Optimal Price of Anarchy of Polynomial and Super-Polynomial Bottleneck Congestion Games

Rajgopal Kannan[1], Costas Busch[1], and Athanasios V. Vasilakos[2]

[1] Dept. of Computer Science, Louisiana State University, Baton Rouge, LA 70803
{rkannan,busch}@csc.lsu.edu

[2] Dept. of Telecomm. Engg., Univ. of Western Macedonia, Greece
vasilako@ath.forthnet.gr

Abstract. We introduce *(super) polynomial bottleneck games*, where the utility costs of the players are (super) polynomial functions of the congestion of the resources that they use, and the social cost is determined by the worst congestion of any resource. In particular, the delay function for any resource r is of the form $C_r^{\mathcal{M}_r}$, where C_r is the congestion measured as the number of players that use r, and the degree of the delay function is bounded as $1 \leq \mathcal{M}_r \leq \log C_r$. The utility cost of a player is the sum of the individual delays of the resources that it uses. The social cost of the game is the worst bottleneck resource congestion: $\max_{r \in R} C_r$, where R is the set of resources. We show that for super-polynomial bottleneck games with $\mathcal{M}_r = \log C_r$, the price of anarchy is $o(\sqrt{|R|})$, specifically $O(2^{\sqrt{\log |R|}})$. We also consider general polynomial bottleneck games where each resource can have a distinct monomial latency function but the degree is bounded i.e $\mathcal{M}_r = O(1)$ with constants $\alpha \leq \mathcal{M}_r \leq \beta$ and derive the price of anarchy as $\min\left(|R|, \max\left(\frac{2\beta}{C^*}, (2|R|)^{\frac{1}{\alpha+1}} \cdot (\frac{2\beta}{C^*})^{\frac{\alpha}{\alpha+1}} \cdot (2\beta)^{\frac{\beta-\alpha}{\alpha+1}}\right)\right)$, where C^* is the bottleneck congestion in the socially optimal state. We then demonstrate matching lower bounds for both games showing that this price of anarchy is tight.

1 Introduction

We consider non-cooperative congestion games with n players, where each player has a *pure strategy profile* from which it selfishly selects a strategy that minimizes the player's utility cost function (such games are also known as *atomic* or *unsplittable-flow* games). We focus on *bottleneck congestion games* where the objective for the social outcome is to minimize C, the maximum congestion on any resource. Typically, the congestion on a resource is a non-decreasing function on the number of players that use the resource; here, we consider the congestion to be simply the number of players that use the resource.

Bottleneck congestion games have been studied in the literature [2,4,3] in the context of routing games, where each player's utility cost is the worst resource

R. Jain and R. Kannan (Eds.): GameNets 2011, LNICST 75, pp. 308–320, 2012.

congestion on its strategy. For any resource r, we denote by C_r the number of users that use r in their strategies. In typical bottleneck congestion games, each player i has utility cost function $C_i = \max_{r \in S_i} C_r$, where S_i is the strategy of the player. The social cost is worst congested resource: $C = \max_i C_i = \max_r C_r$.

In [2] the authors observe that bottleneck games are important in networks for various practical reasons. In networks, each resource corresponds to a network link, each player corresponds to a packet, and a strategy represents a path for the packet. In wireless networks, the maximum congested link is related to the lifetime of the network since the nodes adjacent to high congestion links transmit large number of packets which results to higher energy utilization. High congestion links also result in congestion hot-spots which may slow-down the network throughput. Hot spots also increase the vulnerability of the network to malicious attacks which aim to to increase the congestion of links in the hope to bring down the network. Thus, minimizing the maximum congested edge results to hot-spot avoidance and more load-balanced and secure networks.

In networks, bottleneck games are also important from a theoretical point of view since the maximum resource congestion is immediately related to the optimal packet scheduling. In a seminal result, Leighton *et al.* [16] showed that there exist packet scheduling algorithms that can deliver the packets along their chosen paths in time very close to $C + D$, where D is the maximum chosen path length. When $C \gg D$, the congestion becomes the dominant factor in the packet scheduling performance. Thus, smaller C immediately implies faster packet delivery time.

A natural problem that arises in games concerns the effect of the players' selfishness on the welfare of the whole system measured with the *social cost* C. We examine the consequence of the selfish behavior in pure *Nash equilibria* which are stable states of the game in which no player can unilaterally improve her situation. We quantify the effect of selfishness with the *price of anarchy* (PoA) [15,21], which expresses how much larger is the worst social cost in a Nash equilibrium compared to the social cost in the optimal coordinated solution. The price of anarchy provides a measure for estimating how closely do Nash equilibria of bottleneck congestion games approximate the optimal C^* of the respective coordinated optimization problem.

Ideally, the price of anarchy should be small. However, the current literature results have only provided weak bounds for bottleneck games. In [2] it is shown that if the resource congestion delay function is bounded by some polynomial with degree k (with respect to the packets that use the resource) then $PoA = O(|R|^k)$, where R is the set of links (resources) in the graph. In [4] the authors consider bottleneck routing games for the case $k = 1$ and they show that $PoA = O(L + \log |V|)$, where L is the maximum path length (maximum number of resources) in the players' strategies and V is the set of nodes in the network. This bound is asymptotically tight (within poly-log factors) since it is shown in [4] that there are game instances with $PoA = \Omega(L)$. Since $L = |R|$, the price of anarchy has to be large, $PoA = \Omega(|R|)$.

