e \in P_j^*} \int_{x_e}^{x_e + w_j} f_e(x) dx.
\end{aligned} \tag{14}$$

4.1 Linear cost functions

Theorem 4.1. *In the unsplittable case with linear latency functions, the price of anarchy is at most $3 + 2\sqrt{2}$.*

Proof. Let \mathcal{P} and \mathcal{P}^* be a Nash flow and an optimal unsplittable flow respectively. Writing Equation (14) in the linear case with $f_e(x) = a_e x + b_e$, we obtain

$$\begin{aligned}
C(\mathcal{P}) &\leq \sum_{e \in E} \sum_{j: e \in P_j^*} ((a_e x_e + b_e)w_j + \frac{1}{2}a_e w_j^2) \\
&\leq \sum_{e \in E} ((a_e x_e + b_e)x_e^* + \frac{1}{2}a_e x_e^{*2}) \\
&= \sum_{e \in E} a_e x_e x_e^* + \sum_{e \in E} (\frac{1}{2}a_e x_e^* + b_e)x_e^*.
\end{aligned}$$

We now apply the Cauchy-Schwarz inequality to the first term to obtain

$$\begin{aligned}
C(\mathcal{P}) &\leq \sqrt{\sum_{e \in E} a_e x_e^2 \cdot \sum_{e \in E} a_e x_e^{*2}} + C(\mathcal{P}^*) \\
&\leq \sqrt{2C(\mathcal{P}) \cdot 2C(\mathcal{P}^*)} + C(\mathcal{P}^*).
\end{aligned}$$

Let $\alpha = \frac{C(\mathcal{P})}{C(\mathcal{P}^*)}$. The above gives us $\alpha \leq 2\sqrt{\alpha} + 1$, and so the price of anarchy is at most $3 + 2\sqrt{2} \approx 5.828$. \square

We now provide some matching lower bounds for various game variants. We begin with a weighted congestion game construction. We will require different priority orderings on different edges.

Let the set of items be $I = \{1, 2, 3, \bar{1}, \bar{2}, \bar{3}\}$, and the players be $J = \{1, 2, 3, \bar{1}, \bar{2}, \bar{3}\}$. One should think of the barred items as mirror copies of the originals, and the barred players as reflected copies. We also define $\bar{\bar{1}} = 1$, etc. and $\overline{\{A\}} = \{\bar{A}\}$.

We define the set of strategies for player $j \in J$ as $\mathcal{S}_j = \{S_j, S_j^*\}$ where

$$\begin{aligned}
S_1 &= \{1, 2, 3\} & S_1^* &= \{\bar{1}\} \\
S_2 &= \{1, 2\} & S_2^* &= \{\bar{2}\} \\
S_3 &= \{1, 2\} & S_3^* &= \{\bar{3}\}
\end{aligned}$$

and $\mathcal{S}_{\bar{j}} = \bar{\mathcal{S}}_j$ for $j = 1, 2, 3$.

The player weights w_j are given by

$$\begin{aligned} w_1 = w_{\bar{1}} = w_2 = w_{\bar{2}} &= 1, \\ w_3 = w_{\bar{3}} &= \sqrt{2} - 1. \end{aligned}$$

The priority ordering is

$$\begin{aligned} 1 \succ 2 \succ 3 \succ \bar{1} \succ \bar{2} \succ \bar{3} & \quad \text{for items } 1, 2, 3 \\ \bar{1} \succ \bar{2} \succ \bar{3} \succ 1 \succ 2 \succ 3 & \quad \text{for items } \bar{1}, \bar{2}, \bar{3} \end{aligned}$$

The cost function for item i is $f_i(x) = a_i x$ where

$$a_1 = a_{\bar{1}} = \frac{2\sqrt{2}}{3 + 2\sqrt{2}}, \quad (15)$$

$$a_2 = a_{\bar{2}} = \frac{3}{3 + 2\sqrt{2}}, \quad (16)$$

$$a_3 = a_{\bar{3}} = 2\sqrt{2} - 1. \quad (17)$$

We claim that if all players pick strategy S_j , we have a Nash equilibrium. To show this, we need to show that no player has an incentive to switch to S_j^* . Note that the priority ordering is such that a player would have the lowest priority on an item if they switched.

