

Altruism, Selfishness, and Spite in Traffic Routing

Po-An Chen
University of Southern California
poanchen@usc.edu

David Kempe*
University of Southern California
dkempe@usc.edu

ABSTRACT

In this paper, we study the price of anarchy of traffic routing, under the assumption that users are partially altruistic or spiteful. We model such behavior by positing that the “cost” perceived by a user is a linear combination of the actual latency of the route chosen (selfish component), and the increase in latency the user causes for others (altruistic component). We show that if all users have a coefficient of at least $\beta > 0$ for the altruistic component, then the price of anarchy is bounded by $1/\beta$, for all network topologies, arbitrary commodities, and arbitrary semi-convex latency functions. We extend this result to give more precise bounds on the price of anarchy for specific classes of latency functions, even for $\beta < 0$ modeling spiteful behavior. In particular, we show that if all latency functions are linear, the price of anarchy is bounded by $4/(3 + 2\beta - \beta^2)$.

We next study non-uniform altruism distributions, where different users may have different coefficients β . We prove that all such games, even with infinitely many types of players, have a Nash Equilibrium. We show that if the average of the coefficients for the altruistic components of all users is β , then the price of anarchy is bounded by $1/\beta$, for single commodity parallel link networks, and arbitrary convex latency functions. In particular, this result generalizes, albeit non-constructively, the Stackelberg routing results of Roughgarden and of Swamy. More generally, we bound the price of anarchy based on the class of allowable latency functions, and as a corollary obtain tighter bounds for Stackelberg routing than a recent result of Swamy.

Categories and Subject Descriptors: G.2.3 [Mathematics of Computing]: Discrete Mathematics—Applications

General Terms: Economics, Theory

Keywords: altruism, spite, selfishness, routing, anarchy

*Supported in part by NSF CAREER award 0545855, and NSF grant DDDAS-TMRP 0540420

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’08, July 8–12, 2008, Chicago, Illinois, USA.

Copyright 2008 ACM 978-1-60558-169-9/08/07 ...\$5.00.

1. INTRODUCTION

One of the most basic and important problems in managing networks is to route traffic so as to make the latency experienced by the average user small. This problem can be solved effectively when all the traffic submits to the control of a central authority. However, neither in road networks nor in large-scale decentralized computer networks (such as the Internet) is it feasible to establish such a central authority. Rather, individual users of the network have control over the paths they choose from their origin to their destination. The prevailing assumption is that users will exert this power to choose the route minimizing their individual latency, regardless of the effects that such a choice may have on other users.

A natural question is then how much the average latency increases as a result of such selfish behavior, compared to a central authority balancing the latencies of different users. The ratio between the socially optimal outcome and the outcome of selfish choices has been termed “Price of Anarchy” (PoA) by Koutsoupias and Papadimitriou [23]. Roughgarden and Tardos [36] pioneered the study of the PoA for traffic routing networks. They analyze a model proposed by Wardrop and Beckmann et al. [42, 2], in which edges possess traffic-dependent latency functions. When users choose a certain path, they increase the traffic on all edges of the path, and thus also the latency experienced by all other users sharing the path.

This model of selfishness assumes, in accordance with much of the game theory literature, that users choose their routes completely without regard to the delay that their choice may cause for other users in the system. The assumption of selfishness (as well as that of rationality) has been repeatedly questioned by economists and psychologists. Time and again, experiments have shown that even for simple games in controlled environments, participants do not act selfishly [24, 25]; their behavior can be either altruistic or malicious. Many explanations have been considered for this phenomenon, including an innate sense of fairness [12], reciprocity among agents [17], or spite and altruism [25].

In this work, we investigate the question whether and how the Price of Anarchy in traffic routing will change if users are assumed to be “not entirely selfish”. To this end, we consider a natural model of altruism and spite, extending one proposed by Ledyard [24, p. 154]. Intuitively, we want to model that users will trade off the benefit to themselves against the benefit to others. This can be modeled by assuming that the utility of each player is a linear combination of his own a priori payoff and the payoffs of other players. In

the context of traffic routing, the perceived “cost” of a player is a linear combination of his own latency and the increase in latency the player causes others (precise definitions are given in Section 2). By varying the *altruism coefficient* β , we can smoothly tune the altruism from spiteful ($\beta = -1$), through selfish ($\beta = 0$), to entirely altruistic ($\beta = 1$).

Our first result is that if all users are (at least) β -altruistic, and $\beta > 0$, then the Price of Anarchy is always bounded by $1/\beta$, for all networks, arbitrarily many commodities, and arbitrary semi-convex latency functions on the edges. Thus, if a constant amount of altruism is introduced into the system, then the PoA is bounded by a constant. A more general version of our result characterizes precisely the worst-case PoA for any class of latency functions; from this general result, better bounds can be obtained for more restricted classes of functions. Among others, our result implies a bound of $\frac{4}{3+2\beta-\beta^2}$ on the PoA if all latency functions are linear. The general bound also lets us analyze the *spite resistance* of a class of latency functions: the most spite under which the PoA would still be finite.

We next extend our results beyond uniform altruism, and consider arbitrary distributions of altruism among the players. In that scenario, even the existence of Nash Equilibria is not obvious; we use a theorem of Mas-Colell [28] to prove that such games with infinitely many agents indeed have Nash Equilibria. Even for single-commodity flows in arbitrary graphs, prohibitive lower bounds on the PoA are known [3], so we focus here on parallel link networks, studied for instance by Roughgarden [32].

For parallel link networks, we show that for any non-negative distribution of altruism over the users in the network with *average altruism level* $\bar{\beta}$, the price of anarchy with convex edge latency functions is always bounded by $1/\bar{\beta}$. In the specific case where the distribution of altruism has only completely altruistic or completely selfish users, this matches a bound obtained (with a polynomial-time algorithm) by Roughgarden [32].¹ The bound of $1/\bar{\beta}$ follows from a more general result characterizing the PoA for arbitrary classes of convex functions. In fact, that more general result, when applied to the case of a distribution over entirely selfish and entirely altruistic users, implies tighter bounds for Stackelberg routing compared with a recent result of Swamy [40]. Finally, we show that for the bound we derive, the worst case is in fact attained by the $\{0, 1\}$ altruism distribution, while the best case is when all users are β -altruistic.