1.1 Contributions

The lower bound in [4] suggests that in order to obtain better price of anarchy in bottleneck congestion games (where the social cost is the bottleneck resource C), we need to consider alternative player utility cost functions. Towards this goal, we introduce *(super) polynomial bottleneck games* where the player cost functions are (super) polynomial expressions of the congestions along the resources. In particular, the player utility cost function for player i is: $C'_i = \sum_{r \in S_i} C_r^{\mathcal{M}_r}$, where for each r, the *degree* $\mathcal{M}_r$ ranges between $1 \leq \mathcal{M}_r \leq \log C_r$. Note that the new utility cost is a sum of polynomial or super-polynomial terms on the congestion of the resources in the chosen strategy (instead of the max that we described earlier). The social cost remains the maximum bottleneck congestion C, the same as in typical congestion games.

The new player utility costs have significant benefits in improving both the upper and lower bounds on the price of anarchy for the bottleneck social cost C. Of specific interest are instantiations of the game where the degree $\mathcal{M}_r$ is either a logarithmic function on the congestion, or simply a constant. For *super-polynomial games* with $\mathcal{M}_r = \log C_r$ for each $r \in R$, we prove that the price of anarchy is $o(\sqrt{|R|})$:

$$PoA_{super} = O(2^{\sqrt{\log |R|}}) \tag{1}$$

For general *polynomial games* with $\alpha \leq \mathcal{M}_r \leq \beta$ and constants $1 \leq \alpha \leq \beta$, the price of anarchy is:

$$PoA_{poly} = \min\left(|R|, \max\left(\frac{2\beta}{C^*}, (2|R|)^{\frac{1}{\alpha+1}} \cdot (\frac{2\beta}{C^*})^{\frac{\alpha}{\alpha+1}} \cdot (2\beta)^{\frac{\beta-\alpha}{\alpha+1}}\right)\right) \tag{2}$$

Note that in polynomial games, latency costs are monomials with constant degrees between α and β and thus different resources can have different latency costs even with same congestion. We show that the bounds in Equations 1 and 2 are asymptotically tight by providing specific instances of super-polynomial and polynomial bottleneck games. The price of anarchy bounds above are significant improvements over the price of anarchy from the typical bottleneck games described above.

Polynomial and super-polynomial congestion games are interesting variations of bottleneck games not only because they provide good price of anarchy but also because they represent interesting and important real-life problems. In networks, the overall delay that a packet experiences is directly related with the link congestions along the path and hence the polynomial utility cost function reflects the total delivery delay. In wireless networks, polynomial and super-polynomial player utilities correspond to the total energy that a packet consumes while it traverses the network, and the social cost reflects to the worst energy utilization in any node in the network. Similar benefits from polynomial congestion games appear in the context of job-shop scheduling, where computational tasks require resources to execute. In this context, the social bottleneck cost function C represents the task load-balancing efficiency of the resources, and the player utility costs relate to the makespan of the task schedule. In all the above problems, the degrees $\mathcal{M}_r$ are chosen appropriately to model precisely the involved costs of the resource utilization in each computational environment.

In our analysis, we obtain the price of anarchy upper bound by using two techniques: *transformation* and *expansion*. Consider a game G with a Nash equilibrium S and congestion C. We identify two kinds of players in S: type-A players which use only one resource in their strategies, and type-B players which use two or more resources. In our first technique, transformation, we convert G to a simpler game $\tilde{G}$, having a Nash equilibrium $\tilde{S}$ with congestion $\tilde{C}$, such that $\tilde{C} = O(C)$, and all players in $\tilde{S}$ with congestion above a threshold τ are of type-A; that is, we transform type-B players to type-A players. Having type-A players is easier to bound the price of anarchy. Then, we use a second technique, expansion, which is used to give an upper bound on the price of anarchy of game $\tilde{G}$, which implies an upper bound on the price of anarchy of the original game G.

In [12], we have derived upper bounds for the price of anarchy of games with exponential utility cost functions using similar techniques (transformation and expansion). While exponential cost games have a unique substructure which makes the analysis of Price of Anarchy simpler, we believe these techniques are general enough to adapt in a non-trivial manner for a large class of utility cost functions. For the case of exponential cost games, we obtained logarithmic price of anarchy upper bounds, which was related to the problem structure. Here we obtain tight (optimal) price of anarchy bounds for polynomial and super-polynomial bottleneck games, using a non-trivial application of the general transformation and expansion techniques.

1.2 Related Work

Congestion games were introduced and studied in [20,22], mainly in the context of networks. Typically the social cost is considered to be an aggregate expression on the delay costs of the network edges and the flow that goes through them [5,24,25,26,27]. Specifically, the social cost is $SC = \sum_r C_r \cdot l(C_r)$, where $l(C_r)$ is the delay cost function for resource r, while the player cost is the same as our polynomial game model. Any price of anarchy bounds using social cost SC can be translated to a price of anarchy bound on bottleneck congestion C by amortizing appropriately with the number of resources R. For example, when the latency function is a monomial of the same constant degree d on all resources, the same bounds can be obtained using this method as well as ours (using the $d^{\Theta(d)}$ social cost bound in [1]). However, for atomic congestion games, price of anarchy bounds for SC are not known for super-polynomial delay functions, or for polynomial delay functions of different degrees for the resources, as we consider in this paper. Thus, the techniques that we propose here are useful in providing novel bounds in bottleneck congestion games for a broader range of delay functions.