Let the cost for player j when all players are playing S_j be C_j . Some easy calculations yield:

$$\begin{aligned} C_1 &= \int_0^{w_1} f_1(x) + f_2(x) + f_3(x) dx = \sqrt{2} \\ C_2 &= \int_{w_1}^{w_1+w_2} f_1(x) + f_2(x) dx = \frac{3}{2} \\ C_3 &= \int_{w_1+w_2}^{w_1+w_2+w_3} f_1(x) + f_2(x) dx = \sqrt{2} - \frac{1}{2}. \end{aligned}$$

Let \tilde{C}_j be the cost player j pays upon switching. Then

$$\begin{aligned} \tilde{C}_1 &= \int_{w_{\bar{1}}+w_{\bar{2}}+w_{\bar{3}}}^{w_{\bar{1}}+w_{\bar{2}}+w_{\bar{3}}+w_1} f_{\bar{1}}(x) dx = \sqrt{2} \\ \tilde{C}_2 &= \int_{w_{\bar{1}}+w_{\bar{2}}+w_{\bar{3}}}^{w_{\bar{1}}+w_{\bar{2}}+w_{\bar{3}}+w_2} f_{\bar{2}}(x) dx = \frac{3}{2} \\ \tilde{C}_3 &= \int_{w_{\bar{1}}}^{w_{\bar{1}}+w_3} f_{\bar{3}}(x) dx = \sqrt{2} - \frac{1}{2}. \end{aligned}$$

So none of players 1, 2, 3 have an incentive to switch, and by symmetry neither do players $\bar{1}, \bar{2}, \bar{3}$. So we do have a Nash equilibrium. The optimal strategy is for all players to play S_j^* . Now notice that the utilisation of each item under the Nash is exactly $1 + \sqrt{2}$

times the utilisation under the optimal strategy. It follows that the price of anarchy is $(1 + \sqrt{2})^2 = 3 + 2\sqrt{2}$.

We can turn this into a network game, as shown in Figure 7. The dashed arcs have zero cost, and the remaining arcs are labelled to correspond with the items of the congestion game, and have the same cost functions, and the same priority orderings. The sources s_j and destinations t_j of the players are also labelled. It can easily be verified that this network game reduces to the above congestion game, and so also has a price of anarchy of $3 + 2\sqrt{2}$.




Figure 7: A network game construction with a price of anarchy of $3 + 2\sqrt{2}$.

While the above construction uses different priorities on different edges, we can use the basic idea for constructions with other priority schemes. First, let's go back to the congestion game formulation and consider the global priority game. Let N be some large integer. Let the players be

$$J = \{j_{r,s} : 1 \leq r \leq 3, 1 \leq s \leq N\}$$

and the items be

$$I = \{i_{r,s} : 1 \leq r \leq 3, 1 \leq s \leq N + 1\}.$$

We now set, for $1 \leq s \leq N$,

$$\begin{aligned} S_{j_{1,s}} &= \{i_{1,s}, i_{2,s}, i_{3,s}\} & S_{j_{1,s}}^* &= \{i_{1,s+1}\} \\ S_{j_{2,s}} &= \{i_{1,s}, i_{2,s}\} & S_{j_{2,s}}^* &= \{i_{2,s+1}\} \\ S_{j_{3,s}} &= \{i_{1,s}, i_{2,s}\} & S_{j_{3,s}}^* &= \{i_{3,s+1}\} \end{aligned}$$

The weights are

$$w_{j_{1,s}} = w_{j_{2,s}} = 1, \quad w_{j_{3,s}} = \sqrt{2} - 1.$$

The global priority ordering is

$$j_{1,N} \succ j_{2,N} \succ j_{3,N} \succ j_{1,N-1} \succ j_{2,N-1} \cdots \succ j_{2,1} \succ j_{3,1}.$$

For $s \leq N$, we set the cost functions as before, i.e. $f_{i_r,s} = a_r$ for $r = 1, 2, 3$, with the a_r defined in Equation (15) to (17). The exception is the final group of items, which nobody plays at Nash; thus we have to make it more expensive to ensure that players $j_{1,N}$, $j_{2,N}$ and $j_{3,N}$ do not have an incentive to switch. So simply set

$$f_{i_r,N+1}(x) = C_{j_r,N}(\mathcal{P}).$$

Without this imperfection, the price of anarchy would be exactly as before, since we would simply have N copies instead of two. The addition of the final group reduces the price of anarchy slightly. However, as we increase N to infinity, the effect of this on the total social cost becomes negligible. So we have a construction that yields a price of anarchy of $3 + 2\sqrt{2} - \epsilon$, for any $\epsilon > 0$; thus the upper bound is still tight in the global priority game.

