

The Price of Anarchy and a Priority-Based Model of Routing

Neil Olver

Department of Mathematics and Statistics,

McGill University, Montréal

Québec, Canada

August, 2006

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements of the degree of
Master of Science

Copyright © Neil Olver, 2006



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-32765-4

Our file Notre référence

ISBN: 978-0-494-32765-4

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

Abstract

The *price of anarchy*, a concept introduced by Koutsoupias and Papadimitriou [9], is the main topic of this thesis. It is a measure of the loss of efficiency that occurs when there is no central control over a system consisting of many “selfish” agents. We will be particularly interested in this in the context of network games, which can be used to model congestion in traffic and communication networks.

After an introduction of the relevant concepts and review of related work, we proceed with the new results of this thesis. We provide a new upper bound for the price of anarchy in the case of atomic unsplittable agents with polynomial cost functions, and demonstrate that it is tight by an explicit construction. We then introduce a new model for network routing that introduces priorities; users with a higher priority on a link will be able to traverse the link more quickly. Our model is fairly general, and allows various different priority schemes for modelling different situations. One particularly interesting version, which we have dubbed the *timestamp game*, assigns priorities according to the order of arrival at the start of the link.

We derive tight upper bounds for the price of anarchy in our model in the case of polynomial cost functions and nonatomic agents. We also obtain tight results in the unsplittable case with linear cost functions, and an upper bound with polynomial cost functions.

While we concentrate on network games, most of the results carry through to the more general class of congestion games, which we also discuss.

Résumé

Le prix de l'anarchie, un concept présenté par Koutsoupias et Papadimitriou, est la matière principale de cette thèse. Il s'agit d'une mesure de la perte d'efficacité qui se produit quand il n'y a aucun contrôle central d'un système se composant de beaucoup d'agents "égoïstes". Nous serons particulièrement intéressés par ce concept dans le contexte des jeux de réseau, qui peuvent être employés pour modéliser la congestion dans les réseaux du trafic et de transmission.

Après une introduction des concepts essentiels et une révision de la littérature pertinente, nous introduirons les nouveaux résultats de cette thèse. Nous fournirons un nouveau majorant dans le cas des agents indivisibles avec fonctions polynomiales de coût, et démontrerons, au moyen d'une construction explicite, que ce majorant exact. Nous présenterons alors un nouveau modèle pour le cheminement de réseau qui présente des priorités; les utilisateurs ayant une priorité plus élevée sur un lien pourront traverser le lien plus rapidement. Notre modèle est assez général, et permet divers arrangements prioritaires pour modéliser différentes situations. Une version particulièrement intéressante, que nous avons nommée le "timestamp game", assigne des priorités selon l'ordre d'arrivée au début du lien.

Nous dériverons les majorants exacts pour le prix de l'anarchie dans notre modèle, dans le cas de fonctions polynomiales de coût et des agents non-atomiques. Nous obtiendrons également des résultats exacts dans la situation indivisible avec des fonctions linéaires de coût, et un majorant avec des fonctions polynomiales de coût.

Bien que nous nous concentrons sur les jeux de réseau, la plupart des résultats se généralisent à la catégorie plus large des jeux de congestion.

Acknowledgments

First and foremost, I would like to thank my supervisor, Adrian Vetta. It was through him that I was first exposed to discrete mathematics at a graduate level. Without his input, I think it would have taken me much longer to discover it. I have enjoyed his supervision greatly.

As well as my supervisor, I worked with Babak Farzad on the flow-free model of Chapter 3. It has been a pleasure to work with him.

My officemates, Dave Cottrell, Olivier Dubois and Joel Phillips, have been excellent company during those (too frequent) occasions when work refused to happen. The faculty and students in both the computer science and mathematics departments have made my stay so far at McGill an exciting, stimulating, but most of all fun experience.

It would not have been possible for me to study at McGill University without the financial assistance of the Canadian government, through the Canadian Commonwealth Scholarship program.

The administrative staff of the department have been helpful always; I would like to thank Carmen Baldonado in particular for her patience with my frequent visits about all manner of things.

I would also like to specially thank my family, for their comfort and support from very far away, and Stéfanie, for hers from much closer.

Table of Contents

Abstract	i
Résumé	iii
Acknowledgments	v
Introduction	1
1 Selfish Routing	5
1.1 The model	5
1.2 The price of anarchy	7
1.3 Bounding the price of anarchy	9
1.4 Atomic games	11
1.5 Congestion games	14
2 Unsplittable Flow	17
2.1 The Nash condition	17
2.2 Linear cost functions	18
2.3 Improved bounds for the unsplittable case	20
2.3.1 The lowerbound construction	22
2.3.2 A matching upper bound	25
2.4 Unweighted games	32

3	The Flow-free Model	37
3.1	The model	37
3.2	The price of anarchy of nonatomic agents	45
3.3	Unsplittable atomic agents	52
3.3.1	Linear cost functions	53
3.3.2	Polynomial cost functions	60
3.4	Single-commodity networks	62
	Conclusion	65
A	Calculations	67

Introduction

We begin with an example all too familiar to many: traffic. As cities expand, traffic congestion has become an increasing inconvenience to many motorists, particularly those driving to and from work during heavy rush-hour traffic. A Statistics Canada survey found that the average daily commute for residents of greater Montreal was 76 minutes in 2005, compared to 62 minutes in 1992.¹ Time spent in a traffic jam is typically unproductive; the time could be better spent for work or recreation. Fuel efficiency is also decreased, leading to higher fuel costs (and pollution). So not only is traffic an annoyance, it has a clear-cut economic and environmental costs.

So what strategies should be considered for reducing commute times? An obvious answer is to build more (or wider) roads. But surprisingly, there is increasing evidence suggesting that, quite often, this can increase traffic congestion rather than reduce it. We will return to this in the next section with an emphatic example in the context of a simplified mathematical model.

Another approach is to attempt to change the behaviour of the road users. There are a number of ways to do this, for example by instigating some form of tax or fee. A prominent example of this is the London congestion charge; motorists are charged a fee for entering within the Central London area. Part of the effect of such a fee might be to reduce the total amount of traffic, by encouraging the use of public transport, cycling, etc. But let us assume for the purpose of argument that we cannot reduce the

¹Source: <http://www.statcan.ca/Daily/English/060712/d060712b.htm>

total amount of traffic – perhaps the public transport system is already at capacity. If we can manipulate motorists and somehow control the routes they take, can we increase the overall efficiency of the traffic system? For example, by reducing the traffic on a central artery?

Let's think about this from another perspective; suppose you are driving home from work at the end of a long day. All other things being equal, you will take the quickest way home. Note “quickest”, not “shortest”; you will take a detour if it means avoiding a particularly busy road. When considering the various routes available to you, your fellow motorists are unlikely to be in your thoughts. It's possible that taking a slightly longer, but less popular route would reduce the average commute times, by slightly reducing the congestion on a well-used road, and hence slightly reducing the journey time of many motorists. But you are *selfish*, in the sense that your concern is your own commute time, not the total commute time experienced by everyone. This gives us a more quantitative way of phrasing the question of the previous paragraph. How much of an improvement can we obtain if we are allowed to tell drivers which route to take, compared with the situation obtained if we allow everyone to behave selfishly? Or, to flip this around, how much do we lose if we lack control? This is the essential idea of the so-called *price of anarchy*, which is a central concept used throughout this thesis. A rigorous definition will be provided in the first chapter.

The explosion of telecommunications networks yields another motivation for the study of network routing. As the internet increasingly becomes the communication tool of choice for all types of media, including Voice over IP and streaming video, bandwidth and latency requirements are increasing. Changes to the current routing protocols and systems may be required to keep up with these changes.

Broadly speaking, this thesis will be concerned with the mathematical treatment, using tools such as game theory, of selfishness in network routing. The problem is not a new one; Roughgarden [14] provided a very accessible and comprehensive review of

the subject. However, this thesis contributes a number of novel results.

This thesis is organised as follows. In Chapter 1, we review some fairly recent results, and in particular we will discuss and rigorously define the price of anarchy, and provide proofs of some standard results in the nonatomic case of selfish routing.

In Chapter 2, we discuss the atomic unsplittable case of the classical model. We begin with the linear case, and review the results previously obtained by Awerbuch et al. [1]. We then present some novel results for unsplittable flow with polynomial cost functions:

- We present a new construction which has a price of anarchy higher than previous lower bound constructions.
- We prove an upper bound on the price of anarchy that matches our construction.
- We also consider the unweighted case, and again derive an upper bound and matching lower bound construction.

In Chapter 3, we define a new model which we have called the *flow-free* model. This model tries to capture a very simple observation: a car in traffic causes delay for vehicles behind it, but not in front of it. We define and discuss various properties of the model, and find exact values or bounds for the price of anarchy in various cases:

- In the nonatomic case, we obtain a tight result for polynomial cost functions (including linear cost functions).
- In the atomic case for linear cost functions, we obtain tight results for weighted and unweighted agents.
- In the atomic case for polynomial cost functions, we find an upper bound (although not a tight one).

- In the single-commodity case under certain conditions, we show that the price of anarchy is 1, in contrast to the classical model.

The conclusion deals principally with a discussion of opportunities for further work.

Chapter 1

Selfish Routing

1.1 The model

The model which we describe here has been around for a long time - it was discussed qualitatively by Pigou in 1920 [11], and much work was done in the 1950s by Wardrop [19] and others. The model, while fairly simple, captures a number of properties of road networks as described in the introduction.

We will represent the network as a directed graph G , with vertex set V and arc set E ; we will allow multiple arcs. There will be some set of pairs $\{(s_j, t_j) : 1 \leq j \leq n\}$ which are origin-destination pairs for some motorists (which we will henceforth refer to as users, agents, and later, players). Many users might be making the same journey, originating from the same origin and travelling to the same destination. So let the proportion of traffic corresponding to a particular pair (s_j, t_j) be w_j . Let us suppose that there is sufficient traffic that a single car, on its own, is almost insignificant. We can approximate this by regarding each car as being infinitesimally small; this of course means we must have an infinite number of cars. If the traffic flow really is large enough, this approximation will be fairly good. We will return to the case where the users have non-negligible size later.

We now need to model the congestion on an edge, by defining how long it takes for agents to traverse an edge. Here we will make a significant simplifying assumption: *all players traversing an edge experience the same delay*. It is by no means clear that this assumption is reasonable in all situations, and in fact this is a primary motivation for a new model which we will introduce in Chapter 3.

Given this assumption, it is very reasonable to require that the delay experienced on an edge depends on the amount of traffic using that edge, and nothing else. This can of course be different for different edges; a three lane highway will have very different characteristics from a narrow street. So for each edge $e \in E$, we define the *cost function* $f_e(x)$ which gives the delay on edge e as a function of the amount of traffic on that edge. We will require all cost functions to be non-negative, increasing and continuous. We will use the term “latency function” interchangeably with “cost function”.