In [22], Rosenthal proves that congestion games have always pure Nash equilibria. Koutsoupias and Papadimitriou [15] introduced the notion of price of anarchy in the specific *parallel link networks* model in which they provide the price of anarchy bound 3/2. Roughgarden and Tardos [25] provided the first result for splittable flows in general networks in which they showed that the price of anarchy is bounded by 4/3 for a player cost which reflects to the sum of congestions of the resources of a path. Pure equilibria with atomic flow have been studied in [4,5,17,27] (our work fits into this category), and with splittable flow in [23,24,25,26]. Mixed equilibria with atomic flow have been studied in [7,10,14,15,18,19,21], and with splittable flow in [6,9]. The vast majority of the work on congestion games has been performed for parallel link networks, with only a few exceptions on general network topologies [4,5,6,23]. Our work immediately applies to network topologies.

Basic bottleneck routing games have been studied in [2,4] which consider the maximum congestion metric in general networks, and the player cost is equal to the worst congested edge in the chosen routing path. In [2] the authors show the existence and non-uniqueness of equilibria in both the splittable and atomic flow models. They show that finding the best Nash equilibrium that minimizes the social cost is a NP-hard problem. Further, they show that the price of anarchy may be unbounded for specific resource congestion functions. In [3] the $C + D$ social metric is considered. In [11], the authors prove the existence of strong Nash equilibria, which concern coalitions of players, for games with the lexicographic improvement property. Other variations of basic bottleneck games with player coalitions are studied in [8].

Outline of Paper

In Section 2 we give basic definitions. In Section 3 we convert games with type-B players to games with type-A players. In Section 4 we give a bound on the price of anarchy. We finish with providing a lower bound in Section 5.

2 Definitions

A *congestion game* is a strategic game $G = (\Pi_G, R, \mathbb{S}, (l_r)_{r\in R}, (pc_\pi)_{\pi\in\Pi_G})$ where:

- $\Pi_G = \{\pi_1, \ldots, \pi_n\}$ is a non-empty and finite set of players.
- $R = \{r_1, \ldots, r_z\}$ is a non-empty and finite set of resources.
- $\mathbb{S} = \mathbb{S}_{\pi_1} \times \mathbb{S}_{\pi_2} \times \cdots \times \mathbb{S}_{\pi_n}$, where $\mathbb{S}_{\pi_i}$ is a strategy set for player π_i, such that $\mathbb{S}_{\pi_i} \subseteq powerset(R)$; namely, each strategy $S_{\pi_i} \in \mathbb{S}_{\pi_i}$ is pure, and it is a collection of resources. A *game state* (or *pure strategy profile*) is any $S \in \mathbb{S}$. We consider *finite games* which have finite $\mathbb{S}$ (finite number of states).
- In any game state S, each resource $r \in R$ has a *latency cost* denoted $l_r(S)$.
- In any game state S, each player $\pi \in \Pi_G$ has a player cost $pc_\pi(S) = \sum_{r\in S_\pi} l_r(S)$.

Consider a game G with a state $S = (S_{\pi_1}, \ldots, S_{\pi_n})$. The *(congestion) of a resource* r is defined as $C_r(S) = |\{\pi_i : r \in S_{\pi_i}\}|$, which is the number of players that use r in state S. The *(bottleneck) congestion of a set of resources* $Q \subseteq R$ is defined as $C_Q(S) = \max_{r\in Q} C_r(S)$, which is the maximum congestion over all resources in Q. The *(bottleneck) congestion of state* S is denoted $C(S) = C_R(S)$, which is the maximum congestion over all resources in R. When the context is clear, we will drop the dependence on S. We examine polynomial congestion games:

- *Polynomial games:* The latency cost function for any resource r is $l_r = C_r^{\mathcal{M}_r}$, for some integer constants $\mathcal{M}_l \leq \mathcal{M}_r \leq \mathcal{M}_h$.
- *Super-polynomial games:* The delay cost function for any resource r is $d_r = C_r^{\mathcal{M}_r}$, where $\mathcal{M}_r = \log C_r$.

For any state S, we use the standard notation $S = (S_{\pi_i}, S_{-\pi_i})$ to emphasize the dependence on player π_i. Player π_i is *locally optimal* (or *stable*) in state S if $pc_{\pi_i}(S) \leq pc_{\pi_i}((S'_{\pi_i}, S_{-\pi_i},))$ for all strategies $S'_{\pi_i} \in \mathbb{S}_{\pi_i}$. A greedy move by a player π_i is any change of its strategy from S'_{π_i} to S_{π_i} which improves the player's cost, that is, $pc_{\pi_i}((S_{\pi_i}, S_{-\pi_i})) < pc_{\pi_i}((S'_{\pi_i}, S_{-\pi_i}))$. *Best response dynamics* are sequences of greedy moves by players. A state S is in a *Nash Equilibrium* if every player is locally optimal. Nash Equilibria quantify the notion of a stable selfish outcome. In the games that we study there could exist multiple Nash Equilibria.