1.1 Related Work

The study of the ineffectiveness of Selfish Routing was pioneered within the theory community by the groundbreaking work of Roughgarden and Tardos [36]. It was preceded by work in the economics and traffic engineering communities on congestion models, traffic routing, and the impact of tolls [29, 42, 2]. Since the original paper by Roughgarden and Tardos, a lot of progress has been made on different aspects of the problem, including different objectives [31], Stackelberg strategies in which an altruistic central authority controls a fraction of all traffic [22, 32, 40], the impact of tolls or taxes on the inefficiency [8, 9, 14, 15, 21], atomic games wherein users control non-infinitesimal amounts of

traffic [10, 18, 34], and the effects of network structure on the inefficiency [27, 30, 35]. For an excellent overview of many of these results, see the book by Roughgarden [33].

Among other things, our results draw a connection between Stackelberg strategies and tolls on users, in that the altruistic component of a user’s utility can be considered as a (traffic-dependent) toll, and entirely altruistic users act as though they submitted to the control of a benevolent authority. A more detailed discussions of the connections and related results is given in Section 2.2, after a formal definition of our model of altruism and spite.

Questions about the accuracy of the assumption that users are selfish and rational have been as old as the field of game theory (see, e.g., [24]). Different models have been proposed to model user preferences more accurately [24, 25]. A model somewhat similar to ours was recently studied in the context of contributions to P2P systems by Feldman et al. [13], who posited an intrinsic *generosity* parameter of users, their willingness to contribute to the system. They then study contribution dynamics and their equilibria, akin to many collective behavior scenarios studied by Schelling [37]. Trade-offs between individual optimization and social optimum in the context of traffic routing are also considered by Jahn et al. [19]. They posit that users will be willing to incur latency somewhat exceeding a “lowest possible” baseline if advised by a traffic routing system. They experimentally evaluate how centralized routing of users under this restriction compares with unrestricted centralized routing (which may place very heavy burdens on some users). More recently, Babaioff et al. [1] studied the impact of spiteful behavior on the outcome of routing games. In their model, there are two types of players: selfish rational players, and malicious players, who seek to maximize the average delay experienced by the rational players (while not caring about their own delay). They quantify the impact of malicious players on the equilibrium, and show that the price of anarchy can sometimes be increased, and in fact decreased at other times.

The notion of spiteful behavior by individuals, and models similar to the ones we are proposing here, have recently been studied in the context of auctions. For single-item auctions, Brandt and Weiss [5] study the behavior of “antisocial” agents, whose utility decreases in their competitors’ profit. They analyze optimum bidding strategies in a full-information setting in this model. Following up on this work, Brandt et al. [4] study the Bayesian setting, and derive symmetric Bayesian-Nash equilibria for spiteful agents in first-price and second-price sealed bid auctions. They show that the expected revenue in second-price auctions is higher than the expected revenue in first-price auctions when all agents are neither completely selfish nor completely spiteful. They also prove that in the presence of spite, complete information reduces the revenue in second-price auctions, while it increases the revenue in first-price auctions. Vetsikas and Jennings [41] generalize some of these results for multi-unit auctions, deriving symmetric Bayes-Nash equilibria for spiteful agents in both m^{th} and $(m + 1)^{\text{th}}$ price sealed bid auctions. Similarly, Liang and Qi [26] study the effects of cooperative or vindictive bidding strategies on the revenue of sponsored search auctions and the existence of truthful strategies and equilibria. Finally, mechanism design for spiteful agents in scheduling is considered by Garg et al. [16].

¹Roughgarden’s bound for Stackelberg routing on parallel link networks applies to arbitrary functions, whereas ours requires convexity.

2. PRELIMINARIES

Our model is based on the model of Wardrop [42], as described by Roughgarden and Tardos [33, 36]. We are given a (directed) graph $G = (V, E)$, in which each edge is equipped with a flow-dependent *latency function* $c_e(x)$. The meaning is that if the total flow on the edge e is x , then each user experiences a delay $c_e(x)$ on that edge. We assume that each c_e is a continuously differentiable and monotone non-decreasing function. In addition, for some of our results, we will assume that each c_e is convex, and for others that each c_e is *semi-convex*, i.e., that $x \cdot c_e(x)$ is convex.

We assume that users/agents are *non-atomic*, i.e., infinitesimally small. Thus, we can think of the total traffic as a multi-commodity flow with rates r_i between source-sink pairs (s_i, t_i) , where the total flow from s_i to t_i is r_i . If f_e denotes the total flow on edge e , then the total latency experienced by a user on a path P is $c_P(f) := \sum_{e \in P} c_e(f_e)$. The total latency experienced by all users is thus $C(f) := \sum_e f_e \cdot c_e(f_e)$. An instance of the routing problem is thus a triple (G, r, c) (where r and c are the vectors of flow rates and edge cost functions). The *socially optimum solution* for (G, r, c) is the flow f minimizing $C(f)$, and thus the solution to the convex program

$$\begin{aligned} & \text{Minimize} && \sum_e f_e \cdot c_e(f_e) \\ & \text{subject to} && f \text{ is a feasible multi-commodity flow for } (G, r, c). \end{aligned}$$

The constraints are the standard linear multi-commodity flow constraints; the objective function is convex so long as each c_e is semi-convex. Thus, the optimum can be computed in polynomial time using convex programming.

Selfish users do not care about the cost $C(f)$. Their sole goal is to select a path P minimizing their own latency $c_P(f)$. As the goals of different selfish users in minimizing their latency are conflicting with each other, the traffic routing problem can be considered a game, and the “outcome” of this game will be a *Nash Equilibrium*: a multi-commodity flow f such that, given f , no user has an incentive to choose a different path. Thus, a flow f is at Nash Equilibrium if and only if for each commodity i , all s_i - t_i paths P with $f_P > 0$ have the same latency $c_P(f)$, and all other s_i - t_i paths have at least the same latency. Nash Equilibria, too, can be computed as solutions to a convex program:

PROPOSITION 2.1 ([33], PROPOSITION 2.6.1). *The Nash flows of an instance (G, r, c) are exactly the solutions to the following convex program, and can thus be computed in polynomial time.*

$$\begin{aligned} & \text{Minimize} && \sum_e \int_0^{f_e} c_e(t) dt \\ & \text{subject to} && f \text{ is a feasible multi-commodity flow for } (G, r, c). \end{aligned}$$

If f is a flow at Nash Equilibrium, and f^* the socially optimum flow, then an interesting question, first investigated in detail by Roughgarden and Tardos [36], is how much larger $C(f)$ can be than $C(f^*)$. The ratio $\rho(G, r, c) := C(f)/C(f^*)$ is called the *price of anarchy* of the instance (G, r, c) . Roughgarden and Tardos [36] gave a generalization of Pigou’s example [29], showing that if the cost functions can be arbitrary, then the price of anarchy is unbounded, even for networks consisting of two nodes and two parallel links. On the other hand, they proved that if all functions are linear $c_e(x) = a_e x + b_e$, then the price of anarchy is at most $4/3$.