Let \mathcal{R}_j be the set of $s_j - t_j$ paths, and let $\mathcal{R} = \cup_{j=1}^n \mathcal{R}_j$. For any path $P \in \mathcal{R}$, let x_P be the proportion of agents using path P . Then the total flow on an edge e is simply $x_e = \sum_{P \in \mathcal{R}: e \in P} x_P$. The delay experienced by users on edge e under this flow will then be simply $f_e(x_e)$. The average journey time over all users, which we will denote $C(\mathbf{x})$, is then

$$\begin{aligned} C(\mathbf{x}) &= \sum_{P \in \mathcal{R}} \sum_{e \in P} f_e(x_e) x_P \\ &= \sum_{e \in E} \sum_{P \in \mathcal{R}: e \in P} f_e(x_e) x_P \\ &= \sum_{e \in E} f_e(x_e) x_e. \end{aligned} \tag{1.1}$$

Here, \mathbf{x} is just the vector of all the x_e 's. This seems like a good candidate for a value measuring the overall “social” cost of the routing. Although we have so far measured in terms of the fraction of all agents, this is just a normalisation - setting $\sum_{j=1}^n w_j = 1$. This is not necessary, and we will often for example talk about, for




Figure 1.1: Pigou's example. The edge labels give the cost functions.

example, the “total” social cost; this will differ from the average journey time only by a factor of $\sum w_j$. Equation (1.1) then gives the total social cost; dividing through by $\sum w_j$ yields the average journey time.

1.2 The price of anarchy

Consider a very simple example. Our network consists only of two nodes, s and t , and all users need to travel from s to t . There are two roads available. The first is a very wide highway, which can accommodate a lot of traffic; however, it isn't quite a direct route, and so it takes an hour, no matter what the traffic situation. The second is a small road, with only a single lane; however, the route is much more direct. Let's assume for simplicity that the time taken on this road is simply proportional to the amount of traffic on it, and choose $f_{e_2}(x) = x$ for this edge. Figure 1.1 shows this example, commonly referred to as Pigou's example.

Assuming all the players are acting selfishly, what flow will the system settle down to? Well, if the fraction of players using the narrow road is less than one, that route will be quicker than the highway; thus any players using the highway will switch. On the other hand, if all agents take the narrow road, nobody will have an incentive to switch, since both routes take exactly an hour. In the language of game theory,

this flow is a *Nash equilibrium*. In fact, it is a pure strategy Nash equilibrium (often abbreviated PSNE), since each player picks a specific strategy (the narrow road). It can be shown [2] that a PSNE always exists in this game.

So the Nash equilibrium is obtained when all the players are taking the narrow road, in which case everybody takes an hour to get to work.

Now suppose we are able to force people to take a specific route. What's the most efficient solution? Suppose α is the fraction of traffic that takes the narrow road, so a $1 - \alpha$ fraction takes the highway. The average commute time is then

$$\alpha f_{e_2}(\alpha) + (1 - \alpha)f_{e_1}(1 - \alpha) = \alpha^2 + 1 - \alpha.$$

This is minimised when $\alpha = \frac{1}{2}$, in which case the average commute time is $\frac{3}{4}$ hours.

So for this specific instance of the network game, the ratio between the average commute time in the Nash equilibrium is $\frac{4}{3}$ the average of the optimal solution. We will call this ratio the *price of anarchy*. The concept was first defined in [9], where it was dubbed the coordination ratio; the name “price of anarchy” was coined in [10].

Actually, there may be more than one Nash equilibrium; we define the price of anarchy to be the ratio of the *worst* (highest cost) Nash over the optimum cost.¹

The price of anarchy is $\frac{4}{3}$ for Pigou's very simple example. What about different networks, and different cost functions? How large can it get? Well, if we don't restrict any further the choice of cost functions, it can be arbitrarily large. To see this, consider again Pigou's example, but replace the cost function for the narrow road with $f_{e_2}(x) = x^d$, where d is some integer. The Nash equilibrium remains the same, with all agents taking the narrow road. It can easily be checked that the optimal flow \mathbf{x}^* is obtained by routing a fraction $\alpha = (d + 1)^{-1/d}$ along the lower link, and yields a cost of

$$C(\mathbf{x}^*) = 1 - d(d + 1)^{-1-1/d}.$$

¹Actually, in the network game as defined so far with infinitesimally small agents, all Nash have the same cost [2].

So the price of anarchy is $(1 - d(d+1)^{-1-1/d})^{-1}$, which tends to infinity as $d \rightarrow \infty$. For this reason, we will concentrate on specific classes of cost functions. In this thesis, most of our attention will be on linear cost functions, and polynomial cost functions where some maximum degree is prescribed. We define the price of anarchy of a set of functions (such as linear functions) to be the supremum of the price of anarchy over all network games with cost functions in that set.

1.3 Bounding the price of anarchy

In the original paper [16] of Roughgarden and Tardos, the price of anarchy was derived in the nonatomic case for linear cost functions, and in [12] Roughgarden did the same for polynomial cost functions. The proofs however are somewhat involved. Correa, Schulz and Stier-Moses [5] give a much shorter proof. A variation of these ideas will be used to prove some of the novel results of this thesis later, so we demonstrate their proof here.

Theorem 1.1. *The price of anarchy with linear cost functions and nonatomic agents is $\frac{4}{3}$.*

Proof. The starting point of the proof is the variational form of the requirement for a Nash equilibrium: A flow \mathbf{x} is a Nash equilibrium iff

$$\sum_{e \in E} f_e(x_e)(x_e - x'_e) \leq 0 \quad \text{for all valid flows } \mathbf{x}'. \quad (1.2)$$

This well known result can be found in [18]; we will not rederive it here. So let x^* be




Figure 1.2: A visual proof of Equation 1.3

the optimal flow; we have by the above

$$\begin{aligned}
 C(\mathbf{x}) &\leq \sum_{e \in E} f_e(x_e) x_e^* \\
 &= \sum_{e \in E} f_e(x_e^*) x_e^* + \sum_{e \in E} (f_e(x_e) - f_e(x_e^*)) x_e^* \\
 &\leq C(\mathbf{x}^*) + \sum_{e \in E: x_e^* < x_e} (f_e(x_e) - f_e(x_e^*)) x_e^*,
 \end{aligned}$$

since f_e is increasing. Now consider Figure 1.2. The area of the greyed rectangle is equal to $(f_e(x_e) - f_e(x_e^*))x_e^*$; this is clearly at most $\frac{1}{2}$ of the area of the upper triangle, which is in turn at most $\frac{1}{2}$ the area of the large rectangle. This has area $f_e(x_e)x_e$. So

$$C(\mathbf{x}) \leq C(\mathbf{x}^*) + \sum_{e \in E: x_e^* < x_e} \frac{1}{4} f_e(x_e) x_e \quad (1.3)$$

$$\leq C(\mathbf{x}^*) + \frac{1}{4} C(\mathbf{x}), \quad (1.4)$$

and so

$$\frac{C(\mathbf{x})}{C(\mathbf{x}^*)} \leq \frac{4}{3}, \quad (1.5)$$

as required. □

So we see that the very simple example of Pigou actually yields the maximum possible price of anarchy when only linear cost functions are allowed, *for any network topology*.

The same method applies to polynomial latency functions with non-negative coefficients and maximum degree d ; some calculus is required to calculate the maximum ratio of $f_e(x_e)x_e$ to $(f_e(x_e) - f_e(x_e^*))x_e^*$. The price of anarchy in this case is

$$\rho = (1 - d(d+1)^{-1-1/d})^{-1}, \quad (1.6)$$

again the same as the two-link example with cost functions a and x^d on the two links.

1.4 Atomic games

Up until now, we have assumed that an individual agent is infinitesimally small. What if this is not true? There are two fairly natural ways to modify the model to handle this case, although only the first will apply to the example of traffic flow.

Unsplittable flow

The first variant is quite straightforward. Instead of an infinite number of infinitesimal agents, we have a finite set J of agents (let $n = |J|$). Each agent is associated with a source s_j and sink t_j as before. Now however, each agent has a specified non-zero size w_j . A feasible routing is then defined as $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$, where each P_j is a path from s_j to t_j .

In this model, the agents are *unsplittable*; they take only one route from their source to their destination. This makes sense for traffic modelling, where each agent corresponds to a car. The different sizes w_j might be used to model trucks and buses, or all the w_j 's might be set to 1.

The flow vector \mathbf{x} is easily determined from \mathcal{P} :

$$x_e = \sum_{j:e \in P_j} w_j.$$

We will define the social cost as in the nonatomic model, by Equation (1.1):

$$\begin{aligned} C(\mathcal{P}) &= \sum_{e \in E} f_e(x_e) x_e \\ &= \sum_{e \in E} f_e(x_e) \sum_{j:e \in P_j} w_j \\ &= \sum_{j \in J} \left(\sum_{e \in P_j} f_e(x_e) \right) w_j \\ &= \sum_{j \in J} \ell_j(\mathcal{P}) w_j, \end{aligned} \tag{1.7}$$

where $\ell_j(\mathcal{P})$ is the latency experienced by player j under routing \mathcal{P} . Note that this means that players' contribution to the total cost is proportional to their weight. This is not the only possible option - one could define the social cost to instead be

$$C'(\mathcal{P}) = \sum_{j \in J} \ell_j(\mathcal{P}). \tag{1.8}$$

Which is appropriate depends on exactly what we are modelling and what the agents represent. For example, if agents are cars, and larger w_j 's represent larger vehicles, we probably don't want to give extra consideration to larger vehicles in our cost, so C' would be the appropriate cost. But suppose each agent represents a large bundle of mail, that must be transported as a unit, and the w_j 's represent the number of letters in the bundle. In that case, it makes sense to consider the social cost to be the average time taken per letter, not per bundle, and C is appropriate. We will only consider C as defined in (1.7) for this thesis, following Awerbuch et al. [1]. We will spend a lot of time discussing the nonatomic unsplittable case in Chapters 2 and 3.

Splittable flow

In contrast, the agents in this second atomic variation are allowed to split their flow. This doesn't make much sense in the context of road networks (at least not if we associate each agent with an individual motorist), but it has applications in other areas. Imagine a number of users downloading some large data file over the internet (movies for example). The data is not transferred as a single chunk; rather it is split into packets, each of which can be routed independently through the network. Each user wants to minimise the time required to download the entire file; they are not interested in individual packets. This could have an effect on the optimum routing; it may be better to route some packets along a longer route to reduce congestion for the remaining packets.

The behaviour in this case is somewhat counterintuitive, as can be attested by a number of incorrect results in the literature. Intuitively, it might seem that the price of anarchy in this case should be no larger than in the atomic case; a user is at least trying to enforce the “social optimum” for the flow under her control. It seems plausible that this should push things in the direction of the global social optimum. In fact, Roughgarden [15] and Correa et al. [5] both independently published proofs that the price of anarchy in the atomic splittable case could not exceed that of the nonatomic case. As was demonstrated by Cominetti et al. [4] however, the proofs are incorrect. They give the following example with linear cost functions where the price of anarchy exceeds $4/3$, the upper bound in the nonatomic case.

Consider Figure 1.3. There are infinitely many nonatomic users with total weight 1 routing that want to route from s_1 to t_1 , and there is one player, controlling 1 unit of flow, routing from s_2 to t_2 . (We have a mixture of atomic and nonatomic users in this example; if desired, replacing the nonatomic players with a sufficiently large, but finite, number of atomic players will still provide an example exceeding $4/3$). It can be verified that the routing where all the flow of the nonatomic players, and




Figure 1.3: An atomic splittable game with linear cost functions that has a price of anarchy larger than $4/3$.