For any game G and state S, we will consider a *social cost* (or *global cost*) which is simply the bottleneck congestion $C(S)$. A state S^* is called *optimal* if it has minimum attainable social cost: for any other state S, $C(S^*) \leq C(S)$. We will denote $C^* = C(S^*)$. We quantify the quality of the states which are Nash Equilibria with the *price of anarchy* (PoA) (sometimes referred to as the coordination ratio). Let $\mathcal{P}$ denote the set of distinct Nash Equilibria. Then the price of anarchy of game G is:

$$PoA(G) = \sup_{S \in \mathcal{P}} \frac{C(S)}{C^*},$$

We continue with some more special definitions that we use in the proofs. Consider a game G with a socially optimal state $S^* = (S^*_{\pi_1}, \ldots, S^*_{\pi_n})$, and let $S = (S_{\pi_1}, \ldots, S_{\pi_n})$ denote the equilibrium state.

For any resource $r \in R$, we will let Π_r and Π^*_r denote the set of players with r in their equilibrium and socially optimal strategies respectively, i.e $\Pi_r = \{\pi_i \in \Pi_G | r \in S_{\pi_i}\}$ and $\Pi^*_r = \{\pi_i \in \Pi_G | r \in S^*_{\pi_i}\}$.

Let $G = (\Pi_G, R, \mathbb{S}, l, (pc_\pi)_{\pi \in \Pi_G})$ and $\tilde{G} = (\Pi_{\tilde{G}}, \tilde{R}, \tilde{\mathbb{S}}, \tilde{l}, (\tilde{pc}_\pi)_{\pi \in \Pi_{\tilde{G}}})$ be two games.

Definition 1. *G η-**dominates** $\tilde{G}$ if the following conditions hold between them for the highest cost Nash equilibrium and optimal states: $|\tilde{R}| \leq |R|$, $l = \tilde{l}$, $\tilde{C} \geq C$, $\tilde{C}^* = O(\eta C^*)$, where η is any parameter independent of congestion C. Here C, C^*, $\tilde{C}$ and $\tilde{C}^*$ represent the bottleneck congestions in the highest cost Nash equilibrium and optimal states of G and $\tilde{G}$, respectively.*

Corollary 1. *$PoA(G) \leq \eta \cdot PoA(\tilde{G})$ for an arbitrary game G and dominated game $\tilde{G}$.*

3 Type-B to Type-A Game Transformation

Our approach for obtaining the PoA of an arbitrary game G is to first transform it to a simplified game $\tilde{G}$ with restricted player strategies, obtain the PoA of the restricted game (which should be easier to evaluate than the generic game G) and relate this to the PoA of the unrestricted version G. Transformed game $\tilde{G}$ will consist of players with drastically limited strategies in the equilibrium state. Specifically, for a given game G in equilibrium state S, we consider two special kinds of players with respect to state S:

– *Type-A players:* any player π_i with $|S_{\pi_i}| = 1$.
– *Type-B players:* any player π_i with $|S_{\pi_i}| \geq 2$.

We define type-B games as those containing an arbitrary mix of type-A and type-B players in state S. Thus the type-B label refers to any generic monotonic-bounded congestion game. We define type-A games as those in which highly congested resources (beyond a specific latency-cost dependent threshold that we will define subsequently) are occupied only by type-A players in equilibrium state S. Intuitively, type-A games should be easier to analyze since the equilibrium strategy of players are highly restricted.

Let $\eta > 0$ be a network-related constant (i.e independent of bottleneck congestion). Let τ be an arbitrary congestion threshold such that $\forall r \in R, \forall C_r \geq \tau : \frac{l_r(C_r+1)}{l_r(C_r)} \leq \eta$. For super-polynomial games, where $l_r(C_r) = (C_r)^{\log C_r}$, we can choose τ as any small constant with $\eta = e^2$. For general polynomial games where $l_r(C_r) = (C_r)^{\mathcal{M}_r}$, we can choose $\tau = \max_r \mathcal{M}_r$ with $\eta = e$.

Consider a game $G(S, C, S^*, C^*)$ where S denotes the Nash equilibrium state with the highest social cost (the one having the highest bottleneck congestion) C, and S^* is the socially optimal state with corresponding bottleneck congestion C^*.

Then we have,

Theorem 1. *Every type-B game $G(S, C, S^*, C^*)$ with polynomial or super-polynomial latency costs on resources can be transformed into a type-A game $\tilde{G}(\tilde{S}, \tilde{C}, \tilde{S}^*, \tilde{C}^*)$ in which all resources r with congestion $C_r \geq \tau$ in equilibrium state $\tilde{S}$ are utilized exclusively by type-A players.*

Theorem 2. *Transformed game $\tilde{G}(\tilde{S}, \tilde{C}, \tilde{S}^*, \tilde{C}^*)$ is 7-dominated by $G(S, C, S^*, C^*)$. Specifically, bottleneck congestion in optimal states S^* and $\tilde{S}^*$ of G and $\tilde{G}$ satisfies $C^* \leq \tilde{C}^* \leq 7C^*$ while bottleneck congestion in Nash equilibrium states S and $\tilde{S}$ are the same $C = \tilde{C}$. The Price of Anarchy of G is bounded by $PoA(G) \leq \max(\frac{\tau}{C^*}, 7 \cdot PoA(\tilde{G}))$.*

Proof Sketch of Theorems 1 and 2: We describe a constructive proof of the theorems by iteratively transforming type-B players in G to type-A players in $\tilde{G}$.