2.1 Altruism and Spite

The assumption that users are entirely selfish is simplistic, and not warranted in many scenarios. Indeed, experiments in economics have found time and again that users behave neither rationally nor selfishly, even in the absence of personal interaction or repeated experiments [24, 25]. Different models of such behavior have been proposed, including based on reciprocity [17], an innate notion of fairness [12], or altruism and spite [25]. We base our treatment on a simple and elegant suggestion of Ledyard [24]. In a game with n players, the utility of a player i given an action vector \mathbf{a} is $p_i(\mathbf{a}) + \beta_i \frac{1}{n} \sum_j p_j(\mathbf{a})$, where the p_i are the individuals’ payoff functions. The parameter β_i captures how important the average social welfare is to player i . We modify this approach slightly, and posit that user i ’s utility is the combination $(1 - \beta_i)p_i(\mathbf{a}) + \beta_i \frac{1}{n} \sum_j p_j(\mathbf{a})$, where $\beta_i \in [-1, 1]$ is the user’s *altruism level*. This has the advantage of making all utilities comparable on the same scale, and allowing us to model entirely altruistic behavior by setting $\beta_i = 1$.² We call $p_i(\mathbf{a})$ the *selfish* part of player i ’s utility, and $\frac{1}{n} \sum_j p_j(\mathbf{a})$ the *altruistic* part. If $\beta_i < 0$, then player i derives utility from a decrease in social utility; we call such players *spiteful*.

In order to apply this model to our scenario of traffic routing, we define the payoff of user i on path P as $p_i = -c_P(f)$, where f is the total flow, determined by the actions of all other players. Then, maximizing utility is equivalent to minimizing latency. The traffic routing model assumes that there are infinitely many users, each of whom is infinitesimally small. We can still define the utility function analogously, using the (well-defined) average latency of all users as the altruistic part. However, because users are infinitesimally small and latency functions continuous, the average latency of other users will not depend on an individual user’s action. Thus, as long as $\beta \neq 1$, each partially altruistic user will act exactly like a selfish user. A natural model considering the effect the user has on others should instead be based on the *rate* at which the user’s action will affect other users. We thus use the following definition of a β -altruistic user³:

DEFINITION 2.2. *Each β -altruistic user (for $\beta \in [-1, 1]$) chooses a path P so as to minimize the cost function*

$$c_P^{(\beta)}(f) := (1 - \beta) \sum_{e \in P} c_e(f_e) + \beta \sum_{e \in P} (f_e c_e'(f_e)).$$

The term $\sum_{e \in P} c_e(f_e)$ is the selfish part of the cost, while $\sum_{e \in P} (f_e c_e'(f_e))'$ is the altruistic part. $(f_e c_e'(f_e))'$ denotes the derivative with respect to f_e . Notice that we can rewrite $c_P^{(\beta)}(f) = \sum_{e \in P} c_e(f_e) + \beta \sum_{e \in P} f_e c_e'(f_e)$.

Definition 2.2 is similar to the definition of the valuation of a user with a time/money tradeoff of β in the case of network routing with tolls [8]. However, notice that unlike the standard model for tolls, the “edge toll” τ_e a user incurs in our model is *traffic-dependent*, namely $\tau_e := f_e c_e'(f_e)$.

We say that the users are *uniformly β -altruistic* if all users are β -altruistic. More generally, we allow for the case of arbitrary distributions of altruism among the users. In the

²The restriction to values $\beta_i \geq -1$ is justified in Section 3.

³While our definition is motivated mathematically, there is a “psychological” interpretation of the underlying choice: In order to behave (partially) altruistically, infinitesimally small users must give infinitesimally small weight to their own payoff, which is achieved implicitly by making the altruistic component the derivative of the social welfare.

general case, for each commodity i , we are given an arbitrary *altruism density function* ψ_i on the interval $[-1, 1]$. We only require that all these functions ψ_i be indeed distributions, i.e., forming a Borel measure of total measure 1. If the rate for commodity i is r_i , then the overall altruism density function is $\psi = \frac{1}{\sum_i r_i} \sum_i r_i \psi_i$. The *average altruism* of a distribution ψ is then $\int_{-1}^1 t\psi(t)dt$. An instance of the partially altruistic traffic routing problem is thus the quadruple $(G, r, c, (\psi_i))$. If there is a single commodity with distribution ψ , we write (G, r, c, ψ) , and if the altruism is uniform, we simplify further to (G, r, c, β) .

PROPOSITION 2.3. *Let (G, r, c, β) be an instance with uniform altruism $\beta \geq 0$. Then, the Nash flows are the optima of the convex program*

$$\begin{aligned} & \text{Minimize} && \sum_e \int_0^{f_e} c_e^{(\beta)}(t) dt \\ & \text{subject to} && f \text{ is a feasible multi-commodity flow for } (G, r) \end{aligned}$$

In particular, the instance (G, r, c, β) always possesses a Nash Equilibrium for $\beta \geq 0$.

The proof of this proposition is virtually identical to that of Proposition 2.6.1 from [33]. The proof there only used the fact that each agent was minimizing a sum of monotone increasing functions $\sum_e g_e(f_e)$ to conclude that the Nash Equilibrium was the flow minimizing the (convex) objective $\sum_e \int_0^{f_e} g_e(t)dt$. Thus, it applies equally to $g_e(t) := c_e^{(\beta)}(t)$.

The situation is not quite as straightforward for the case of non-uniform altruism distributions ψ , or for negative β . Even for two different values of altruism, there appears to be no natural convex programming formulation for Nash Equilibria. However, using a theorem of Mas-Collel [28], we can still prove the existence of Nash Equilibria.

THEOREM 2.4. *Each instance $(G, r, c, (\psi_i))$ has a Nash Equilibrium.*

Proof. Theorem 1 of Mas-Collel [28] proves that each game of infinitely many players has a Nash Equilibrium. A game is characterized by a distribution (Borel measure) over utility functions which are continuous in the action of the player, and the distribution of actions by the remaining players. It is easy to see that each player in the routing game has a utility function $-c_p^{(\beta)}(f)$ continuous in the choice of path P (trivially, since the space of all simple s_i - t_i paths is finite) and in the distribution of other players' strategies f (by continuity of each c_e). The utility for paths not connecting s_i to t_i is $-\infty$ (or an appropriately negative constant). The distribution of altruism values β implies a corresponding distribution over utility functions. Thus, the theorem of Mas-Collel implies the existence of Nash Equilibria for routing games. ■

The proof by Mas-Collel is inherently non-constructive; accordingly, Theorem 2.4 does not imply any algorithm for finding such equilibria. Since there always exists a Nash Equilibrium of instances $(G, r, c, (\psi_i))$, we can again define the Price of Anarchy (PoA), as $\rho(G, r, c, (\psi_i)) = C(f)/C(f^*)$, where f is a Nash flow for $(G, r, c, (\psi_i))$, and f^* a socially optimal flow for (G, r, c) .