0.9 of the flow from the atomic player, is routed through the central edge is a Nash. This has cost 3.89. On the other hand, the optimum flow is obtained by routing all of the nonatomic players along the left edge, and all of the flow from the atomic player through the centre. This has a total cost of 2.9, yielding a price of anarchy of approximately 1.341, which is slightly larger than $4/3$.

1.5 Congestion games

The games we have considered up until this point have all been tied to a network structure. Is there a way to generalise these models to versions that don't have this underlying structure? There is, and the generalised games are called congestion games. We will define the congestion game equivalent of the atomic unsplittable network game; it will be clear how to generalise other variants.

So again we have a set J of players, each with weight w_j . We also have a set of items I - these should be considered the equivalent of the edges in the network model. A cost function $f_i(x)$ is associated with each item $i \in I$, exactly as in the

network game. 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 players allowed strategies, or how many strategies a player may have. Now 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. These games are normally referred to as *weighted congestion games* in the literature, since congestion games were first considered in the unweighted case where $w_j = 1$ for all j .

Essentially all of the results we obtain in this thesis will apply to both network games and the more general congestion game equivalents. We will generally prove results in terms of network games, but it will be clear that the network structure is not used.

Chapter 2

Unsplittable Flow

In this chapter, we consider in detail the atomic unsplittable model defined in the previous chapter. After deriving a useful inequality that holds for all Nash flows, we will review the linear case, where tight results were already available. The bulk of the chapter will then be devoted to demonstrating a tight upper bound in the polynomial case, which is a new result.

We will only consider pure Nash equilibria in this chapter. If the cost functions are linear, a pure Nash always exists, as was demonstrated using a potential function approach [6]. Allowing for more general cost functions, a pure Nash need not exist; in Goemans et al. [7] a construction with no pure Nash is given using quadratic cost functions. In [1], both pure and mixed strategy Nash equilibria are considered. We will only consider pure Nash equilibria; most likely, the results for the polynomial case could be generalised to mixed equilibria without too much difficulty.

2.1 The Nash condition

Suppose \mathcal{P} is a Nash flow. This means that any agent $j \in J$ has no incentive to switch, since the latency along any $s_j - t_j$ path is no smaller than the latency along

P_j , the path taken at Nash. If player j tried to switch to some other path P' , his latency would be

$$\ell'_j = \sum_{e \in P_j \cap P'} f_e(x_e) + \sum_{e \in P' \setminus P_j} f_e(x_e + w_j),$$

since he would increase the flow along any new edges that weren't already part of P_j . Applying this with $P' = P_j^*$, the path of agent j at some optimal solution \mathcal{P}^* , we obtain

$$\begin{aligned} \ell_j(\mathcal{P}) &\leq \sum_{e \in P_j \cap P_j^*} f_e(x_e) + \sum_{e \in P_j^* \setminus P_j} f_e(x_e + w_j) \\ &\leq \sum_{e \in P_j^*} f_e(x_e + w_j). \end{aligned}$$

Now using (1.7),

$$\begin{aligned} C(\mathcal{P}) &= \sum_{j \in J} \ell_j(\mathcal{P}) w_j \\ &\leq \sum_{j \in J} \sum_{e \in P_j^*} f_e(x_e + w_j) w_j. \end{aligned} \tag{2.1}$$

We will sometimes use the notation $C_j(\mathcal{P})$ for the portion of the total cost attributable to player j ; this is equal to $\ell_j(\mathcal{P}) w_j$.

2.2 Linear cost functions

A tight result in the case of linear cost functions was obtained by Awerbuch et al. [1], which we briefly review here.

Theorem 2.1 (Awerbuch et al. [1]). *The price of anarchy of an atomic unsplittable network game with linear cost functions is no more than $(3 + \sqrt{5})/2$.*

Proof. Let $f_e(x) = a_e x + b_e$.

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

Now since $x_e^* = \sum_{j: e \in P_j^*} w_j$, and so also $x_e^{*2} \geq \sum_{j: e \in P_j^*} w_j^2$,

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

The second term is just $C(\mathcal{P}^*)$; apply Cauchy-Schwartz to the first:

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

Write $\alpha = \sqrt{C(\mathcal{P})/C(\mathcal{P}^*)}$; then dividing the above by $C(\mathcal{P}^*)$,

$$\alpha^2 \leq \alpha + 1.$$

Thus $\alpha \leq (\sqrt{5} + 1)/2$, and so finally

$$\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} = \alpha^2 \leq \frac{3 + \sqrt{5}}{2}.$$

□

This upper bound is tight, as is confirmed by the construction shown in Figure 2.1. There are four players; the four (s_j, t_j, w_j) triplets are, in order, (u, v, ϕ) , (u, w, ϕ) ,




Figure 2.1: A network game with the largest possible price of anarchy when restricted to linear cost functions.

$(v, w, 1)$ and $(w, v, 1)$. Here, ϕ is the golden ratio, $(1 + \sqrt{5})/2$. It can easily be checked that the strategy profile defined by the paths

$$P_1 = uwv, \quad P_2 = uvw, \quad P_3 = vuw, \quad P_4 = wuv$$

is a Nash equilibrium, and has cost $4\phi^2 + 4\phi + 2$. The optimum solution is

$$P_1^* = uv, \quad P_2^* = uw, \quad P_3^* = vw, \quad P_4^* = wv,$$

which has cost $2\phi^2 + 2$. The price of anarchy of this game is thus

$$\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} = \frac{4\phi^2 + 4\phi + 2}{2\phi^2 + 2} = \frac{\phi^2 + 3\phi + 2}{\phi + 2} = \phi + 1 = \frac{3 + \sqrt{5}}{2}.$$

2.3 Improved bounds for the unsplittable case

In this section, we consider the unsplittable atomic case with polynomial latency functions of maximum degree d . Actually, we consider cost functions in \mathcal{C}_d , defined as the set of polynomials of maximum degree d with non-negative coefficients. The requirement that the coefficients be non-negative will be assumed from here on.

Awerbuch et al. [1] showed that the price of anarchy is $\Omega(d^{d/2})$ and $\mathcal{O}(2^d d^{d+1})$. The lower bound by Christodoulou and Koutsoupias [3] of $\Omega(d^{d(1-o(1))})$ for finite congestion

games can also be modified to provide a lower bound for unsplittable network flow. Let $\varphi(d)$ be the positive real root of the equation $(x+1)^d = x^{d+1}$, so that

$$(\varphi(d) + 1)^d = \varphi(d)^{d+1} \quad (2.2)$$

is satisfied. We will write simply φ if there is no confusion over the value of d . We will show that φ^{d+1} is a tight upper bound for the value of the price of anarchy, and also give constructions that obtain this upper bound. It is interesting to compare with the nonatomic case, where the asymptotic behaviour of Equation (1.6) is easily found to be

$$\rho_{\text{nonatomic}} = \Theta\left(\frac{d}{\ln d}\right). \quad (2.3)$$

Although computing the exact asymptotic behaviour of φ^{d+1} seems to be somewhat problematic, it can easily be shown that for any $\epsilon > 0$,

$$\frac{d}{\ln d} < \varphi < (1 + \epsilon) \frac{d}{\ln d}$$

for d sufficiently large. It follows that the price of anarchy in the unsplittable case satisfies

$$\rho = o\left((1 + \epsilon)^{d+1} (d/\ln d)^{d+1}\right) \quad \text{and} \quad \rho = \omega\left((d/\ln d)^{d+1}\right).$$

So, roughly speaking, ρ is about the $(d+1)$ 'th power of $\rho_{\text{nonatomic}}$ - a considerable difference.

We will require the following identities that follow from the definition of φ :

$$\varphi + 1 = \varphi^{1+1/d} \quad (2.4)$$

$$1 + \varphi^{-1} = \varphi^{1/d} \quad (2.5)$$

$$\varphi^{-1/d} + \varphi^{-1-1/d} = 1 \quad (2.6)$$

2.3.1 The lowerbound construction

We construct a congestion game with price of anarchy $\varphi(d)^{d+1}$.

Let the set of items be $I = I_0 \cup \bar{I}_0$, where $I_0 = \{0, 1, \dots, d\}$, and $\bar{I}_0 = \{\bar{0}, \bar{1}, \dots, \bar{d}\}$ is a disjoint copy of I_0 . Let the set of players be $J = J_0 \cup \bar{J}_0$, where $J_0 = \{0, 1, \dots, d\}$ and $\bar{J}_0 = \{\bar{0}, \bar{1}, \dots, \bar{d}\}$. We also define a bar operation in the obvious way, so $\bar{\bar{0}} = 0$, etc. and $\overline{\{A\}} = \{\bar{A}\}$.

The player weights w_j are defined by

$$\begin{aligned} w_j &= w_{\bar{j}} = \varphi^{-(j+1)/d} & \forall j \in \{0, 1, \dots, d-1\} \\ w_d &= w_{\bar{d}} = 1. \end{aligned}$$

The cost functions f_i are defined as $f_i(x) = a_i x^d$, where

$$\begin{aligned} a_0 &= a_{\bar{0}} = 1 \\ a_i &= a_{\bar{i}} = (\varphi + 1)^{i-1} & \forall i \in \{1, 2, 3, \dots, d\}. \end{aligned}$$

The set of allowed strategies for player $j \in J$ is $\mathcal{S}_j = \{S_j, S_j^*\}$ where

$$\begin{aligned} S_j &= \{0, 1, \dots, j\} & \forall j \in J_0 \\ S_{\bar{j}} &= \{\bar{0}, \bar{1}, \dots, \bar{j}\} & \forall \bar{j} \in \bar{J}_0 \\ S_j^* &= \{\overline{j+1}\} & \forall j \in J_0 \setminus \{d\} \\ S_{\bar{j}}^* &= \{j+1\} & \forall \bar{j} \in \bar{J}_0 \setminus \{\bar{d}\} \\ S_d^* &= \{\bar{0}\} \\ S_{\bar{d}}^* &= \{0\}. \end{aligned}$$

Let $\mathcal{P} = \{S_j : j \in J\}$ and $\mathcal{P}^* = \{S_j^* : j \in J\}$. Once our construction is complete, we will show that \mathcal{P} is a Nash equilibrium; \mathcal{P}^* will be the optimal solution.