We initialize $\tilde{G}$, the input to our transformation algorithm as a restricted version of game G with exactly two strategies per player: $\tilde{S}_\pi = S_\pi$ and $\tilde{S}^*_\pi = S^*_\pi$. We will iteratively transform $\tilde{G}$ into a type-A game by converting all type-B players of cost $\geq l(\tau) + 1$ into type-A players, in phases in decreasing order of player costs. We add and delete players/resources from $\tilde{G}$ iteratively and have a working set of players. However $\tilde{G}$ will always remain in equilibrium state $\tilde{S}$ at every step of the transformation process. When we add a new player π_k to

$\tilde{\Pi}$ we will assign two strategy sets to π_k: an 'equilibrium' strategy $\tilde{S}_{\pi_k}$ and an optimal strategy $\tilde{S}^*_{\pi_k}$. Thus $\tilde{S} = \tilde{S} \bigcup \tilde{S}_{\pi_k}$ and $\tilde{S}^* = \tilde{S}^* \bigcup \tilde{S}^*_{\pi_k}$.

First we convert $\tilde{G}$ into a 'clean' version in which every type-B player $\pi \in \tilde{\Pi}$ has *distinct* resources in its equilibrium and optimal strategies i.e $\tilde{S}_\pi \bigcap \tilde{S}^*_\pi = \emptyset$. If not already true, this can be achieved by creating $|\tilde{S}_\pi \bigcap \tilde{S}^*_\pi|$ new type-A players with identical and one type-B player with disjoint equilibrium and optimal strategies for each original player π. The new type-B player has $\tilde{S}_\pi - \tilde{S}^*_\pi$ and $\tilde{S}^*_\pi - \tilde{S}_\pi$ as its equilibrium and optimal strategy respectively while the new type-A players each use one resource from $|\tilde{S}_\pi \bigcap \tilde{S}^*_\pi|$ as their identical equilibrium and optimal strategies. Note that the new players are also in equilibrium in $\tilde{S}$.

Let π_i be an arbitrary type-B-player using k resources $r_1, r_2, \ldots, r_k$ in its equilibrium strategy $\tilde{S}_{\pi_i}$ that are distinct from the m resources $r^*_1, \ldots, r^*_m$ in its optimal strategy $\tilde{S}^*_{\pi_i}$. Let C_{r_j}, $C_{r^*_j}$ denote the congestion on these resources in equilibrium state $\tilde{S}$. Without loss of generality, assume the resources in $\tilde{S}_{\pi_i}$ and $\tilde{S}^*_{\pi_i}$ have been sorted in non-increasing order of congestion i.e $C_{r_1} \geq C_{r_2} \ldots \geq \ldots C_{r_m}$ and $C_{r^*_1} \leq C_{r^*_2} \ldots \leq \ldots C_{r^*_m}$. Then we have the following:

Lemma 1. *$\tilde{S}_{\pi_i}$ and $\tilde{S}^*_{\pi_i}$ can be partitioned into t pairs $(L_1, L^*_1), (L_2, L^*_2), \ldots, (L_t, L^*_t)$ where*

$$\sum_{r \in L^*_j} l(C_r + 1) \geq \sum_{r \in L_j} l(C_r), \quad 1 \leq j \leq t \tag{3}$$

and further

1. *The L_j's form a disjoint resource partition of $\tilde{S}_{\pi_i}$ i.e $L_j \bigcap L_k = \emptyset$ with $\bigcup_{j=1}^t L_j = \tilde{S}_{\pi_i}$.*
2. *$|L^*_j \bigcap L^*_{j+1}| \leq 1$, for $1 \leq j \leq t$. If $|L^*_j \bigcap L^*_{j+1}| = 1$ then the last resource in L^*_j is the first resource in L^*_{j+1}.*
3. *$\forall j : 1 \leq j \leq t$, either $|L_j| = 1$ or $|L^*_j| = 1$ or both. If $|L_j| > 1$ and $|L^*_j| = 1$ with $L^*_j = \{r^*_p\}$ we must have $C_{r^*_p} \geq \max\{C_r | r \in L_j\}$.*
4. *r^*_m appears at most once in a partition (specifically L^*_t) while r^*_1 appears in at most two partitions. If r^*_1 appears in two partitions then at least one of the partitions contains only one resource (i.e r^*_1). Every other resource $r^*_p \in \tilde{S}^*_{\pi_i}$, $2 \leq p \leq m-1$ appears in at most three partitions. If r^*_p appears in three partitions then two of the partitions contains only r^*_p. If r^*_p appears in two partitions then it is the last resource in the first partition and the first resource in the second partition.*