2.2 Taxes and Stackelberg Strategies

Our definition of partial altruism naturally relates to two strategies that have been proposed in the literature for dealing with the selfishness of users: Pigou taxes and Stackelberg strategies.

The idea of taxes or tolls on edges is to charge users a fee for using an edge. The assumption is that money and latency can be measured on the same scale, and users will minimize the (weighted) sum of the two. It is well-known [29] that if the toll charged on each edge e equals the *marginal cost* to others ($f_e^* c_e'(f_e^*)$) at the optimum solution, then the Nash Equilibrium will minimize $C(f)$, i.e., be optimal. Our model of partial altruism can thus be interpreted as charging users a *traffic-dependent* constant fraction of the marginal tax, i.e., with respect to the *current* flow. When the altruism is not uniform, different users will be charged different taxes $\beta_i f_e c_e'(f_e)$ on edges. Our model can thus be considered as investigating the price of anarchy when different users have different tradeoffs between taxes and latency, but their tradeoff stays constant across different edges. Similar models were considered, e.g., in [11, 39]. Cole et al. [9, 8] also study optimization problems arising from non-uniform taxation in networks. However, their goal is to minimize the total tolls, subject to forcing the flow to optimal, whereas we study the price of anarchy given the taxation scheme of charging a (user-dependent) fraction of the marginal tax on each edge.

A different strategy for lowering the price of anarchy is available when a benevolent central authority controls a λ fraction of the total traffic. The central authority's goal is to route this fraction so as to minimize the total cost $C(f)$, subject to the fact that the remaining users will subsequently route their traffic selfishly. Algorithms for routing flows with this objective are called *Stackelberg strategies*, and the corresponding asymmetric games *Stackelberg games* (see, e.g., [32]).

When the altruism distribution has support $\{0, 1\}$, and the cumulative distribution function of ψ is the step function whose value at 0 is $1 - \lambda$, and whose value at 1 is 1, the altruistic users can be interpreted as a central authority, and their flow as a Stackelberg strategy with the corresponding price of anarchy. When the central authority controls a λ fraction of the traffic, then the average altruism is exactly λ , and thus, any bound on the price of anarchy for average altruism λ gives rise to the same bound for Stackelberg routing. Notice that the converse is not necessarily true: at the moment, it is not known if every optimal Stackelberg strategy gives rise to a Nash Equilibrium of the routing game with altruism support $\{0, 1\}$.

Such Stackelberg routing strategies have been studied extensively. In general, the price of anarchy can still be unbounded, even for single-commodity flows where a central authority controls a large constant fraction of the traffic [3]. For linear latency functions, Karakostas and Kolliopoulos [22] recently showed an upper bound of $(4 - X)/3$ on the Price of Anarchy (where $X = \frac{(1 - \sqrt{1 - \lambda})(3\sqrt{1 - \lambda} + 1)}{2\sqrt{1 - \lambda} + 1}$) for arbitrary networks and commodities in which a central authority controls a λ fraction of traffic. For arbitrary latency functions in series-parallel networks, Swamy [40] bounds the price of anarchy by $1 + 1/\lambda$. For parallel link networks with latency functions from a class \mathcal{C} with an upper bound $\rho(\mathcal{C})$ on the price of anarchy in Pigou examples, he shows an up-

per bound of $\lambda + (1 - \lambda)\rho(\mathcal{C})$. In the context of Stackelberg routing, a converse direction has been studied by Sharma and Williamson [38] and Kaporis and Spirakis [20]. They ask how much traffic needs to be controlled by a central authority to guarantee any improvement in average latency [38] (called *Stackelberg threshold*) or to guarantee optimality of the resulting Nash Equilibrium [20] (called *Price of Optimum*).

3. UNIFORM ALTRUISM

In this section, we focus on the model of uniformly altruistic users: each user is β -altruistic for $-1 \leq \beta \leq 1$. Thus, the perceived cost of an edge e to the user is $c_e^{(\beta)}(x) = (1 - \beta)c_e(x) + \beta \frac{d}{dx}(xc_e(x)) = c_e(x) + \beta xc_e'(x)$. (Notice that for $\beta = 0$, this coincides with selfishness; $\beta = 1$ corresponds to complete altruism, and $c_e^{(1)}(x)$ is exactly the marginal cost of e . For $\beta = -1$, the users are completely spiteful.) Our first result follows directly from the definitions of flows at Nash equilibrium and optimum, and gives a (tight) upper bound on the Price of Anarchy for arbitrary networks, commodities, and arbitrary semi-convex cost functions.⁴

PROPOSITION 3.1. *If all cost functions c_e are nondecreasing and semi-convex, then for all networks G and flow rates r , and any altruism level $\beta \in (0, 1]$,*

$$\rho(G, r, c, \beta) \leq 1/\beta.$$

Proof. Let \hat{f} be a Nash Equilibrium flow, minimizing the potential function $\Phi(f) = \sum_e \int_0^{f_e} c_e^{(\beta)}(t) dt$, the objective function of the convex program in Proposition 2.3. Also, let f^* the optimum flow, minimizing the total cost $C(f) = \sum_e \int_0^{f_e} (tc_e(t))' dt$. Simply from the definition of $c_e^{(\beta)}(t)$, it follows that for any flow f , we have $\Phi(f) \leq C(f) \leq \frac{1}{\beta}\Phi(f)$.

Applying the first inequality to f^* and the second to \hat{f} , and using the optimality of \hat{f} for Φ , we obtain $C(\hat{f}) \leq \frac{1}{\beta}\Phi(\hat{f}) \leq \frac{1}{\beta}\Phi(f^*) \leq \frac{1}{\beta}C(f^*)$. ■

More generally, we derive a result bounding the price of anarchy when all cost functions c_e are drawn from a given class of cost functions. Our characterization will be in terms of the *anarchy value* $\alpha^{(\beta)}(\mathcal{C})$ of a set \mathcal{C} of functions for β -altruistic users, which is defined as a generalization of the anarchy value of functions in [33].

DEFINITION 3.2. 1. *For any cost function c , the anarchy value $\alpha^{(\beta)}(c)$ of c for β -altruistic users is defined as*

$$\alpha^{(\beta)}(c) = \sup_{r, x \geq 0} \frac{r \cdot c(r)}{x \cdot c(x) + (r-x) \cdot c^{(\beta)}(r)},$$

where $0/0$ is defined to 1.