This construction can be turned into a network game fairly easily, in much the same way as before (we omit the details); once we have this, we can also obtain a timestamp game construction by judicious choice of starting times. In particular, if we set the start times as

$$\tau_{i_1,i} = (N - i)K, \quad \tau_{i_2,i} = (N - i)K + 1, \quad \tau_{i_3,i} = (N - i)K + 2,$$

where K is sufficiently large, we clearly end up with the same priority ordering.

Next consider the unweighted case, where $w_j = 1$ for all players j . We give a tight result here also.

Theorem 4.2. *For unweighted agents and linear cost functions, the price of anarchy is at most $17/3$.*

Proof. We need the following easily proven lemma:

Lemma 4.3. *Let $i, j \geq 0$ be integers. Then*

$$(2i + 1)j \leq \frac{2}{5}i^2 + \frac{17}{5}j^2. \quad \square$$

Now:

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{e \in E} (a_e x_e + b_e) x_e^* + \sum_{e \in E} \sum_{j: e \in P_j^*} \frac{1}{2} a_e w_i^2 \\ &= \sum_{e \in E} a_e \left(x_e + \frac{1}{2}\right) x_e^* + \sum_{e \in E} b_e x_e^* && \text{(using } w_i = w_i^2\text{)} \\ &\leq \sum_{e \in E} \frac{1}{2} a_e \left(\frac{2}{5} x_e^2 + \frac{17}{5} x_e^{*2}\right) + \sum_{e \in E} b_e x_e^* && \text{(using Lemma 4.3)} \\ &\leq \frac{2}{5} C(\mathcal{P}) + \frac{17}{5} C(\mathcal{P}^*). \end{aligned}$$

Thus

$$\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} \leq \frac{17/5}{1 - 2/5} = \frac{17}{3}.$$

□

The following construction shows that this upper bound is tight. Let

$$I = \{1, 2, 3, 4, \bar{1}, \bar{2}, \bar{3}, \bar{4}\} \quad \text{and} \quad J = \{1, 2, 3, \bar{1}, \bar{2}, \bar{3}\}$$

be the items and players respectively. The strategies are

$$\begin{aligned} S_1 &= \{1, 2, 3, 4\} & S_1^* &= \{\bar{2}, \bar{3}\} \\ S_2 &= \{1, 2, 3, 4\} & S_2^* &= \{\bar{4}\} \\ S_3 &= \{1, 2\} & S_3^* &= \{\bar{1}\} \end{aligned}$$

and $\mathcal{S}_j = \bar{\mathcal{S}}_j$ for $j = 1, 2, 3$. The priority ordering is

$$\begin{aligned} 1 \succ 2 \succ 3 \succ \bar{1} \succ \bar{2} \succ \bar{3} & \quad \text{on items } 1, 2, 3, 4 \\ \bar{1} \succ \bar{2} \succ \bar{3} \succ 1 \succ 2 \succ 3 & \quad \text{on items } \bar{1}, \bar{2}, \bar{3}, \bar{4} \end{aligned}$$

The cost function for item i is $f_i(x) = a_i x$ where

$$a_1 = \frac{5}{7}, \quad a_2 = \frac{2}{7}, \quad a_3 = \frac{1}{5} \quad \text{and} \quad a_4 = \frac{9}{5}$$

(and symmetrically for the remaining items).

Defining \mathcal{P} and \mathcal{P}^* as usual, it can easily be verified that

$$\begin{aligned} C_1(\mathcal{P}) &= \frac{3}{2} = C_1(\mathcal{P}^*) \\ C_2(\mathcal{P}) &= \frac{9}{2} = C_2(\mathcal{P}^*) \\ C_3(\mathcal{P}) &= \frac{5}{2} = C_3(\mathcal{P}^*) \end{aligned}$$

Hence \mathcal{P} is a Nash equilibrium. The price of anarchy is then easily calculated to be $17/3$, as required.

Again, it is straightforward to convert this to a network game. Variations for more restrictive priority schemes are possible using the same approach as for the weighted case.