We label the procedure implementing lemma 1 as Procedure PMS−Partition(). This procedure is used to create new players and forms the basic step in our transformation algorithm. We ensure the equilibrium of these new players in $\tilde{G}$ using the key constructs of *exact matching sets* and *potential matching sets.*

A set of resources $\tilde{R}$ in $\tilde{G}$ forms an exact matching set for a newly created player π_k with newly assigned equilibrium strategy $\tilde{S}_{\pi_k}$ if $\sum_{r \in \tilde{R}} (C_r + 1)^{\mathcal{M}} \geq pc_{\pi_k}(\tilde{S}_{\pi_k}, \tilde{S}_{-\pi_k}) = \sum_{r \in \tilde{S}_{\pi_k}} C_r^{\mathcal{M}}$. Clearly, $\tilde{R}$ can be assigned as the new optimal strategy $\tilde{S}^*_{\pi_k}$ in game $\tilde{G}$ without violating the equilibrium of π_k.

Potential matching sets are defined for newly created type-B players. A potential matching set $\tilde{R}$ is an exact matching set that can 'potentially' be added to the optimal set of resources $\tilde{S}^*_{\pi_k}$ of a type-B player $\pi_k \in \tilde{G}$ without increasing the optimal bottleneck congestion in $\tilde{G}$ from original game G by an η factor i.e $\tilde{C}^* = O(\eta C^*)$.

Now consider a type-B player π_i to be transformed. We partition the resources in its equilibrium and optimal strategies $\tilde{S}_{\pi_i}$ and $\tilde{S}^*_{\pi_i}$ according to PMS−Partition(π_i) and remove it from $\tilde{G}$, i.e $\tilde{S} = \tilde{S} - \tilde{S}_{\pi_i}$ and $\tilde{S}^* = \tilde{S}^* - \tilde{S}^*_{\pi_i}$.

Consider those partition-pairs (L_j, L_j^*) with $|L_j| = 1$. We can create a new type-A player π_k and add it to to $\tilde{G}$ with an equilibrium strategy $\tilde{S}_{\pi_k}$ that is the singleton resource in L_j. Due to the condition in Eq. 3, the set of resources in L_j^* forms an exact matching set for π_k and can therefore be assigned to $\tilde{S}^*_{\pi_k}$. π_k is in equilibrium in $\tilde{G}$ and the equilibrium and optimal congestion on resources in $\tilde{S}_{\pi_k}$ and $\tilde{S}^*_{\pi_k}$ are now the same as before. This forms the 'easy' part of the transformation process.

Consider however, those partitions (L_j, L_j^*) with $1 < |L_j| \leq |R|$ and $L_j^* = \{r_l^*\}$. Similar to the above, we can create $|L_j|$ new type-A players and assign a distinct resource in L_j to each such players equilibrium strategy. However if, as above, we assign r_l^*, the single resource in L_j^*, to each players optimal strategy, we might increase the socially optimal congestion $\tilde{C}^*$ of $\tilde{G}$ to as much as $C^* + |R|$, thereby violating the domination of G over $\tilde{G}$. Thus we need to find an appropriate potential matching set from among existing resources and assign them to these players, without increasing the optimal congestion beyond $O(\eta C^*)$. Finding such a set is the 'hard' part of the transformation process.

We form game $\tilde{G}$ by transforming type-B players in distinct phases corresponding to decreasing values of player cost functions $l_r(C_r)$ from $max_r l_r(C)$ down to $min_r l_r(\tau)$. During phase i with congestion index C_i, we transform all type-B players with costs in the range $l_{r_j}(C_i)$ down to $l_{r_k}(C_{i-1})+1$ into type-A players, where $l_{r_j}()$ and $l_{r_k}()$ are respective maximizers in initial game G. To find new potential matching sets for all type-B players without increasing the optimal congestion in $\tilde{G}$ beyond a constant factor, we utilize the set of resources with congestion exactly C_i to generate For details of the proof, please see [13].

4 Price of Anarchy

4.1 Price of Anarchy for Type-A Player Games

By Theorem 2 and Corollary 1, for every type-B game G, we can find a dominated type-A game $\tilde{G}$ such that

$$PoA(\hat{G}) \leq 7 \cdot PoA(\tilde{G}) \tag{4}$$

Thus we only need to find an upper bound on the PoA of type-A games. Consider a generic type-A game G with optimal solution $S^* = (S^*_{\pi_1}, \ldots, S^*_{\pi_n})$, optimal congestion C^*, and a Nash equilibrium state $S = (S_{\pi_1}, \ldots, S_{\pi_n})$ with the highest congestion C among all Nash equilibria of G. Define a congestion threshold ψ such that $\forall r \forall C_r : C_r \geq \psi$, we have $C_r \geq \frac{l_r(C_r+1)}{l_r(C_r)} \cdot C^*$. For general polynomial games with $l_r(C_r) = C_r^{\mathcal{M}_r}$, we can choose $\psi = 2 \max_r(\mathcal{M}_r, C^*)$. For superpolynomial games with $l_r(C_r) = C_r^{\log C_r}$, we can choose $\psi = e^2 \cdot C^*$. We specify G as a type-A game in which all resources r with $C_r > \psi$ are utilized only by type-A players in in state S. As a consequence of Eq. 4, we can bound the PoA of arbitrary type-B games by bounding the ratio C/C^*.