2. *For any class \mathcal{C} of cost functions, the anarchy value for β -altruistic users $\alpha^{(\beta)}(\mathcal{C})$ is $\sup_{c \in \mathcal{C}, c \neq 0} \alpha^{(\beta)}(c)$.*

The motivation for this definition of $\alpha^{(\beta)}(c)$ is that it captures the price of anarchy for uniformly β -altruistic users in a two-node two-link network, where one link has latency function c and the other has a worst-case constant. Indeed, we will prove this to be the case in Lemma 3.7 below. Notice that Lemma 3.7 immediately implies that $\alpha(\mathcal{C})$ is a lower

⁴We thank an anonymous reviewer for the simplified proof.

bound on the price of anarchy in the worst case when all edge latency functions are chosen from \mathcal{C} . Our main theorem in this section shows that it is also an upper bound for all networks and arbitrary commodities.

We are mostly interested in $\alpha^{(\beta)}(\mathcal{C})$ when it is finite. In particular, this suggests defining the *spite resistance* of \mathcal{C} as *the least altruistic behavior that \mathcal{C} could support*. Formally, $b_c = \inf\{\beta \mid \alpha^{(\beta)}(c) < \infty\}$, and $b_{\mathcal{C}} = \inf_{c \in \mathcal{C}} b_c$. It is not difficult to show that $b_c = -\inf_r \frac{c(r)}{rc'(r)}$, and that $\alpha^{(\beta)}(c) = \infty$ for $\beta \leq b_c$. Using L'Hôpital's rule, one sees that the monotonicity and convexity of c imply that $b_c \geq -\lim_{r \rightarrow \infty} \frac{c(r)}{rc'(r)} \geq -1$ for all c , which also motivates our earlier restriction to altruism values $\beta \geq -1$.

THEOREM 3.3. *Let \mathcal{C} be a set of cost functions, and (G, r, c) an instance with cost functions $c_e \in \mathcal{C}$. Then,*

$$\rho(G, r, c, \beta) \leq \alpha^{(\beta)}(\mathcal{C}).$$

Proof. Fix an instance (G, r, c) with cost functions $c_e \in \mathcal{C}$. Let f^* be an optimal flow and f a Nash flow for β -altruistic users. By rearranging Definition 3.2, we obtain the bound $x \cdot c_e(x) \geq \frac{r \cdot c_e(r)}{\alpha^{(\beta)}(\mathcal{C})} + (x-r) \cdot c_e^{(\beta)}(r)$ for any $x, r \geq 0$. Applying this bound to each edge e , with $x = f_e^*$ and $r = f_e$, we bound

$$\begin{aligned} C(f^*) &= \sum_{e \in E} f_e^* c_e(f_e^*) \\ &\geq \frac{1}{\alpha^{(\beta)}(\mathcal{C})} \cdot \sum_{e \in E} f_e c_e(f_e) + \sum_{e \in E} (f_e^* - f_e) \cdot c_e^{(\beta)}(f_e) \\ &= \frac{C(f)}{\alpha^{(\beta)}(\mathcal{C})} + \sum_{e \in E} (f_e^* - f_e) \cdot c_e^{(\beta)}(f_e). \end{aligned}$$

It remains to show that $\sum_e f_e^* \cdot c_e^{(\beta)}(f_e) \geq \sum_e f_e \cdot c_e^{(\beta)}(f_e)$. To this end, recall that f is a Nash flow for β -altruistic users if and only if it minimizes $c_P^{(\beta)}(f) f_P$ over all feasible flows \hat{f} . In particular, applying this variational inequality to f and f^* proves the desired inequality. ■

As a corollary of Theorem 3.3, we can obtain a tight bound in the case where the cost functions are polynomials of degree at most p with non-negative coefficients. We denote this class by \mathcal{C}_p .

THEOREM 3.4. *If (G, r, c) has cost functions in \mathcal{C}_p , then for any altruism value $\beta \in (-1/p, 1]$,*

$$\rho(G, r, c, \beta) \leq \left(\left(\frac{1+\beta p}{1+p} \right)^{1/p} \left(\frac{1+\beta p}{1+p} - 1 - \beta p \right) + 1 + \beta p \right)^{-1}.$$

Proof. First, notice that $b_{\mathcal{C}_p} = -1/p$. It can be easily verified that all subsequent calculations stay valid for $\beta > -1/p$, while for $\beta \leq -1/p$, the price of anarchy is unbounded.

As observed in [33], it suffices to focus only on polynomials $c(x) = ax^i$ with $x \leq p$. For any instance (G, r, c) with arbitrary polynomials can be equivalently transformed into one with only such monomials, by replacing each edge with cost function $c_e(x) = \sum_{i=0}^p a_i x^i$ by a directed path of $p+1$ edges, the i^{th} edge of which has cost function $\tilde{c}_{e,i}(x) = a_i x^i$. In order to compute the anarchy value $\alpha(c)$ of a nonzero polynomial function $c(x) = ax^i$, we use the equivalent characterization that

$$\alpha^{(\beta)}(c) = \sup_{r \geq 0} \left(\frac{\lambda c(\lambda r)}{c(r)} + (1 - \lambda) \left(1 + \frac{\beta r c'(r)}{c(r)} \right) \right)^{-1},$$

where $\lambda \in [0, 1]$ solves $c^{(1)}(\lambda r) = c^{(\beta)}(r)$, and $0/0$ is defined to 1. To prove this equivalent characterization, we first observe that

$$\frac{d}{d\lambda} (c(\lambda r) \lambda r + c^{(\beta)}(r)(r - \lambda r)) = c^{(1)}(\lambda r) r - c(r) r = 0,$$

so there is indeed a value of $\lambda \in [0, 1]$ solving $c^{(1)}(\lambda r) = c^{(\beta)}(r)$. By Lemma 3.7 below, $\alpha^{(\beta)}(c)$ is the price of anarchy in a two-node two-link network, one of whose links has the cost function $c(x)$, the other link having constant cost $c^{(\beta)}(r)$. Routing λr units of flow on the link with cost $c(x)$, and the rest on the link with cost $c^{(\beta)}(r)$, provides an optimal flow, while the Nash Equilibrium has all of its flow on the link with cost $c(x)$. Thus, the ratio of the cost of a Nash flow to that of an optimal flow is

$$= \frac{r \cdot c(r)}{c(\lambda r)\lambda r + (c(r) + \beta r c'(r))(r - \lambda r)} = \left(\frac{\lambda c(\lambda r)}{c(r)} + (1 - \lambda) \left(1 + \frac{\beta r c'(r)}{c(r)} \right) \right)^{-1}.$$