4.2 Polynomial cost functions

We will give only an upper bound for the polynomial case. For the lower bound, we will simply note that the $(d+1)^{d+1}$ value obtained in the nonatomic case still applies by using the same construction with sufficiently small agents. Clearly a better construction is possible, and the upper bound is also unlikely to be tight. We will not discuss the unweighted variations here. The proof uses a combination of the techniques used earlier, and is given in the appendix.

Theorem 4.4. *The price of anarchy is $O(2^d d^d)$ in the unsplittable atomic case with polynomial cost functions of maximum degree d .*

Further work In the atomic version of the model, a pure Nash equilibrium need not always exist. It should be easy to extend our results to handle mixed strategy Nash equilibria. It is not clear however if these are of much interest in our model; another avenue would be

to investigate in these games the so called “price of sinking”, introduced by Goemans et al. [8].

We have considered only linear and polynomial latency functions. Other cost functions are of course possible, and may be of interest. All of the cost functions we have considered are convex. This is generally assumed for traffic congestion, but may not apply for all applications of this model. However, utilising concave cost functions in the service game could sometimes be appropriate, e.g. manufacturers facing decreasing marginal costs. This would also allow for the modelling of other complex interactions between companies and manufacturers; for example, manufacturers could pass on the gains from decreasing marginal costs to more favoured customers.

As noted earlier, it would be very interesting to consider selfish routing in a dynamic flow model, in order to obtain a much more realistic version of the timestamp game.

Acknowledgements We would like to thank George Karakostas, Mohammad Mahdian and Nicolás Stier-Moses for interesting discussions on this topic, and Tobias Harks for pointing out an error in an earlier version of this manuscript. We would also like to thank the anonymous referees for their very helpful comments and suggestions.

References

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 218–229, 2006.
- [3] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.
- [4] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [5] G. Christodoulou and E. Koutsoupias. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.
- [6] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games. In *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 167–181, 2005.
- [7] L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6(3):419–433, 1958.
- [8] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

- [9] T. Harks, S. Heinz, and M. E. Pfetsch. Competitive online multicommodity routing. In *Proceedings of the 4th Workshop on Approximation and Online Algorithms (WAOA)*, pages 240–252, 2006.
- [10] T. Harks and L. A. Végh. Selfish routing with online demands. In *Proceedings of the 4th Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN)*, pages 27–45, 2007.
- [11] E. Köhler and M. Skutella. Flows over time with load-dependent transit times. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 174–183, 2002.
- [12] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
- [13] N. Olver. The price of anarchy and a priority-based model of routing. Master’s thesis, McGill University.
- [14] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.
- [15] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, V2(1):65–67, 1973.
- [16] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 428–437, 2002.
- [17] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [18] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [19] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.
- [20] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt II*, volume 1, pages 325–378, 1952.

Appendix

Proof of Theorem 3.7. Recall the network shown in Figure 6. Each link of the form (s_i, s_{i+1}) has latency zero, and each link $e_i = (s_i, t)$ has latency $l_e(x) = x^d/i$. For each i , there is one unit of traffic going from s_i to t , and for all $j < i$, all agents originating at s_i have higher priority than agents originating at s_j . Agents originating at the same vertex are indistinguishable, except for some fixed priority ordering among them.

Any agent is unaffected by the choices of lower priority agents, so we can calculate the Nash by working from the highest priority agents (i.e those starting from s_n) to the lowest (starting at s_1). Define $x_{i,j}$ to be the flow on the edge (s_i, t) after all the players with

origins in $\{s_j, s_{j+1}, \dots, s_n\}$ have played; in addition, define $x_{i,n+1} = 0$. Let $y_j = f_{e_j}(x_{j,j})$. The Nash condition implies that

$$f_{e_i}(x_{i,j}) = f_{e_j}(x_{j,j}) = y_j \quad \text{for all } i \leq j.$$

Inverting this gives

$$x_{i,j} = (iy_j)^{1/d} \quad \text{for all } i \leq j.$$

Now since the total flow from s_j is 1, we have $\sum_{i=1}^j (x_{i,j} - x_{i,j+1}) = 1$, so

$$\sum_{i=1}^j \left((iy_j)^{1/d} - (iy_{j+1})^{1/d} \right) = 1.$$

Define $h_k := \sum_{i=1}^k i^{1/d}$. Then

$$y_j^{1/d} = h_j^{-1} + y_{j+1}^{1/d}.$$

Thus

$$y_j^{1/d} = \sum_{k=j}^n h_k^{-1},$$

as $y_{n+1} = 0$.