Let x_i and x_i^* be the total utilisation of item i under \mathcal{P} and \mathcal{P}^* respectively. We

first calculate x_i for $i \in I_0$:

$$\begin{aligned}
 x_i &= \sum_{k=i}^{d-1} w_k + w_d \\
 &= \sum_{k=i+1}^d \varphi^{-k/d} + 1 \\
 &= \varphi^{-(i+1)/d} \left(\frac{1 - \varphi^{-(d-i)/d}}{1 - \varphi^{-1/d}} \right) + 1 \\
 &= \varphi^{-(i+1)/d} \left(\frac{1 - \varphi^{-1+i/d}}{\varphi^{-1-1/d}} \right) + 1 \quad \text{from (2.6)} \\
 &= \varphi^{1-i/d}
 \end{aligned} \tag{2.7}$$

Thus for $j \in J_0$,

$$\begin{aligned}
 f_j(x_j) &= f_j(x_j) = a_j(\varphi^{1-j/d})^d \\
 &= a_j \varphi^{d-j} \\
 &= \begin{cases} \varphi^d & j = 0 \\ (\varphi + 1)^{j-1} \varphi^{d-j} & j \geq 1, \end{cases}
 \end{aligned}$$

and hence

$$\begin{aligned}
 C_j(\mathcal{P}) &= w_j \sum_{k=0}^j f_k(x_k) \\
 &= w_j \left[\sum_{k=1}^j \varphi^{d-1+(k-1)/d} + \varphi^d \right] \\
 &= w_j \left[\varphi^{d-1} \left(\frac{\varphi^{j/d} - 1}{\varphi^{1/d} - 1} \right) + \varphi^d \right] \\
 &= w_j \left[\varphi^{d-1} \left(\frac{\varphi^{j/d} - 1}{\varphi^{-1}} \right) + \varphi^d \right] \quad \text{using (2.5)} \\
 &= w_j \varphi^{d+j/d}.
 \end{aligned}$$

Define $\mathcal{P}^{(j)} = \mathcal{P} \setminus S_j \cup S_j^*$. For $j \in \{0, 1, \dots, d-1\}$:

$$\begin{aligned}
 C_j(\mathcal{P}^{(j)}) &= w_j f_{j+1}(x_{j+1} + w_j) \\
 &= w_j f_{j+1}(x_j) \\
 &= w_j (\varphi + 1)^j \varphi^{d-j} \\
 &= w_j \varphi^{j(d+1)/d} \varphi^{d-j} \quad \text{by (2.2)} \\
 &= w_j \varphi^{d+j/d}.
 \end{aligned}$$

Also,

$$\begin{aligned}
 C_d(\mathcal{P}^{(d)}) &= w_d f_0(x_0 + w_d) \\
 &= w_d (\varphi + 1)^d \\
 &= w_d \varphi^{d+1} \quad \text{by (2.2)}.
 \end{aligned}$$

Thus $C_j(\mathcal{P}) = C_j(\mathcal{P}^{(j)})$ for all $j \in J_0$, and so (by symmetry) for all $j \in J$. So at \mathcal{P} , no player has an incentive to switch to S_j^* , and so it is a Nash.

Finally, notice that for all $i \in \{1, 2, \dots, d\}$,

$$\frac{x_i}{x_i^*} = \frac{x_i}{w_{i-1}} = \frac{\varphi^{1-i/d}}{\varphi^{-i/d}} = \varphi.$$

Additionally,

$$\frac{x_0}{x_0^*} = \frac{\varphi}{1} = \varphi.$$

So (by symmetry) $x_i/x_i^* = \varphi$ for all $i \in I$, and so the price of anarchy is

$$\rho = \frac{\sum_{i \in I} a_i x_i^{d+1}}{\sum_{i \in I} a_i x_i^{*d+1}} = \frac{\sum_{i \in I} a_i \varphi^{d+1} x_i^{*d+1}}{\sum_{i \in I} a_i x_i^{*d+1}} = \varphi^{d+1}.$$

The construction can easily be modified to give a network game. Figure 2.2 shows a possible construction for $d = 2$. Arcs labelled $1, \bar{1}$ etc refer to the items of the original congestion game, and have the same cost functions; dashed arcs have zero cost. Each player $j \in J$ requires a flow of w_j from s_j to t_j . It is easy to check that




Figure 2.2: A network game version of the lower bound construction for $d = 2$.

the $s_j - t_j$ paths available to player j correspond exactly to the allowed strategies in the congestion game construction. It is also easy to see how this can be generalised to arbitrary d .

2.3.2 A matching upper bound

The proof of the matching upper bound requires two steps; we first show that the price of anarchy does not decrease if we restrict the cost functions to be in the set

$$\tilde{\mathcal{C}}_d = \{ax^d : a \geq 0\};$$

we then use a sequence of inequalities partially based on the upper bound proof in [1].

Theorem 2.2. *Given an arbitrary weighted finite congestion game G with cost functions in \mathcal{C}_d which has a pure Nash equilibrium, there exists another congestion game \hat{G} with cost functions in $\tilde{\mathcal{C}}_d$ where the price of anarchy is at least as large.*

Proof. Denote the set of items by E , the players by J , and the set of strategies available to player $j \in J$ by \mathcal{S}_j . As usual, denote the cost functions by $f_e(x)$ and the player weights by w_j .

Let the price of anarchy of G be ρ . Pick a Nash equilibrium of G with maximal cost, i.e. ρ times the cost of the optimal solution. Let $S_j \in \mathcal{S}_j$ be the strategy that player j plays in this Nash equilibrium, and let S_j^* be the strategy player j plays in the optimal solution. Note that any other strategies in G are superfluous—discarding them does not affect the price of anarchy. As usual, let $\mathcal{P} = \{S_j : j \in J\}$, $\mathcal{P}^* = \{S_j^* : j \in J\}$ and $\mathcal{P}^{(j)} = \mathcal{P} \setminus S_j \cup S_j^*$. By Equation (3.2), the Nash requirement yields $C_j(\mathcal{P}) \leq C_j(\mathcal{P}^{(j)})$.

We now define a new game \hat{G} . Define

$$F = \{e \in E : x_e < x_e^*\},$$

$$F_j^s = \{e \in F : x_e^{(j)} \leq x_e^*\}$$

and

$$F_j^l = F \setminus F_j^s.$$

Intuitively, at the Nash flow, player j makes a “small” increase (or even a decrease) to elements of F_j^s upon switching to S_j^* , but a “large” increase to elements of F_j^l (which also implies that $F_j^l \subseteq S_j^*$). Also let

$$F' = \{e' : e \in F\}$$

be a disjoint copy of F ; for any $U \subseteq F$, we will use U' to denote $\{e' \in F' : e \in U\}$. Also, define a new item t_j for every $j \in J$. The item set for \hat{G} will be

$$\hat{E} = E \cup F' \cup \{t_j : j \in J\}.$$

The cost functions $\hat{f}_e(x)$ are defined as follows:

$$\hat{f}_e(x) = f_e(x_e) \left(\frac{x}{x_e} \right)^d \quad \forall e \in E \setminus F, \quad (2.8)$$

$$\hat{f}_e(x) = f_e(x_e^*) \left(\frac{x}{x_e^*} \right)^d \quad \forall e \in F, \quad (2.9)$$

$$\hat{f}_{e'}(x) = f_e(x_e) \left(\frac{x}{x_e} \right)^d - \hat{f}_e(x) \quad \forall e' \in F', \quad (2.10)$$

$$\hat{f}_{t_j} = \sum_{e \in F_j^s} f_e(x_e^{(j)}) \left(\frac{x}{w_j} \right)^d \quad \forall j \in J. \quad (2.11)$$

We should verify that $\hat{f}_{e'} \in \tilde{\mathcal{C}}_d$; to see this, we first define

$$h_e(x) = \frac{f_e(x)}{x^d} = \sum_{i=0}^d a_{e,i} x^{i-d}.$$

$h_e(x)$ is clearly a non-increasing function. Now for any $e' \in F'$,

$$\begin{aligned} \hat{f}_{e'}(x) &= \left(\frac{f_e(x_e)}{x_e^d} - \frac{f_e(x_e^*)}{x_e^{*d}} \right) x^d \\ &= (h_e(x_e) - h_e(x_e^*)) x^d \\ &\in \tilde{\mathcal{C}}_d, \end{aligned}$$

since $x_e \leq x_e^*$ for $e \in F$ and so $h_e(x_e) \geq h_e(x_e^*)$.

Define the strategy set \hat{S}_j for player j in \hat{G} as follows: $\hat{S}_j = \{\hat{S}_j, \hat{S}_j^*\}$, where

$$\hat{S}_j = S_j \cup (F \cap S_j)',$$

$$\hat{S}_j^* = S_j^* \setminus F_j^s \cup \{t_j\}.$$

We now claim that $\hat{\mathcal{P}} = \{\hat{S}_j : j \in J\}$ is a Nash for the new game. We use the notation \hat{C} to refer to costs in \hat{G} , and \hat{x} for the Nash flow vector.

$$\begin{aligned} \hat{C}_j(\hat{\mathcal{P}}) &= w_j \sum_{e \in \hat{S}_j} \hat{f}_e(\hat{x}_e) \\ &= w_j \left[\sum_{e \in S_j \setminus F} \hat{f}_e(x_e) + \sum_{e \in S_j \cap F} \hat{f}_e(x_e) + \sum_{e' \in (S_j \cap F)'} \hat{f}_{e'}(x_e) \right]. \end{aligned}$$

Combine the last two terms and apply (2.10):

$$\begin{aligned}
\hat{C}_j(\hat{\mathcal{P}}) &= w_j \left[\sum_{e \in S_j \setminus F} f_e(x_e) \left(\frac{x_e}{x_e} \right)^d + \sum_{e \in S_j \cap F} f_e(x_e) \left(\frac{x_e}{x_e} \right)^d \right] \\
&= w_j \sum_{e \in S_j} f_e(x_e) \\
&= C_j(\mathcal{P}).
\end{aligned}$$

Define $\hat{\mathcal{P}}^{(j)} = \hat{\mathcal{P}} \setminus \hat{S}_j \cup \hat{S}_j^*$ analogously to $\mathcal{P}^{(j)}$; the strategy profile where player j plays the optimal strategy, and all other players the Nash strategy.

$$\begin{aligned}
\hat{C}_j(\hat{\mathcal{P}}^{(j)}) &= w_j \sum_{e \in \hat{S}_j^*} \hat{f}_e(\hat{x}_e^{(j)}) \\
&= w_j \left[\sum_{e \in S_j^* \setminus F_j^s} \hat{f}_e(x_e^{(j)}) + \hat{f}_{t_j}(w_j) \right] \\
&= w_j \left[\sum_{e \in S_j^* \setminus F_j^s} \hat{f}_e(x_e^{(j)}) + \sum_{e \in F_j^s} f_e(x_e^{(j)}) \left(\frac{w_j}{w_j} \right)^d \right] \\
&= w_j \left[\sum_{e \in S_j^* \setminus F} f_e(x_e) \left(\frac{x_e^{(j)}}{x_e} \right)^d + \sum_{e \in F_j^l} f_e(x_e^*) \left(\frac{x_e^{(j)}}{x_e^*} \right)^d + \sum_{e \in F_j^s} f_e(x_e^{(j)}) \right],
\end{aligned}$$

by the definition of \hat{f}_e and since $S_j^* \setminus F_j^s = (S_j^* \setminus F) \cup F_j^l$. Now since $x_e^{(j)} \geq x_e$ for all $e \in S_j^*$ (and hence $S_j^* \setminus F$), and $x_e^{(j)} \geq x_e^*$ for all $e \in F_j^l$, and $h_e(x)$ is decreasing,

$$\begin{aligned}
\hat{C}_j(\hat{\mathcal{P}}^{(j)}) &\geq w_j \left[\sum_{e \in S_j^* \setminus F} f_e(x_e^{(j)}) + \sum_{e \in F_j^l} f_e(x_e^{(j)}) + \sum_{e \in F_j^s} f_e(x_e^{(j)}) \right] \\
&= C_j(\mathcal{P}^{(j)}).
\end{aligned}$$

Since $C_j(\mathcal{P}^{(j)}) \geq C_j(\mathcal{P})$ in G by the Nash requirement, $\hat{C}_j(\hat{\mathcal{P}}^{(j)}) \geq \hat{C}_j(\hat{\mathcal{P}})$ as required.