Solving for λ in the special case $c(x) = ax^i$, we obtain $\lambda = \left(\frac{1 + \beta i}{1 + i} \right)^{1/i}$, and thus $\frac{c(\lambda r)}{c(r)} = \frac{1 + \beta i}{1 + i}$ and $\frac{c'(r)}{c(r)} = \frac{i}{r}$. Then,

$$\alpha^{(\beta)}(c) = \left(\left(\frac{1 + \beta i}{1 + i} \right)^{1/i} \left(\frac{1 + \beta i}{1 + i} - 1 - \beta i \right) + 1 + \beta i \right)^{-1},$$

which is independent of a and increasing in i (by a derivative test). Hence, the largest $\alpha(c)$ is attained for $c = x^p$, giving

$$\alpha^{(\beta)}(C_p) = \left(\left(\frac{1 + \beta p}{1 + p} \right)^{1/p} \left(\frac{1 + \beta p}{1 + p} - 1 - \beta p \right) + 1 + \beta p \right)^{-1},$$

as claimed. \blacksquare

It is not difficult to verify that the previous bound converges to $\frac{1}{\beta}$ as $p \rightarrow \infty$; the worst case behavior is in fact attained with polynomials of high degree. However, for $p = 1$, Theorem 3.4 also allows us to obtain a tighter bound in the special case that all latency functions are linear.

COROLLARY 3.5. *If (G, r, c) has linear cost functions, then for any $\beta \in (-1, 1]$,*

$$\rho(G, r, c, \beta) \leq \frac{4}{3 + 2\beta - \beta^2}.$$

Notice that for any $\beta > 0$, this bound improves on the bound by Roughgarden and Tardos [36] of $4/3$ when all users are completely selfish. As the bound can also be shown to be tight, it thus characterizes exactly the gain by partial positive altruism with linear cost functions, and the spite resistance of linear cost functions. In particular, it shows that linear costs have the highest spite resistance among all classes of cost functions.

REMARK 3.6. *Our results in this section extend straightforwardly to general non-atomic congestion games (not necessarily network congestion games), so long as all cost functions are nondecreasing. (See, for instance, [7].) In a general congestion game, each player's strategy consists of a set of resources, and the cost of the strategy depends simply on the number of players using each resource. Thus, the perceived cost of a player's strategy S with altruism β is $c^{(\beta)}(S) = \sum_{e \in S} c_e^{(\beta)}(x_e) = \sum_{e \in S} c_e(x_e) + \beta x_e c'_e(x_e)$, where x_e is the total measure of players using resource e , and c_e is a nondecreasing function. With the same definitions of $\alpha^{(\beta)}(C)$, the proofs of the above proposition and theorems naturally carry over to this more general setting.*

Finally, we show that the bounds derived in Theorem 3.3 are indeed tight, even for two-node two-link networks:

LEMMA 3.7. *Consider a two-node two-link network with flow rate $r = 1$, and cost functions $c_1(x) = c(x)$ on the first*

link, and constant cost function $c_2(x) = c^{(\beta)}(r) = c(r) + \beta r c'(r)$ for the second link. Then, the price of anarchy of this instance is $\alpha^{(\beta)}(c)$.

Proof. It is easy to observe from the definition of c_2 that all β -altruistic users will end up using link 1, so that the total cost of the Nash Equilibrium is $c(1) = rc(r)$, while the socially optimum solution has total cost $\inf_{x \leq 1} (x \cdot c(x) + (r - x) \cdot c(r) + \beta(r - x)c'(r))$. Hence, the price of anarchy is exactly $\alpha(c)$. \blacksquare

By applying this characterization together with Theorem 3.4 and letting the degree of the polynomial go to ∞ , we obtain instances (G, r, c) whose price of anarchy approaches $1/\beta$ arbitrarily closely. Similarly, by choosing $p = 1$, we obtain that the bound in Corollary 3.5 is tight.

4. NON-UNIFORM ALTRUISM

In this section, we extend our results to the more general and realistic case where different users can have different altruism levels. In the most general case, we are given a distribution ψ of altruism. The existence of Nash Equilibria in this model was shown non-constructively as Theorem 2.4. Even for a single commodity and an altruism distribution with support $\{0, 1\}$, and arbitrarily large constant $\bar{\beta}$, a recent result on Stackelberg routing due to Bonifaci et al. [3] shows that the price of anarchy can become unbounded.

Thus, we focus here on the case of single-commodity traffic in *parallel link networks*. Parallel link networks have been studied by Roughgarden [32]; among others, they naturally model the assignment of infinitesimally small jobs to machines with load-dependent latencies. Formally, a parallel link network has two nodes s, t , and m parallel edges e_1, \dots, e_m from s to t . Our main theorem in this section gives a (tight) upper bound on the Price of Anarchy in the presence of partial altruism for single commodity parallel link networks and arbitrary (convex) cost functions.

THEOREM 4.1. *If all cost functions c_e are convex and non-decreasing, then for parallel link networks G and flow rates r , and any overall altruism density function ψ with nonnegative support and average altruism $\bar{\beta}$,*

$$\rho(G, r, c, \psi) \leq 1/\bar{\beta}.$$

We will prove Theorem 4.1 as a corollary of the following more general result, bounding the price of anarchy in terms of the set of functions permissible as edge latencies.

THEOREM 4.2. *If all cost functions c_e are convex and non-decreasing, then for parallel link networks G and flow rates r , and any overall altruism density function ψ with nonnegative support,*

$$\rho(G, r, c, \psi) \leq \left(\int_0^1 \psi(t) \frac{1}{\alpha^{(t)}(C)} dt \right)^{-1}.$$

Proof. Let f denote the flow at Nash Equilibrium. We first show that without loss of generality, we can assume that each link e contains only one type of users (i.e., if users have different altruism values β, β' , then they do not share a link) and that the support of ψ is finite. To see this, assume that f has users of altruism values $\beta < \beta'$ sharing an edge e . Now replace all users on e with altruism β by users with altruism β' . f must still be a flow at Nash Equilibrium for

the new instance (because β' -altruistic users are on link e in Nash Equilibrium). By repeating this process, we eventually obtain an instance with altruism density ψ' which stochastically dominates ψ and has finite support. For this new ψ' , the bound on the price of anarchy for f provided by the right-hand side of Theorem 4.2 can only be smaller, giving us an even better bound than required. Thus, we can from now on focus on the case described above.