We first define a resource graph $\mathcal{N}$ for state S. There are $V = V_1 \bigcup V_2$ nodes in $\mathcal{N}$. Each resource $r \in R$ with $C_r > \psi$ ($C_r \leq \psi$, resp.) corresponds to the equivalent node $r \in V_1$ ($r \in V_2$). Henceforth we will use the term resource and node interchangeably. For every player π using a resource $x \in V_1$ in equilibrium, there is a directed edge (x, y) between node x and all nodes $y \in V$, where $y \neq x$ is in the optimal strategy set of π i.e $S_\pi = x$ and $y \in S^*_\pi$. We use the notation $\mathrm{Ch}(x)$ to denote the set $\bigcup_{\pi : S_\pi = x} S^*_\pi$. Note that there could be multiple incoming links to a node x from the same node, however there are no self-loops and x can be the child of at most C^* nodes. Also note that nodes in V_2 are terminal nodes that have no outgoing links.

Recursively counting the number of descendants of the root node in $\mathcal{T}$ will help us relate the number of resources $|R|$ with the parameters C and C^* and obtain our PoA bound. However since $\mathcal{N}$ can have cycles, we instead modify $\mathcal{N}$ to remove cycles and replace it with a Directed Acyclic Graph (DAG) $\mathcal{T}$ (without increasing the size of $\mathcal{N}$).

Lemma 2. *Resource graph $\mathcal{N}$ can be transformed into expansion DAG $\mathcal{T}$ without affecting the equilibrium state S and optimal congestion C^*, where $|\mathcal{T}| \leq |\mathcal{N}|$.*

Since $\mathcal{T}$ is a DAG we know that it has sink nodes (with outdegree 0). Every node in V_1 is an internal node (with non-zero indegree and outdegree) since it has congestion $> C^*$ and hence the sink nodes in $\mathcal{T}$ are nodes from V_2. Consider the DAG starting at the root node with congestion C. Let κ denote the PoA of G.

Lemma 3. *For DAG $\mathcal{T}$ with root node x and congestion $C_x = C$, it holds that*

$$\sum_{r \in \mathrm{Descendants}(x) \bigcap V_2} l_r(\psi) \geq (\kappa - 1) l_x(C)$$

Using the fact that $|V_2| \leq |R|$, and denoting $l_{max}()$ and $l_{min}()$ as the respective maximizers/minimizers among the latency cost functions, we get

$$\kappa l_{min}(C) \leq 2|R| l_{max}(\psi) \tag{5}$$

Note that the the PoA of superpolynomial and general polynomial games is at least ψ/C^*.

For superpolynomial games using the fact that $\forall r \in R : l_r(C_r) = C_r^{\log C_r}$ and $\psi = e^2 \cdot C^*$ and simplifying, we get that $\log \kappa + (\log \kappa)^2 \leq \log |R|$. Hence for type-$A$ superpolynomial games the Price of Anarchy is bounded by

$$PoA_{super-poly} = O(2^{\sqrt{\log |R|}})$$

For general polynomial games with $l_r(C_r) = C_r^{\mathcal{M}_r}$, let β and α respectively, denote the largest and smallest degrees in $\mathcal{M}_r$. Substituting in Eq. 5, we get $\kappa C^\alpha \leq 2|R| \cdot \psi^\beta$ or $\kappa^{\alpha+1} \leq 2|R| \cdot (\frac{\psi}{C^*})^\alpha \cdot \psi^{\beta-\alpha}$. Hence for type-$A$ general polynomial games the Price of Anarchy is bounded by

$$PoA_{poly} = \min\left(|R|, \max\left(\frac{\psi}{C^*}, (2|R|)^{\frac{1}{\alpha+1}} \cdot (\frac{\psi}{C^*})^{\frac{\alpha}{\alpha+1}} \cdot \psi^{\frac{\beta-\alpha}{\alpha+1}}\right)\right)$$

Substituting $\psi = 2\max(\beta, C^*)$, we get the result.

5 Lower Bounds on Price of Anarchy

We demonstrate a simple lower bound to show the tightness of the above upper bound. Consider a type-A superpolynomial game G in which C players in Nash equilibrium state utilize the same resource r, i.e $C_r = C$. No other resource is being utilized in this state. The optimal strategy of each player has $C^{\log C}$ unique resources, i.e the congestion in optimal state S^* is $C^* = 1$. Clearly G is in equilibrium and the PoA is $\kappa = C$. The total number of resources R is given by $|R| = C \cdot C^{\log C} + 1$ and hence we have

$$\kappa = 2^{\sqrt{\log(|R|-1)}}$$

Similarly consider a type-A general polynomial game G in which C players in Nash equilibrium state are utilizing only one resource x, where $l_x(C_x) = (C_x)^\alpha$. Let N be any integer, $2 \leq N \leq C$. Each distinct subset of N out of these C players are sharing $\lceil \frac{C^\alpha}{N^\beta} \rceil$ unique resources in optimal state S^*. These resources are used by $N-1$ other players in equilibrium as well as optimal states S and S^*. Hence the equilibrium congestion of these resources is $N-1$ while congestion in optimal state S^* is $C^* = 2N-1$. Note thate there are $\lceil \frac{C}{N} \rceil$ sets of these resources. The latency cost on each of these resources r is $l_r(C_r) = (C_r)^\beta$.