So $\hat{\mathcal{P}}$ is a Nash.

Since $\hat{C}(\hat{\mathcal{P}}) = \hat{C}(\mathcal{P})$, to show that the price of anarchy of \hat{G} is no smaller than ρ we must show that $\hat{C}(\hat{\mathcal{P}}^*) \leq C(\mathcal{P}^*)$.

$$\begin{aligned}
\hat{C}_j(\hat{\mathcal{P}}^*) &= w_j \sum_{e \in \hat{S}_j^*} \hat{f}_e(\hat{x}_e^*) \\
&= w_j \left[\sum_{e \in S_j^* \setminus F_j^s} \hat{f}_e(\hat{x}_e^*) + \hat{f}_{t_j}(w_j) \right] \\
&= w_j \left[\sum_{e \in S_j^* \setminus F} f_e(x_e) \left(\frac{x_e^*}{x_e} \right)^d + \sum_{e \in F_j^l} f_e(x_e^*) + \sum_{e \in F_j^s} f_e(x_e^{(j)}) \right].
\end{aligned}$$

Now $x_e \geq x_e^*$ for $e \in S_j^* \setminus F$ and $x_e^* \geq x_e^{(j)}$ for $e \in F_j^s$, and so since $h_e(x)$ is decreasing we have

$$\begin{aligned}
C_j(\hat{\mathcal{P}}^*) &\leq w_j \left[\sum_{e \in S_j^* \setminus F} f_e(x_e^*) + \sum_{e \in F_j^l} f_e(x_e^*) + \sum_{e \in F_j^s} f_e(x_e^*) \right] \\
&= w_j \sum_{e \in S_j^*} f_e(x_e^*) \\
&= C_j(\mathcal{P}^*).
\end{aligned}$$

Thus we have achieved the required reduction. \square

This means an upper bound on the price of anarchy with cost functions in $\tilde{\mathcal{C}}_d$ will apply to general polynomial cost functions in \mathcal{C}_d , so we restrict our attention to cost functions in $\tilde{\mathcal{C}}_d$ from now on. We begin with a useful lemma:

Lemma 2.1. *For $a, b \geq 0$, $d \geq 1$ and $0 < \gamma < 1$,*

$$(a + b)^d \leq \gamma^{1-d} a^d + (1 - \gamma)^{1-d} b^d. \quad (2.12)$$

Proof.

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

□

Theorem 2.3. *For any weighted congestion game G with cost functions in $\tilde{\mathcal{C}}_d$, the price of anarchy is at most φ^{d+1} .*

Proof. We begin with Equation (2.1):

$$C(\mathcal{P}) \leq \sum_{e \in E} \sum_{j: e \in P_j^*} f_e(x_e + w_j) w_j \quad (2.13)$$

$$= \sum_{e \in E} \sum_{j: e \in P_j^*} a_e (x_e + w_j)^d w_j. \quad (2.14)$$

We now apply Lemma 2.1, with $a = x_e$ and $b = w_j$, and γ to be determined later:

$$\begin{aligned}
C(\mathcal{P}) &\leq \sum_{e \in E} \sum_{j: e \in P_j^*} [a_e \gamma^{1-d} x_e^d w_j + a_e (1-\gamma)^{1-d} w_j^{d+1}] \\
&\leq \gamma^{1-d} \sum_{e \in E} a_e x_e^d x_e^* + (1-\gamma)^{1-d} \sum_{e \in E} a_e x_e^{*d+1}.
\end{aligned}$$

We now apply Hölder's inequality,

$$\sum_e u_e^\alpha v_e^\beta \leq \left(\sum_e u_e \right)^\alpha \left(\sum_e v_e \right)^\beta, \quad \alpha + \beta = 1$$

to the first term, with $\alpha = d/(d+1)$, $\beta = 1/(d+1)$, $u_e = a_e x_e^{d+1}$ and $v_e = a_e x_e^{*d+1}$ to obtain

$$\begin{aligned}
C(\mathcal{P}) &\leq \gamma^{1-d} \left(\sum_{e \in E} a_e x_e^{d+1} \right)^{d/(d+1)} \left(\sum_{e \in E} a_e x_e^{*d+1} \right)^{1/(d+1)} + (1-\gamma)^{1-d} C(\mathcal{P}^*) \\
&= \gamma^{1-d} C(\mathcal{P})^{d/(d+1)} C(\mathcal{P}^*)^{1/(d+1)} + (1-\gamma)^{1-d} C(\mathcal{P}^*). \quad (2.15)
\end{aligned}$$

Now let $z = \left(\frac{C(\mathcal{P})}{C(\mathcal{P}^*)}\right)^{1/(d+1)}$ and divide by $C(\mathcal{P}^*)$ to obtain

$$z^{d+1} \leq \gamma^{1-d} z^d + (1 - \gamma)^{1-d};$$

let $\zeta(z) = z^{d+1} - \gamma^{1-d} z^d - (1 - \gamma)^{1-d}$.

We now choose $\gamma = \frac{\varphi}{\varphi+1}$. Then

$$\begin{aligned} \zeta(\varphi) &= \varphi^{d+1} - \left(\frac{\varphi+1}{\varphi}\right)^{d-1} \varphi^d - (\varphi+1)^{d-1} \\ &= \varphi^{d+1} - (\varphi+1)^{d-1} \varphi - (\varphi+1)^{d-1} \\ &= \varphi^{d+1} - (\varphi+1)^d \\ &= 0 \end{aligned}$$

by the definition of φ , and so φ is a root of ζ . Also,

$$\begin{aligned} \zeta'(z) &= (d+1)z^d - \left(\frac{\varphi+1}{\varphi}\right)^{d-1} dz^{d-1} \\ &= z^{d-1} \left((d+1)z - \frac{\varphi^2}{\varphi+1} d \right) \\ &> dz^{d-1} (z - \varphi) \\ &\geq 0 \quad \text{for } z \geq \varphi. \end{aligned}$$

Thus φ is the largest root, and in order to satisfy (2.15) we must have

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

□

Combining the two theorems, we obtain

Corollary 2.1. *Any weighted congestion game with polynomial cost functions has a price of anarchy of at most φ^{d+1} .*

□

2.4 Unweighted games

We now consider the case where we restrict $w_j = 1$ for all players j . Our bound is more precise than the $d^{\Theta(d)}$ bound of Christodoulou and Koutsoupias [3], and is again tight. The construction is actually quite different from the weighted version, which is perhaps surprising.

By writing (2.2) in the form $(1 + \varphi^{-1})^d = \varphi$, it can easily be shown that φ is never an integer. Let $k = \lfloor \varphi \rfloor$. Then define α, β by

$$\alpha = (k + 1)^d - k^{d+1} \tag{2.16}$$

$$\beta = (k + 1)^{d+1} - (k + 2)^d. \tag{2.17}$$

Since $x^{-d} \cdot (x^{d+1} - (x + 1)^d) = x - (1 + x^{-1})^d$ is increasing and has $\varphi \notin \mathbb{N}$ as a root, α and β are both strictly positive. We will show the following:

Theorem 2.4. *A tight upper bound for the price of anarchy for unweighted network (or congestion) games with cost functions in \mathcal{C}_d is*

$$\frac{\alpha}{\alpha + \beta} (k + 1)^{d+1} + \frac{\beta}{\alpha + \beta} k^{d+1}. \tag{2.18}$$

Proof. We begin with the lower bound construction. Let the players be $J = J_0 \cup \bar{J}_0$, with

$$J_0 = \{p_1, p_2, \dots, p_k, q_1, q_2, \dots, q_{k+1}\}.$$

Let the items be $I = I_0 \cup \bar{I}_0$, with

$$I_0 = \{u_1, u_2, \dots, u_{k+1}, v_1, v_2, \dots, v_{k+1}\}.$$

Define the cost functions as $f_i = a_i x^d$, for all $i \in I_0$ where:

$$\begin{aligned} a_{u_i} &= \beta \left(\frac{k}{k+1} \right)^d & i \leq k \\ a_{u_{k+1}} &= \alpha \beta \left(\frac{1}{k+1} \right)^d \\ a_{v_i} &= \alpha \left(\frac{k+1}{k+2} \right)^d & i \leq k \\ a_{v_{k+1}} &= \alpha \left(\frac{(k+1)^d - \beta}{(k+2)^d} \right). \end{aligned}$$

The cost functions for items in \bar{I}_0 are defined symmetrically.

We define the strategies for players in J_0 as follows:

$$\begin{aligned} S_{p_j} &= \{u_i : i \leq k+1\} & j \leq k \\ S_{q_j} &= \{v_i : i \leq k+1\} & j \leq k+1 \\ S_{p_j}^* &= \{\bar{u}_j\} & j \leq k \\ S_{q_j}^* &= \{\bar{v}_j\} & j \leq k \\ S_{q_{k+1}}^* &= \{\bar{u}_{k+1}, \bar{v}_{k+1}\} \end{aligned}$$

Again, for players in \bar{J}_0 we define the strategies in a symmetric fashion.

We now show that the strategy distribution $\mathcal{P} = \{S_j : j \in J\}$ is a Nash. First, for all $j \leq k$,

$$\begin{aligned} C_{p_j}(\mathcal{P}) &= \sum_{i=1}^{k+1} a_{u_i} k^d \\ &= \frac{1}{(k+1)^d} (\beta k^d \cdot k + \alpha \beta) k^d \\ &= \beta k^d, \quad \text{using (2.16).} \\ C_{p_j}(\mathcal{P}^{(p_j)}) &= \beta \left(\frac{k}{k+1} \right)^d (k+1)^d \\ &= \beta k^d. \end{aligned}$$

Also,

$$\begin{aligned}
C_{q_j}(\mathcal{P}) &= \sum_{i=1}^{k+1} a_{v_i} (k+1)^d \\
&= \alpha (k+1)^d, \quad \text{using (2.17).} \\
C_{q_j}(\mathcal{P}^{(q_j)}) &= \alpha \left(\frac{k+1}{k+2} \right)^d (k+2)^d \\
&= \alpha (k+1)^d.
\end{aligned}$$

Finally,

$$\begin{aligned}
C_{q_{k+1}}(\mathcal{P}) &= \alpha (k+1)^d. \\
C_{q_{k+1}}(\mathcal{P}^{(q_{k+1})}) &= a_{u_{k+1}} (k+1)^d + a_{v_{k+1}} (k+2)^d \\
&= \alpha \beta + \alpha ((k+1)^d - \beta) \\
&= \alpha (k+1)^d.
\end{aligned}$$

So the Nash requirements are satisfied. The optimal flow is obtained by taking all of the S_p^* strategies, where every item is used only once. Thus the price of anarchy is

$$\begin{aligned}
\rho &= \frac{\sum_{i=1}^{k+1} a_{u_i} k^{d+1} + \sum_{i=1}^{k+1} a_{v_i} (k+1)^{d+1}}{\sum_{i=1}^{k+1} a_{u_i} + \sum_{i=1}^{k+1} a_{v_i}} \\
&= \frac{\beta k^{d+1} + \alpha (k+1)^{d+1}}{\beta + \alpha},
\end{aligned}$$

as required.