Let $0 \leq \beta_1 < \beta_2 < \dots < \beta_k \leq 1$ be the (finite) support of ψ , where the rate of β_i -altruistic users is r_i (so $\sum_{i=1}^k r_i = r$). We need to show that for all flows g of rate r (in particular the optimum flow), we have

$$C(g) \geq \left(\sum_{i=1}^k \frac{r_i}{r} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})} \right) \cdot C(f), \quad (1)$$

which we will do by induction on k . The base case $k = 0$ is of course trivial.

For the inductive step, let f be a Nash Equilibrium flow, and g any flow of rate r . For each i , let E_i be the set of edges with positive flow of β_i -altruistic users under f . Notice that by our assumption, the sets E_i are disjoint. For any set E' of edges, let $f(E') = \sum_{e \in E'} f_e$ (similarly, $g(E')$) denote the total flow on E' . Let $E' := E \setminus E_1$ denote the set of all edges not used by β_1 -altruistic users.

Intuitively, because the more altruistic users prefer the edge set E' over E_1 , we would expect a “good” flow g to do the same. Indeed, we first show that the latency under f on all edges in E_1 is no larger than in E' , while the derivative is no larger in E' than in E_1 . Let $e \in E_1, e' \in E_j, j > 1$ be arbitrary links with positive flow f . Thus, all users on e have altruism β_1 , while all users on e' have altruism $\beta_j > \beta_1$. Because f is at Nash Equilibrium,

$$c_e(f_e) + \beta_1 f_e c'_e(f_e) \leq c_{e'}(f_{e'}) + \beta_1 f_{e'} c'_{e'}(f_{e'}), \quad (2)$$

$$c_e(f_e) + \beta_j f_e c'_e(f_e) \geq c_{e'}(f_{e'}) + \beta_j f_{e'} c'_{e'}(f_{e'}). \quad (3)$$

Combining appropriately scaled versions of Inequality (2) and Inequality (3) gives us that

$$c_e(f_e) \leq c_{e'}(f_{e'}), \quad (4)$$

$$\begin{aligned} & (1 - \xi)c_e(f_e) + (\beta_j - \xi\beta_1)f_e c'_e(f_e) \\ & \geq (1 - \xi)c_{e'}(f_{e'}) + (\beta_j - \xi\beta_1)f_{e'} c'_{e'}(f_{e'}), \end{aligned} \quad (5)$$

where ξ for $0 \leq \xi \leq 1$ is a scalar, which we can set later.

Our high-level strategy will be to bound the Nash Equilibrium flow on E' against a restriction g' of g of rate $r - r_1$ on E' by induction, and use a comparison argument for the flow on E_1 . We will construct a flow h of rate r_1 whose cost is cheaper than a component of g of the same rate, and which is optimal for modified “residual” edge costs. We can thus compare it against the flow f on E_1 using Theorem 3.3.

Define f' to be the restriction of f to the edge set E' , i.e., $f'_e = f_e$ for $e \in E'$, and $f'_e = 0$ for $e \in E_1$. Thus, f' is a flow of rate $r' := r - r_1$. Define the modified cost function $\tilde{c}_e(x) := c_e(f'_e + x) + \beta_1 f'_e c'_e(f'_e)$ for all edges e . Thus, $\tilde{c}_e(x)$ is the cost incurred by flow on e if f'_e is unalterable, but not considered part of the actual flow, plus a suitable constant term to “mimic” the altruistic component. This definition of $\tilde{c}_e(x)$ implies that the perceived cost of edge e to β_1 -altruistic users is $\tilde{c}_e^{(\beta_1)}(x) = c_e(f'_e + x) + \beta_1 x c'_e(f'_e + x) + \beta_1 f'_e c'_e(f'_e)$. Thus, for $e \in E'$, we have that $\tilde{c}_e^{(\beta_1)}(x) \geq c^{(\beta_1)}(f'_e)$ for all $x \geq 0$, while for $e \in E_1$, because $f'_e = 0$, $\tilde{c}_e^{(\beta_1)}(x) = c^{(\beta_1)}(x + f_e)$.

In particular, this implies that the β_1 -altruistic users are at Nash Equilibrium with respect to the modified cost functions $\tilde{c}_e(x)$. Hence, by Theorem 3.3, and because $\tilde{c}_e(x) = c_e(x)$ for all $e \in E_1$, we get $C(f - f') = \tilde{C}(f - f') \leq \alpha^{(\beta_1)}(\mathcal{C}) \cdot \tilde{C}(\tilde{f})$ where \tilde{f} is an optimum flow of rate r_1 with respect to the modified edge cost functions \tilde{c}_e .

In order to compare f' against the part of g on the edge set E' , it will be useful to assume that $g(E') \geq f(E')$. We will show next that we can make this assumption w.l.o.g. For assume that it did not hold. Then, let $e \in E_1, e' \in E'$ be edges with $g_e > f_e > 0$ and $g_{e'} < f_{e'}$. (The existence of e, e' follows from the assumption $g(E') < f(E')$). By the bound on the derivatives in Inequality (5), and using the convexity of the edge latency functions, we show that $(g_e c_e(g_e))' \geq (f_e c_e(f_e))' \geq (f_{e'} c_{e'}(f_{e'}))' \geq (g_{e'} c_{e'}(g_{e'}))'$ in the following. The first and last inequalities hold simply by the semi-convexity of c_e and $c_{e'}$, and the second equality is obtained by setting $\xi = \frac{1 - \beta_j}{1 - \beta_1}$ to get $1 - \xi = \beta_j - \xi\beta_1$ so $(f_e c_e(f_e))' = c_e(f_e) + f_e c'_e(f_e) \geq c_{e'}(f_{e'}) + f_{e'} c'_{e'}(f_{e'}) = (f_{e'} c_{e'}(f_{e'}))'$. Thus, g can be made cheaper by moving some of its flow from e to e' . By repeating this process, we can thus assume that $g(E') \geq f(E')$.

Let γ be such that $C(f - f') = \gamma C(f)$. Because f' and $f - f'$ use disjoint edge sets, we get $C(f') = (1 - \gamma)C(f)$. (Notice that the assumption of disjoint edge sets is indeed crucial here. Due to the non-constant cost of edges, in general, it does not hold that $C(f) + C(f') = C(f + f')$.)