Since the sequence $(i^{1/d})_{i=1}^j$ is increasing, we have the bound

$$h_k \leq \int_0^{k+1} x^{1/d} dx = \frac{d}{d+1} (k+1)^{1+1/d}.$$

Hence

$$\begin{aligned} y_j^{1/d} &\geq \frac{d+1}{d} \sum_{k=j}^n (k+1)^{-(1+1/d)} \\ &\geq \frac{d+1}{d} \int_j^{n+1} (x+1)^{-(1+1/d)} dx \\ &= (d+1) \left((j+1)^{-1/d} - (n+2)^{-1/d} \right) \end{aligned}$$

We can now get a lower bound on the cost of the Nash flow \mathcal{P}_n . Since the flow from s_1, s_2, \dots, s_{j-1} does not use edge e_j , the total flow along edge e_j at Nash is $x_{j,j}$. Thus

$$\begin{aligned} C(\mathcal{P}_n) &= \sum_{j=1}^n \int_0^{x_{j,j}} f_{e_j}(x) dx \\ &= \sum_{j=1}^n \frac{1}{j(d+1)} x_{j,j}^{d+1} \\ &= \frac{1}{d+1} \sum_{j=1}^n j^{1/d} y_j^{1+1/d} \\ &\geq \frac{1}{d+1} \sum_{j=1}^n j^{1/d} \left((d+1) \left((j+1)^{-1/d} - (n+2)^{-1/d} \right) \right)^{d+1} \\ &= (d+1)^d \sum_{j=1}^n j^{1/d} \left((j+1)^{-1/d} - (n+2)^{-1/d} \right)^{d+1}. \end{aligned} \tag{18}$$

We can rewrite the statement of Lemma 3.6 as

$$a^r \geq \gamma^{r-1}(a+b)^r - \left(\frac{\gamma}{1-\gamma}\right)^{r-1} b^r.$$

Apply this to (18) with $a = (j+1)^{-1/d} - (n+2)^{-1/d}$, $b = (n+2)^{-1/d}$ and $r = d+1$ to obtain, for any constant $0 < \gamma < 1$,

$$C(\mathcal{P}_n) \geq (d+1)^d \left(\gamma^d \sum_{j=1}^n j^{-1} \left(1 + \frac{1}{j}\right)^{-1-1/d} - \left(\frac{\gamma}{1-\gamma}\right)^d (n+2)^{-1-1/d} \sum_{j=1}^n j^{1/d} \right).$$

We deal with each term separately. We have

$$\begin{aligned} (n+2)^{-1-1/d} \sum_{j=1}^n j^{1/d} &= (n+2)^{-1} \sum_{j=1}^n \left(\frac{j}{n+2}\right)^{1/d} \\ &< (n+2)^{-1} \cdot n, \end{aligned}$$

and so the second term is $O(1)$. For the first term, note that

$$\begin{aligned} j^{-1} \left(1 - \left(1 + \frac{1}{j}\right)^{-1-1/d}\right) &\leq j^{-1} \left(1 - \left(1 + \frac{1}{j}\right)^{-2}\right) \\ &\leq \frac{3}{(j+1)^2}. \end{aligned}$$

It follows that $\sum j^{-1} \left(1 - \left(1 + \frac{1}{j}\right)^{-1-1/d}\right)$ is a convergent series, and hence that

$$\sum_{j=1}^n j^{-1} \left(1 + \frac{1}{j}\right)^{-1-1/d} - H_n = O(1),$$

where H_n is the harmonic series. Thus there exists a constant D_γ , depending on γ but not n , such that

$$C(\mathcal{P}_n) \geq \gamma^d (d+1)^d H_n - D_\gamma.$$

Let \mathcal{P}'_n be the solution obtained by sending all flow from s_i through arc e_i for each i . This yields a cost of

$$C(\mathcal{P}'_n) = \frac{1}{d+1} \sum_{i=1}^n \frac{1}{i} = \frac{H_n}{d+1},$$

which is an upper bound on the cost of the optimal solution \mathcal{P}_n^* .