There is some intuition behind this construction. There are two groups of items; the u_i 's are used by k players at Nash, and the v_i 's by $k+1$ players. Because $k+1 > \varphi$, the q_j players can't quite fit their optimal strategies within the v_i group; on the other hand, because $k < \varphi$, the u_i group has some "extra space" available. The cost functions are scaled in the different groups to match the "excess" of the v_i 's with the amount of "extra space" from the u_i group.

We now provide a matching upper bound. First a lemma; the proof is annoyingly technical, and has been relegated to the appendix.

Lemma 2.2. *For all $r, s \in \mathbb{N}$,*

$$(r+1)^d s \leq \mu r^{d+1} + \nu s^{d+1},$$

where

$$\mu = \frac{(k+2)^d - (k+1)^d}{(k+1)^{d+1} - k^{d+1}}, \quad \nu = (k+1)^d - \mu k^{d+1}.$$

We will again assume that the transformation of the previous section has been applied, so that all cost functions can be written $f_e(x) = a_e x^d$. We begin again with (2.13):

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{j \in J} \sum_{e \in P_j^*} f_e(x_e + w_j) w_j \\ &= \sum_{e \in E} a_e (x_e + 1)^d x_e^* \\ &\leq \sum_{e \in E} a_e (\mu x_e^{d+1} + \nu x_e^{*d+1}) \\ &= \mu C(\mathcal{P}) + \nu C(\mathcal{P}^*). \end{aligned}$$

Now

$$\begin{aligned} \mu &= \frac{(k+2)^d - (k+1)^d}{(k+1)^{d+1} - k^{d+1}} \\ &= \frac{(k+1)^{d+1} - \beta - (k+1)^d}{(k+1)^{d+1} - k^{d+1}} \\ &= 1 - \frac{\alpha + \beta}{(k+1)^{d+1} - k^{d+1}}. \end{aligned}$$

Thus

$$\begin{aligned}
\frac{C(\mathcal{P})}{C(\mathcal{P}^*)} &\leq \frac{\nu}{1-\mu} \\
&= \frac{(k+1)^d - \mu k^{d+1}}{1-\mu} \\
&= \frac{\alpha + (1-\mu)k^{d+1}}{1-\mu} \\
&= \alpha \frac{(k+1)^{d+1} - k^{d+1}}{\alpha + \beta} + k^{d+1} \\
&= \frac{\alpha}{\alpha + \beta} (k+1)^{d+1} + \frac{\beta}{\alpha + \beta} k^{d+1}.
\end{aligned}$$

This matches the lower bound construction. □

Chapter 3

The Flow-free Model

In this chapter, we will consider a new variation on the “classical” model that we have considered so far. We will first motivate and define the model, and then proceed to some results on the price of anarchy, first for nonatomic agents and then for atomic agents. We will consider linear and polynomial cost functions, and most of our results will be tight.

3.1 The model

The classical model considered in the previous chapter has the property that all players using an edge experience the same latency. This makes sense if we think of the players as continuously routing flow. For example, the model is a reasonable one for users streaming audio or video across the internet. But suppose we are interested in rush hour traffic. Here, the assumption of equal latency is no longer as reasonable; cars that use a road earlier will cause congestion to later traffic, but not the reverse.

We provide a very simple modification to the classical model which incorporates this effect. We will first give an informal description of the model in the nonatomic case, where all players are infinitesimally small. In the classical model, on an edge




Figure 3.1: The cost is defined by the area of the greyed region in our new model, as opposed to the area of the large rectangle.

with utilisation x_e , all players pay a rate of $f_e(x_e)$, giving a total cost of $f_e(x_e)x_e$ for that edge. In our new model, the players on an edge have some priority ranking. For example, the priority could be based on the time at which the agents arrived at the start of the edge; agents who arrive earlier get higher priority. For any agent j , let $x_e^{(j)}$ be the amount of flow with higher priority than j along edge e ; this depends on the current routing of the players. We now dictate that player j 's latency on edge e is not $f_e(x_e)$ as in the classical model, but $f_e(x_e^{(j)})$. The total contribution to the social cost by edge e will then be $\int_0^{x_e} f_e(z)dz$. See Figure 3.1 for a pictorial representation. The area of the greyed region is exactly the above integral.

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).

The atomic case

As before, we have a directed network $G = (V, E)$, and n players $J = \{1, 2, \dots, n\}$. Player j wishes to route traffic of size w_j from vertex s_j to vertex t_j . As before, each edge has an associated cost functions $f_e(x)$. But in addition, we must define some

kind of priority scheme on the edges. We will allow this to be very general - the priority ordering on an edge can depend arbitrarily on the current routing \mathcal{P} . Later we will look at special cases for the priority scheme. If player i has lower priority than player j on edge e under routing \mathcal{P} , we write $i \prec_{e,\mathcal{P}} j$. For a fixed e and \mathcal{P} , the relation $\prec_{e,\mathcal{P}}$ 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 to avoid notational clutter.

Let us now consider some particular priority schemes that seem natural.

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.

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.

The timestamp game: This particular variant was the inspiration for the flow-free model. 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. So it is not perhaps completely clear that we have enough information to uniquely determine the priorities. To see that we do, imagine simulating the game. If we take the player j with the smallest value of τ_j , that player must have the highest priority on the first edge e_1 of her path. So we can imagine moving her to u , the second vertex of her path P_j . Her timestamp at u will be $\tau_j + \int_0^{w_j} f_{e_1}(z) dz$. Now we continue by again taking the player

with the lowest timestamp (this could be the same player as in the first step, if the updated timestamp is still lower than that of the other players). When a player moves in our simulation, she will always have a priority lower than any players who have already moved along that edge, but higher than those yet to move along it. When the simulation terminates, we have the priorities we require.

There is also the small difficulty of ties - two agents taking the same edge who happen to have the same timestamp. We resolve this by simply prescribing a tie-breaking order on the players. Alternatively, perturbing the starting times by sufficiently small values will break the ties without modifying the orderings.

We can also consider the congestion game generalisation of this flow-free model. The generalisation is exactly analogous to the classical case - we simply remove the network structure and allow strategies to be arbitrarily specified subsets of items. Of the three specific models mentioned above, the first two generalise to this context; the timestamp game does not have a natural generalisation.

As in the informal discussion, 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.

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

$C_j(\mathcal{P})$ is again the portion of the total cost attributable to player j , which is

$$C_j(\mathcal{P}) = \sum_{e \in E} \int_{x_e^{(j)}(\mathcal{P})}^{x_e^{(j)}(\mathcal{P}) + w_j} f_e(x) dx. \quad (3.1)$$

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

$$C_j(\mathcal{P}) \leq C_j(\mathcal{P}') \quad (3.2)$$

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.

The nonatomic case

In the limit as $\max w_j \rightarrow 0$, we obtain the nonatomic case, where the agents are negligibly small. Intuitively, there is not much difficulty here, and all of the specific models defined above for the atomic case would seem to carry over easily to the nonatomic case. We have to be somewhat careful if we are to define this rigorously however. Our treatment, which requires a little measure theory, is inspired by a 1973 paper by Schmeidler [17], which first discussed nonatomic games.

We denote the set of players R by an interval $[0, N]$, where $N \in \mathbb{R}^+$. The choice of N is completely irrelevant, and could be normalised to 1 if desired. We also have two measurable functions $s, t : R \rightarrow V$, which specify the origin and destination of each player respectively. Finally, we define a priority scheme exactly as before; for a specified routing \mathcal{P} , $\prec_{e,\mathcal{P}}$ is a relation defining a total ordering on edge e .

The value $x_e^{(r)}(\mathcal{P})$, the amount of flow on edge e with higher priority than player r under routing \mathcal{P} , still makes sense in the nonatomic case, although we must define it differently:

$$x_e^{(r)}(\mathcal{P}) = \int_{L_r} f_e d\mu, \quad (3.3)$$

where $L_r = \{q \in R : q \succ_{e,\mathcal{P}} r\}$. The integral is a Lebesgue integral.¹ The total latency experienced by player r , i.e. the time taken for the player to traverse from the source to the sink, is

$$\ell_r(\mathcal{P}) = \sum_{e \in P_r} f_e (x_e^{(r)}(\mathcal{P})). \quad (3.4)$$

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' ,

$$\ell_r(\mathcal{P}) \leq \ell_r(\mathcal{P}') \quad (3.5)$$

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

¹We need L_r to be Lebesgue measurable, which implies a requirement on the ordering \succ ; any even remotely reasonable ordering will satisfy this very technical requirement however.

The above is more complicated than the definition of the nonatomic case of the classical model. There is a good reason for this. In the classical model, all agents with the same origin-destination pair are essentially indistinguishable. Because of this, essentially everything of importance can be defined in terms of the flow vector \mathbf{x} . In our model, this is not the case; there is no way to formulate the Nash condition, for instance, in terms of just the flow, since the priority scheme can encode very complicated dependence on the routing. Hence our model is more reminiscent of general nonatomic games.

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

Existence of pure 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. In particular, the following is an example in the fixed priority game. Consider the network shown in Figure 3.2. 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.

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 3.3.

Now for a positive result: in the global priority model, even in the atomic un-




Figure 3.2: A fixed-priority game with no pure Nash equilibria.




Figure 3.3: Widget to imitate an undirected edge $e = (v, w)$ with directed edges (dashed arcs have zero cost).

splittable case, there is always a pure Nash (as long as the cost functions are at least non-negative 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 a shortest path given the congestion effects of the higher priority agents that have already been routed.

It should be possible to show existence in the nonatomic case under some weak assumptions on the priority scheme. In particular, we conjecture that a pure Nash always exists, as long as the priority scheme is such that $x_e^{(r)}$ depends continuously on the flow \mathcal{P} , for every edge e and player r .

A correspondence with the classical model

The optimal flows in the flow-free model can be linked quite nicely with the classical model, especially in the nonatomic case:

Lemma 3.1. *Given an instance $G = (V, E)$ of the (atomic or nonatomic) flow-free 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*

$$\hat{f}_e(x) = \frac{1}{x} \int_0^x f_e(z) dz. \quad (3.6)$$

Proof. This follows by noting that the cost of a flow \mathbf{x} in the flow-free 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 3.1. *In a nonatomic flow-free game, the optimal flows are exactly the Nash 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 game, an old result [2, 11] also discussed in Roughgarden's book [14, Section 2.4]: A flow \mathbf{x} is optimal for a classical nonatomic game with continuously differentiable, semiconvex² cost functions \hat{f}_e iff it is a Nash for a game on the same network, where the cost functions are replaced by

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

But if the \hat{f}_e 's are defined as in Equation (3.6), then $f_e^*(y) = f_e(y)$, and the result follows. □

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

3.2 The price of anarchy of nonatomic agents

We begin by considering the case of nonatomic agents, i.e. where there are an infinite number of players, each controlling an infinitesimal amount of flow. We will obtain tight bounds for linear and polynomial cost functions. First a useful inequality:

Theorem 3.1. *For any Nash flow \mathcal{P} , under any priority scheme,*

$$C(\mathcal{P}) \leq \sum_{e \in E} f_e(x_e) x_e^*. \quad (3.7)$$

Proof. Let $\mathcal{P}^* = \{P_r^* : r \in R\}$ be some assignment of paths to players that obtains an optimum flow \mathbf{x}^* (so formally, it is a valid routing such that $\int_{r: e \in P_r^*} 1 \, d\mu = x_e^*$). Apply (3.5) 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 (3.4). 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

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

□

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 more transparent.