Clearly G is in equilibrium and the PoA is $\kappa = \frac{C}{2N-1}$. The total number of resources $|R| = \frac{C}{N} \cdot \frac{C^\alpha}{N^\beta} + 1$ and hence we have

$$\kappa = \min\left(|R|, \max\left(\frac{C}{2N-1}, (|R|-1)^{\frac{1}{\alpha+1}} \cdot N^{\frac{\beta-\alpha}{\alpha+1}} \cdot \frac{N}{2N-1}\right)\right), 2 \leq N \leq C.$$

Note that this has the same form as the upper bound above.

6 Conclusions

We have considered bottleneck congestion games with polynomial and superpolynomial resource delay cost functions. The price of anarchy result for superpolynomial functions is $o(\sqrt{|R|})$ with respect to the number of resources. We also demonstrate two novel techniques, B to A player conversion and expansion which help us obtain this result. These techniques which enable us to simplify games for analysis are sufficiently general. In future work, we plan to use these techniques to analyze the PoA of games with arbitrary player cost functions.

References

1. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), Baltimore, MD, USA, pp. 57–66. ACM (May 2005)
2. Banner, R., Orda, A.: Bottleneck routing games in communication networks. IEEE Journal on Selected Areas in Communications 25(6), 1173–1179 (2007); also appears in INFOCOM 2006
3. Busch, C., Kannan, R., Vasilakos, A.V.: Quality of routing congestion games in wireless sensor networks. In: Proc. 4th International Wireless Internet Conference (WICON), Maui, Hawaii (November 2008)
4. Busch, C., Magdon-Ismail, M.: Atomic routing games on maximum congestion. Theoretical Computer Science 410(36), 3337–3347 (2009)
5. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), Baltimore, MD, USA, pp. 67–73. ACM (May 2005)
6. Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Computational Complexity, Fairness, and the Price of Anarchy of the Maximum Latency Problem. In: Bienstock, D., Nemhauser, G.L. (eds.) IPCO 2004. LNCS, vol. 3064, pp. 59–73. Springer, Heidelberg (2004)
7. Czumaj, Vocking: Tight bounds for worst-case equilibria. ACM Transactions on Algorithms (TALG) 3 (2007)
8. de Keijzer, B., Schäfer, G., Telelis, O.A.: On the Inefficiency of Equilibria in Linear Bottleneck Congestion Games. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 335–346. Springer, Heidelberg (2010)
9. Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Selfish unsplittable flows. Theoretical Computer Science 348(2-3), 226–239 (2005)
10. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: Computing Nash equilibria for scheduling on restricted parallel links. In: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC), Chicago, Illinois, USA, pp. 613–622 (June 2004)
11. Harks, T., Klimm, M., Möhring, R.H.: Strong Nash Equilibria in Games with the Lexicographical Improvement Property. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 463–470. Springer, Heidelberg (2009)
12. Kannan, R., Busch, C.: Bottleneck Congestion Games with Logarithmic Price of Anarchy. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 222–233. Springer, Heidelberg (2010)
13. Kannan, R., Busch, C.: Optimal price of anarchy for polynomial and super-polynomial bottleneck congestion games. LSU-CSC Tech Report 11-03 (2011)
14. Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: Approximate equilibria and ball fusion. Theory Comput. Syst. 36(6), 683–693 (2003)
15. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)

16. Leighton, F.T., Maggs, B.M., Rao, S.B.: Packet routing and job-scheduling in $O(congestion + dilation)$ steps. Combinatorica 14, 167–186 (1994)
17. Libman, L., Orda, A.: Atomic resource sharing in noncooperative networks. Telecomunication Systems 17(4), 385–409 (2001)
18. Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. Theoretical Computer Science 406(3), 187–206 (2008)
19. Mavronicolas, Spirakis: The price of selfish routing. Algorithmica 48 (2007)
20. Monderer, D., Shapely, L.S.: Potential games. Games and Economic Behavior 14, 124–143 (1996)
21. Papadimitriou, C.: Algorithms, games, and the Internet. In: ACM (ed.) Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), Hersonissos, Crete, Greece, pp. 749–753 (July 2001)
22. Rosenthal, R.W.: A class of games possesing pure-strategy Nash equilibria. International Journal of Game Theory 2, 65–67 (1973)
23. Roughgarden, T.: The maximum latency of selfish routing. In: Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), New Orleans, Louisiana (USA), pp. 980–981 (January 2004)
24. Roughgarden, T.: Selfish routing with atomic players. In: Proc. 16th Symp. on Discrete Algorithms (SODA), pp. 1184–1185. ACM/SIAM (2005)
25. Roughgarden, T., Tardos, É.: How bad is selfish routing. Journal of the ACM 49(2), 236–259 (2002)
26. Roughgarden, T., Tardos, É.: Bounding the inefficiency of equilibria in nonatomic congestion games. Games and Economic Behavior 47(2), 389–403 (2004)
27. Suri, S., Toth, C.D., Zhou, Y.: Selfish load balancing and atomic congestion games. Algorithmica 47(1), 79–96 (2007)