We thus get a bound for the price of anarchy for any given n :

$$\frac{C(\mathcal{P}_n)}{C(\mathcal{P}_n^*)} \geq \frac{\gamma^d (d+1)^d H_n - D_\gamma}{(d+1)^{-1} H_n} = \gamma^d (d+1)^{d+1} - \frac{D_\gamma (d+1)}{H_n}.$$

Consequently, letting $n \rightarrow \infty$, we find that $\gamma^d (d+1)^{d+1}$ is a lower bound for the price of anarchy. Finally, since γ was an arbitrary constant strictly less than 1, we send $\gamma \rightarrow 1$ to obtain $(d+1)^{d+1}$ as a lower bound. \square

Proof of Theorem 4.4. Let $f_e(x) = \sum_{i=0}^d a_{e,i}x^i$. Begin from Equation (14):

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{e \in E} \sum_{j: e \in P_j^*} \int_{x_e}^{x_e + w_j} f_e(x) dx \\ &\leq \sum_{e \in E} \sum_{j: e \in P_j^*} f_e(x_e + w_j) w_j \\ &= \sum_{e \in E} \sum_{j: e \in P_j^*} a_{e,0} w_j + \sum_{e \in E} \sum_{i=1}^d \sum_{j: e \in P_j^*} a_{e,i} (x_e + w_j)^i w_j. \end{aligned}$$

Now apply Lemma 3.6, with $a = x_e$, $b = w_j$, $r = i$ and γ to be determined later:

$$C(\mathcal{P}) \leq \sum_{e \in E} \sum_{j: e \in P_j^*} a_{e,0} w_j + \sum_{e \in E} \sum_{i=1}^d \sum_{j: e \in P_j^*} \left(a_{e,i} \gamma^{1-i} x_e^i w_j + a_{e,i} (1-\gamma)^{1-i} w_j^{i+1} \right).$$

Now since $\sum_{j: e \in P_j^*} w_j = x_e^*$, and hence $\sum_{j: e \in P_j^*} w_j^i \leq x_e^{*i}$ for $i \geq 1$,

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{e \in E} a_{e,0} x_e^* + \sum_{e \in E} \sum_{i=1}^d \left(a_{e,i} \gamma^{1-d} x_e^i x_e^* + a_{e,i} (1-\gamma)^{1-d} x_e^{*i+1} \right) \\ &\leq \sum_{e \in E} \sum_{i=0}^d \left(a_{e,i} \gamma^{1-d} x_e^i x_e^* + a_{e,i} (1-\gamma)^{1-d} x_e^{*i+1} \right) \\ &= \gamma^{1-d} \sum_{e \in E} f_e(x_e) x_e^* + (1-\gamma)^{1-d} \sum_{e \in E} f_e(x_e^*) x_e^* \\ &\leq \gamma^{1-d} \sum_{e \in E} f_e(x_e) x_e^* + (1-\gamma)^{1-d} (d+1) C(\mathcal{P}^*). \end{aligned}$$

The technique used for the nonatomic case is applicable to the first term (see the proof of Theorem 3.5). We thus obtain, for any $\alpha \geq 1$ and $0 < \gamma < 1$,

$$C(\mathcal{P}) \leq \gamma^{1-d} \left(\alpha(d+1) C(\mathcal{P}^*) + d(\alpha(d+1))^{-1/d} C(\mathcal{P}) \right) + (1-\gamma)^{1-d} (d+1) C(\mathcal{P}^*).$$

Thus

$$\rho \leq (d+1) \cdot \frac{\gamma^{1-d} \alpha + (1-\gamma)^{1-d}}{1 - \gamma^{1-d} d (\alpha(d+1))^{-1/d}}.$$

Now set $\alpha = 2^d d^d$ and $\gamma = 1 - \frac{1}{2d}$. Then

$$\begin{aligned} \gamma^{1-d} &= \left(1 - \frac{1}{2d}\right)^{1-d} \leq (e^{-1/2d})^{1-d} \leq e^{1/2} \\ \frac{d}{(\alpha(d+1))^{1/d}} &= 2^{-1} d^{1/d} (d+1)^{-1/d} \leq \frac{1}{2}. \end{aligned}$$

Thus

$$\begin{aligned}\rho &\leq (d+1) \cdot \frac{\sqrt{e}2^d d^d + 2^{d-1} d^{d-1}}{1 - \frac{1}{2}\sqrt{e}} \\ &= \left(\frac{\sqrt{e} + \frac{1}{2}}{1 - \frac{1}{2}\sqrt{e}} \right) 2^d d^{d-1} (d+1)\end{aligned}$$

So we have that $\rho = O(2^d d^d)$.

□