Theorem 3.2. *In the nonatomic case with linear cost functions, and for any priority scheme, 4 is an upper bound on the price of anarchy.*

Proof. Let $f_e(x) = a_e x + b_e$ for all $e \in E$. Note the following, for any flow vector \mathbf{x}' :

$$\begin{aligned} \sum_{e \in E} f_e(x'_e) x'_e &= 2 \sum_{e \in E} \frac{1}{2} a_e x_e'^2 + \frac{1}{2} b_e x'_e \\ &\leq 2 \sum_{e \in E} \int_0^{x'_e} a_e x + b_e dx \\ &= 2C(\mathbf{x}'). \end{aligned} \tag{3.8}$$

Beginning with the result of Theorem 3.1, we again use the technique from [5].

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{e \in E} f_e(x_e) x_e^* \\ &= \sum_{e \in E} f_e(x_e^*) x_e^* + \sum_{e \in E} (f_e(x_e) - f_e(x_e^*)) x_e^* \\ &\leq 2C(\mathcal{P}^*) + \sum_{e \in E: x_e \geq x_e^*} (f_e(x_e) - f_e(x_e^*)) x_e^* \quad \text{from (3.8)}. \end{aligned}$$

Following exactly Equation (1.3), 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 (3.8) 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. This requires a generalisation of the technique used for Theorem 3.2 and Theorem 1.1.

Theorem 3.3. *For the nonatomic case with polynomial cost functions of maximum degree d , $(d+1)^{d+1}$ is an upper bound for the price of anarchy.*

Proof. The proof uses a generalisation of the technique used to prove Theorem 3.2.

Let $\alpha \geq 1$ be a constant to be chosen later. We have

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

Now consider:

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

Let $g(\phi) = \phi - \alpha \phi^{d+1}$. Since $g'(\phi) = 1 - \alpha(d+1)\phi^d$, the maximum value of g occurs at $\phi_m = (\alpha(d+1))^{-1/d}$, giving

$$g(\phi) \leq \frac{d}{d+1} \cdot \frac{1}{(\alpha(d+1))^{1/d}}.$$

Thus

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

and since $\sum_{e \in E} f_e(x_e) x_e \leq (d+1)C(\mathcal{P})$,

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




Figure 3.4: Lower bound construction for polynomial cost functions.

yielding

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

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

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

completing the proof. \square

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.

Theorem 3.4. *For the nonatomic case with polynomial cost functions of maximum degree d , $(d+1)^{d+1}$ is a lower bound on the price of anarchy in the global priority model.*

Proof. Consider a network of the form shown in Figure 3.4. 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.

Let us calculate the price of anarchy for this network. 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). Let $x_{i,j}$ 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 ($x_{i,n+1} := 0$). Let $y_j = f_{e_j}(x_{j,j})$. It is easy to see that the Nash condition implies that

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

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

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

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

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

We can now get a lower bound on the cost of the Nash flow \mathcal{P} . 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

$$\begin{aligned}
C(\mathcal{P}) &= \sum_{j=1}^n \int_0^{x_{e_j}} f_{e_j}(x) dx \\
&= \sum_{j=1}^n \frac{1}{j(d+1)} x_{e_j}^{d+1} \\
&= \frac{1}{d+1} \sum_{j=1}^n \frac{1}{j} (j y_j)^{1+1/d} \\
&= \frac{1}{d+1} \sum_{j=1}^n j^{1/d} y_j^{1+1/d} \\
&\geq \frac{1}{d+1} \sum_{j=1}^n j^{1/d} [(d+1)((j+1)^{-1/d} - (n+2)^{-1/d})]^{d+1} \\
&= (d+1)^d \sum_{j=1}^n j^{1/d} [(j+1)^{-1/d} - (n+2)^{-1/d}]^{d+1}. \tag{3.9}
\end{aligned}$$

We can rewrite the statement of Lemma 2.1 as

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

Apply this to (3.9) with $a = (j+1)^{-1/d} - (n+2)^{-1/d}$, $b = (n+2)^{-1/d}$ and $d = d+1$ to obtain, for any constant $0 < \gamma < 1$,

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

We deal with each term separately. We have

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

and so the second term is $\mathcal{O}(1)$. For the first term, we have

$$\begin{aligned} \sum_{j=1}^n j^{-1} \left(1 + \frac{1}{j}\right)^{-1-1/d} &= \sum_{j=1}^n j^{-1} - \sum_{j=1}^n j^{-1} \left(1 - \left(1 + \frac{1}{j}\right)^{-1-1/d}\right) \\ &= H_n + \mathcal{O}(1). \end{aligned}$$

The proof that the second term is convergent is given in the appendix. Thus

$$C(\mathcal{P}) = \gamma^{d-1} (d+1)^d H_n + \mathcal{O}(1).$$

The optimal flow \mathcal{P}^* is clearly obtained by routing all of the flow from s_i through e_i for each i . This yields a cost of

$$C(\mathcal{P}^*) = \frac{1}{d+1} \sum_{i=1}^n \frac{1}{i} = \frac{H_n}{d+1}.$$

We thus get a bound for the price of anarchy:

$$\begin{aligned} \frac{C(\mathcal{P})}{C(\mathcal{P}^*)} &\geq \frac{\gamma^d (d+1)^d H_n + \mathcal{O}(1)}{(d+1)^{-1} H_n} \\ &= \gamma^d (d+1)^{d+1} + o(1). \end{aligned}$$

Thus 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 time-stamp 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 time-stamps 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 price of anarchy of our model with polynomial latency functions.

3.3 Unsplittable atomic agents

In this section, 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.

We begin with a useful inequality that holds for any Nash flow \mathcal{P} . As usual let \mathcal{P} be a Nash flow, \mathcal{P}^* be an optimal flow, and define $\mathcal{P}^{(j)} = \mathcal{P} \setminus P_j \cup P_j^*$, where everyone follows \mathcal{P} except player j . Using equations (3.1) and (3.2),

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

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

$$\begin{aligned} C(\mathcal{P}) &\leq \sum_{j \in J} \sum_{e \in P_j^*} \int_{x_e}^{x_e + w_j} f_e(x) dx \\ &= \sum_{e \in E} \sum_{j: e \in P_j^*} \int_{x_e}^{x_e + w_j} f_e(x) dx. \end{aligned} \quad (3.10)$$

3.3.1 Linear cost functions

Theorem 3.5. *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 flow respectively. Writing Equation (3.10) in the linear case with $f_e(x) = a_e x + b_e$, we obtain

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

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 $\alpha^2 \leq 2\alpha + 1$, whence 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

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

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

$$a_1 = a_{\bar{1}} = \frac{2\sqrt{2}}{3 + 2\sqrt{2}}, \quad (3.11)$$

$$a_2 = a_{\bar{2}} = \frac{3}{3 + 2\sqrt{2}}, \quad (3.12)$$

$$a_3 = a_{\bar{3}} = 2\sqrt{2} - 1. \quad (3.13)$$

We claim that if all players pick strategy S_j , we have a Nash. 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 . Then:

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

$$\begin{aligned} C_2 &= \int_{w_1}^{w_1+w_2} f_1(x) + f_2(x) dx \\ &= \int_1^2 x dx \\ &= \frac{3}{2}. \end{aligned}$$

$$\begin{aligned} C_3 &= \int_{w_1+w_2}^{w_1+w_2+w_3} f_1(x) + f_2(x) dx \\ &= \int_2^{\sqrt{2}+1} x dx \\ &= \sqrt{2} - \frac{1}{2}. \end{aligned}$$

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

$$\begin{aligned} \tilde{C}_1 &= \int_{w_1+w_2+w_3}^{w_1+w_2+w_3+w_1} f_1(x) dx \\ &= \int_{\sqrt{2}+1}^{\sqrt{2}+2} \frac{2\sqrt{2}}{3+2\sqrt{2}} x dx \\ &= \sqrt{2}. \end{aligned}$$

$$\begin{aligned} \tilde{C}_2 &= \int_{w_1+w_2+w_3}^{w_1+w_2+w_3+w_2} f_2(x) dx \\ &= \int_{\sqrt{2}+1}^{\sqrt{2}+2} \frac{3}{3+2\sqrt{2}} x dx \\ &= \frac{3}{2}. \end{aligned}$$




Figure 3.5: A network game construction with a price of anarchy of $3 + 2\sqrt{2}$.

$$\begin{aligned}
 \tilde{C}_3 &= \int_{w_1}^{w_1+w_3} f_3(x) dx \\
 &= \int_1^{\sqrt{2}} (2\sqrt{2} - 1)x dx \\
 &= \sqrt{2} - \frac{1}{2}.
 \end{aligned}$$

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 3.5. 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 $3 + 2\sqrt{2}$.

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 (3.11) through (3.13). 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 will not demonstrate the exact construction here, since it is somewhat complicated and not very edifying. Once we have this, we also can obtain a timestamp game construction by judicious choice of starting times. In particular, if we set the start times as

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

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

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

Theorem 3.6. *For unweighted agents, the price of anarchy is at most $\frac{17}{3}$.*

Proof. We need the following lemma:

Lemma 3.2. *Let $i, j \geq 0$ be integers. Then*

$$(2i + 1)j \leq \frac{2}{5}i^2 + \frac{17}{5}j^2. \quad (3.14)$$

Proof. Note that

$$2(i - \frac{5}{2}j)^2 \geq 5j - \frac{9}{2}j^2. \quad (3.15)$$

To see this, note that if $j \geq 2$ the right hand side is negative; if $j = 1$ it is $\frac{1}{2}$, and the left hand side is at least $\frac{1}{2}$ because i is an integer. Simplifying this equation yields the result. \square

Now:

$$\begin{aligned}
C(\mathcal{P}) &\leq \sum_{e \in E} (a_e x_e + b_e) x_e^* + \sum_{e \in E} \sum_{j: e \in P_j^*} \frac{1}{2} a_e w_i^2 \\
&= \sum_{e \in E} a_e \left(x_e + \frac{1}{2}\right) x_e^* + \sum_{e \in E} b_e x_e^* \quad \text{using } w_i = w_i^2 \\
&\leq \sum_{e \in E} \frac{1}{2} a_e \left(\frac{2}{5} x_e^2 + \frac{17}{5} x_e^{*2}\right) + \sum_{e \in E} b_e x_e^* \quad \text{using Lemma 3.2} \\
&\leq \frac{2}{5} C(\mathcal{P}) + \frac{17}{5} C(\mathcal{P}^*).
\end{aligned}$$

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 $S_{\bar{j}} = \bar{S}_j$ 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

$$C_1(\mathcal{P}) = \frac{3}{2} = C_1(\mathcal{P}^*)$$

$$C_2(\mathcal{P}) = \frac{9}{2} = C_2(\mathcal{P}^*)$$

$$C_3(\mathcal{P}) = \frac{5}{2} = C_3(\mathcal{P}^*)$$

Hence \mathcal{P} is a Nash. The price of anarchy is

$$\frac{(a_1 + a_2)\frac{1}{2} \cdot 3^2 + (a_3 + a_4)\frac{1}{2} \cdot 2^2}{(a_1 + a_2 + a_3 + a_4)\frac{1}{2} \cdot 1^2} = \frac{\frac{9}{2} + 4}{\frac{3}{2}} = \frac{17}{3},$$

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.

3.3.2 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, simply 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.

Theorem 3.7. *The price of anarchy is $\mathcal{O}(2^d d^d)$ in the unsplittable atomic case with polynomial cost functions of maximum degree d .*

Proof. Begin from Equation (3.10):

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

Now apply Lemma 2.1, with $a = x_e$ and $b = w_j$, and γ to be determined later:

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

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 $i \geq 1$,

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

The technique used for the nonatomic case is applicable to the first term (see the proof of Theorem 3.3). We thus obtain, for any $\alpha \geq 1$ and $0 < \gamma < 1$,

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

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 $\rho = \mathcal{O}(2^d d^d)$. □

3.4 Single-commodity networks

Single-commodity networks refer to the case where all agents have the same source s and destination t . It is reasonable to expect the price of anarchy to perhaps be reduced in this case; this is indeed true, at least for some choices of priority schemes. We only require the cost functions be continuous, non-negative and increasing for the following results.

First, a lemma:

Lemma 3.3. *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 an optimal flow. In other words, for any $s-t$ path P with $x_P > 0$, and all $s-t$ paths P' ,*

$$\sum_{e \in P} f_e(x_e) \leq \sum_{e \in P'} f_e(x_e). \quad (3.16)$$

Proof. This follows directly from Corollary 3.1, since \mathbf{x} is clearly a Nash in the classical version of the game. □

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 3.8. *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. Consider Equation (3.5). In our case, it can be written

$$\ell_r(\mathcal{P}) \leq f_{e'}(x_{e'}^{(r)}) \quad \text{for all } e' \in E, \quad (3.17)$$

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 (3.17) then yields

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

Hence the result follows by Lemma 3.3. \square

We can obtain a similar result for general single-commodity networks if we restrict the priority scheme:

Theorem 3.9. *The nonatomic versions of both the global priority and timestamp games have a price of anarchy of one in single-commodity networks.*

Proof. We first show that in the single-commodity case, the timestamp game is exactly equivalent to the global priority game. Take any two players $r, s \in R$ whose routes in the Nash routing \mathcal{P} intersect, and where the start times satisfy $\tau_r < \tau_s$. Then for any edge $e \in P_r \cap P_s$, r must arrive at the start of this edge earlier than s —for if not, r could change her route to be the same as s 's route until edge e , hence arriving earlier and contradicting the Nash requirement.

So we need consider only the global priority game. Take any path P on which \mathcal{P} has non-zero flow. Consider player r , the lowest priority agent that takes path




Figure 3.6: A single-commodity game with price of anarchy larger than one.

P . Since we are at Nash, this player has no incentive to switch; hence in particular, $\ell_r(\mathcal{P}) \leq \sum_{e \in P'} f_e(x_e)$ for any $s - t$ path P' . Thus Lemma 3.3 applies, and \mathcal{P} is an optimal routing. \square

The previous theorem does not hold if we consider fixed priority games instead. Consider the simple network shown in Figure 3.6. Take $R = [0, 1]$ for the set of players, and set all the cost functions to x . For those edges marked $>$, the priority ordering is defined by $r \succ s$ iff $r > s$; for the edge marked $<$, $r \succ s$ iff $r < s$. It can easily be checked that the routing which sends all players in $[0, 1/3]$ along the bottom path and all players in $(1/3, 1]$ along the top path is a Nash. This has a social cost larger than the optimum obtained by splitting the flow evenly between the two paths.

These results are in strong contrast to the classical model, where Pigou's two-link network yields the largest possible price of anarchy in most cases [12].

Conclusion

In this thesis, we have proved some new results on the price of anarchy of the classical game in the case of atomic and unsplittable flow. We have also introduced and analysed a new model that takes priorities of agents into account.

Plenty of questions remain, and there is plenty of scope for further work. Here we mention some of the more interesting avenues for exploration.

We have very little in the way of existence or uniqueness results in our model. It should be possible to prove existence in the nonatomic case under some fairly general restrictions on the priority scheme. It would be interesting to know if the timestamp game always has a Nash in the atomic case. In those cases where a pure Nash need not exist, it may be possible to extend our results to handle mixed strategy Nash equilibria. Another avenue would be to investigate the so called “price of sinking” introduced by Goemans et al. [7].

We have not considered the atomic splittable case in the flow-free model. There are some indications that, unlike for the classical model, the price of anarchy may be no larger than for the nonatomic case.

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; for certain applications, concave cost functions might be interesting.

In the nonatomic case, we have shown that the price of anarchy is one for single-

commodity flow in the global priority and timestamp games. On the other hand, our construction that maximised the price of anarchy required taking a limit where the number of sources tends to infinity. So a natural question to ask is whether we can obtain bounds that depend on the number of source-destination pairs.

We have captured the simple idea that a car only causes congestion to traffic behind it, rather than before it, in the timestamp game. But this is very crude; a car will only cause congestion to cars *a short time* behind it; rush-hour traffic in the morning has no effect on rush-hour traffic in the evening. It may be possible to model a car's effect as some kind of "hump" which decays with time.

While we have concentrated almost exclusively on calculations of the price of anarchy in this thesis, there are other aspects of our model that could be investigated. For instance, in some situations the addition of an arc or arcs to a network can *increase* the price of anarchy. This rather counterintuitive behaviour is called *Braess's paradox*. This feature has received substantial attention in the classical model, and it would be interesting to see how similar (or different) the behaviour is in the flow-free model. Roughgarden [14] has a thorough review on the topic.

Appendix A

Calculations

Proof of Lemma 2.2. Let

$$h(r, s) = (r + 1)^d s - \mu r^{d+1} - \nu s^{d+1},$$

where

$$\mu = \frac{(k + 2)^d - (k + 1)^d}{(k + 1)^{d+1} - k^{d+1}},$$

$$\nu = (k + 1)^d - \mu k^{d+1}.$$

We must show $h(r, s) \leq 0$ for all $r, s \in \mathbb{N}$. We first consider the case $s = 1$, so let $h_1(r) = h(r, 1)$. Then

$$\begin{aligned} h'_1(x) &= d(x + 1)^{d-1} - (d + 1)\mu x^d \\ &= x^{d-1} \left[d\left(1 + \frac{1}{x}\right)^{d-1} - (d + 1)\mu \right]. \end{aligned}$$

Thus $h'_1(x)/x^{d-1}$ is clearly decreasing for $x > 0$, positive at $x = 0$, and negative for x sufficiently large; let x_0 be the unique positive zero. Then (by multiplying through by x^{d-1}) it is clear that $h'_1(x_0) = 0$, $h'_1(x) > 0$ for $x \in (0, x_0)$, and $h'_1(x) < 0$ for $x \in (x_0, \infty)$. Now notice that by the choice of μ and ν ,

$$h_1(k) = h_1(k + 1) = 0$$

(recall $k = \lfloor \varphi \rfloor$). It follows by the intermediate value theorem that there is a turning point of h_1 in the interval $(k, k+1)$, i.e. $x_0 \in (k, k+1)$. We now have that for $r \leq k$, $h_1(r) \leq h_1(k) = 0$ since $h_1(x)$ is increasing on $[0, k]$, and for $r \geq k+1$, $h_1(r) \leq h_1(k+1) = 0$ since $h_1(x)$ is decreasing on $[k+1, \infty)$. So $h(r, 1) \leq 0$ for all $r \in \mathbb{N}$.

Now consider $s \geq 2$. Let $\tilde{h}(r, s) = h(r, s)/s^{d+1}$.

$$\begin{aligned} \tilde{h}(r, s) &= \left(\frac{r+1}{s}\right)^d - \mu \left(\frac{r}{s}\right)^{d+1} - \nu \\ &\leq \left(\frac{r}{s} + 1\right)^d - \mu \left(\frac{r}{s}\right)^{d+1} - \nu \\ &= h_1\left(\frac{r}{s}\right) \\ &\leq 0 \quad \text{for } \frac{r}{s} \in [0, k] \cup [k+1, \infty). \end{aligned}$$

On the other hand, suppose $k < \frac{r}{s} < k+1$. Then $\frac{r+1}{s} \leq k+1$ (since, $r, s \in \mathbb{N}$), and so

$$\begin{aligned} \tilde{h}(r, s) &\leq (k+1)^d - \mu \left(\frac{r}{s}\right)^{d+1} - \nu \\ &= (k+1)^d - \mu \left(\left(\frac{r}{s}\right)^{d+1} - k^{d+1} \right) - (k+1)^d \\ &\leq 0. \end{aligned}$$

Thus $h(r, s) = \tilde{h}(r, s)s^{d+1} \leq 0$ for all $r, s \in \mathbb{N}$.

□

Lemma A.1. *The series*

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

is convergent for all $d \geq 1$.

Proof.

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

Hence

$$0 \leq \frac{1}{j} \left(1 - \left(1 + \frac{1}{j}\right)^{-1-1/d}\right) \leq \frac{3}{(j+1)^2},$$

and so since

$$\sum_{j=1}^{\infty} \frac{3}{(j+1)^2}$$

converges, the result follows by the comparison test. □

Bibliography

- [1] B. Awerbuch, Y. Azar, and A. Epstein. Large the price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 57–66, New York, NY, USA, 2005. ACM Press.
- [2] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [3] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, New York, NY, USA, 2005. ACM Press.
- [4] R. Cominetti, J. Correa, and N. Stier-Moses. Network games with atomic players. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, volume 4051 of *Lecture Notes in Computer Science*, pages 525–536, January 2006.
- [5] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games extended abstract. In *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3509 of *Lecture Notes in Computer Science*, pages 167–181, January 2005.
- [6] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoret. Comput. Sci.*, 348(2-3):226–239, 2005.

-
- [7] 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.
 - [8] S. Kontogiannis and P. Spirakis. Atomic selfish routing in networks: A survey. In *Proceedings of the 1st Workshop on Internet and Network Economics*, 2005.
 - [9] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, January 1999.
 - [10] C. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, New York, NY, USA, 2001. ACM Press.
 - [11] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.
 - [12] 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, New York, NY, USA, 2002. ACM Press.
 - [13] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, May 2002.
 - [14] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, May 2005.
 - [15] T. Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1184–1185, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
 - [16] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, March 2002.

- [17] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, April 1973.
- [18] M. J. Smith. The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological*, 13:295–304, December 1979.
- [19] 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.