By Lemma 4.3 below, we can decompose $g = h + g'$, where g' is a flow of rate r' entirely on E' , and h is a flow of rate r_1 satisfying the property (7), namely $\tilde{C}(\tilde{f}) \leq \sum_e h_e c_e(g_e) + \sum_e g'_e (c_e(g_e) - c_e(g'_e))$. We can thus apply induction on the flows f' and g' of rate r' on the modified graph with edge set E' . Notice that while f' may not be an Equilibrium flow on E' , it is indeed an Equilibrium flow on E' . Thus, we obtain that

$$\begin{aligned} C(g) &= C(g') + \sum_e h_e c_e(g_e) + \sum_e g'_e (c_e(g_e) - c_e(g'_e)) \\ &\geq \left(\sum_{i=2}^k \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})} \right) \cdot C(f') + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} C(f - f') \quad (6) \\ &= \left(\sum_{i=2}^k \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})} \right) \cdot (1 - \gamma) + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} \cdot \gamma \cdot C(f). \end{aligned}$$

We next show that $\gamma \leq \frac{r_1}{r}$. By Inequality (4), every user on E_1 incurs lower delay than every user on E_j , and consequently on E' . Thus, the average delay $\frac{1}{r_1} C(f - f')$ of users on E_1 is at most the average delay $\frac{1}{r} C(f)$ of all users, so $C(f - f') \leq \frac{r_1}{r} C(f)$.

The lower bound (6) is a convex combination of the non-negative terms $\sum_{i=2}^k \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}$ and $\frac{1}{\alpha^{(\beta_1)}(\mathcal{C})}$, with coefficients $(1 - \gamma)$ and γ . The anarchy value $\alpha^{(\beta)}(\mathcal{C})$ is a monotone non-increasing function of β , so the weighted average reciprocal anarchy value for altruism levels β_2, \dots, β_k is at least the reciprocal for β_1 . Thus, the convex combination is minimized when the coefficient γ of the smaller term $\frac{1}{\alpha^{(\beta_1)}(\mathcal{C})}$ is as large as possible, i.e., when $\gamma = r_1/r$. Substituting this bound,

$$\begin{aligned} C(g) &\geq \left(\sum_{i=2}^k \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})} \right) \cdot \frac{r'}{r} + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} \cdot \frac{r_1}{r} \cdot C(f) \\ &= \left(\sum_{i=1}^k \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})} \right) \cdot C(f), \end{aligned}$$

completing the inductive step, and thus the proof. \blacksquare

LEMMA 4.3. *Let f' be a flow of rate r' using only edges from E' , and define $\tilde{c}_e(x) := c_e(f'_e + x) + \beta_1 f'_e c'_e(f'_e)$. Let g be any flow of rate $r = r' + r_1$, with $g(E') \geq r'$. Let \tilde{f} be the optimum flow of rate r_1 with respect to edge costs \tilde{c}_e . Then, g can be decomposed as $g = h + g'$, where g' is a flow of rate r' on E' , satisfying*

$$\tilde{C}(\tilde{f}) \leq \sum_e h_e c_e(g_e) + \sum_e g'_e (c_e(g_e) - c_e(g'_e)). \quad (7)$$

Proof. Let $\Delta := g(E') - r' \geq 0$ be the amount of “excess flow” that g sends on E' , compared to f' . We begin by setting $h_e = g_e$ for all edges $e \in E_1$, giving us a flow of rate $r_1 - \Delta$. So we need to add Δ more units of flow to h . Let $E'' := \{e \in E' \mid g_e \geq f'_e\}$ be the set of edges in E' on which g sends more flow than f' . Thus, we have that $g(E'') - f'(E'') \geq g(E') - f'(E') = \Delta$. In particular, we can define a flow h of total rate Δ on E'' , such that $h_e \leq g_e - f'_e$ for all $e \in E''$. For all other edges e , we set $h_e = 0$, and thus obtain a flow h of rate r_1 , such that $h_e \leq g_e$ for all edges e . We then have that

$$\begin{aligned} \sum_e h_e c_e(g_e) &= \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(g_e) \\ &\geq \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(f'_e + h_e), \end{aligned}$$

where the inequality follows from the monotonicity of the latencies c_e . Next, because $g'_e \geq f'_e$ for all $e \in E'$, and the latency functions are convex, $\frac{c_e(g_e) - c_e(g'_e)}{h_e} \geq c'_e(f'_e)$ for all $e \in E''$ with $h_e > 0$. Combining this bound with the fact that $\beta_1 \leq 1$, we obtain that

$$\begin{aligned} \sum_e g'_e (c_e(g_e) - c_e(g'_e)) &\geq \sum_{e \in E''} g'_e (c_e(g_e) - c_e(g'_e)) \\ &\geq \sum_{e \in E''} f'_e \beta_1 h_e c'_e(f'_e). \end{aligned}$$

Summing the previous two inequalities now gives us

$$\begin{aligned} &\sum_e h_e c_e(g_e) + \sum_e g'_e (c_e(g_e) - c_e(g'_e)) \\ &\geq \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(f'_e + h_e) \\ &\quad + \sum_{e \in E''} h_e \beta_1 f'_e c'_e(f'_e) \\ &= \sum_e h_e \tilde{c}_e(h_e) \\ &\geq \tilde{C}(\tilde{f}) \end{aligned}$$

where the final inequality follows from the optimality of \tilde{f} with respect to the cost functions \tilde{c}_e . \blacksquare

Proof of Theorem 4.1. If \mathcal{C} is specifically the set of all increasing semi-convex functions, Proposition 3.1 implies that $\frac{1}{\alpha^{(t)}(\mathcal{C})} \geq t$. Substituting this bound into the integral gives us that

$$\rho(G, r, c, \psi) \leq \left(\int_0^1 \psi(t) t dt \right)^{-1} = 1/\bar{\beta}. \quad \blacksquare$$

It would of course be desirable to extend Theorems 4.1 and 4.2 to distributions including negative support. However, such an extension is in general not possible. One can construct scenarios in which almost all of the latency is incurred by a small fraction of spiteful users who together congest a link with very steep increase. At the same time, all altruistic users use links with very small constant latency. Then, the PoA is much larger than 1, while the bounds of both theorems would require it to be close to 1.

An immediate corollary of Theorem 4.1 can be obtained by choosing the distribution with a rate of λ users being completely altruistic, and $1 - \lambda$ users being completely selfish.

Since $\bar{\beta} = \lambda$ for this distribution, Theorem 4.1 immediately implies

COROLLARY 4.4. *In parallel link networks, the price of anarchy under Stackelberg routing with a λ -fraction of traffic being controlled by a central authority is at most $1/\lambda$.*

This result was of course already proved constructively (and giving efficient algorithms) by Roughgarden [32]; nevertheless, it is interesting that it follows directly from our general result. More generally, by using the same distribution with support $\{0, 1\}$ in Theorem 4.2, we obtain the following corollary:

COROLLARY 4.5. *In parallel link networks, the price of anarchy under Stackelberg routing with a λ -fraction of traffic controlled by a central authority is at most $(\frac{1-\lambda}{\alpha(\mathcal{C})} + \lambda)^{-1}$.*

Notice that Corollary 4.5 improves (albeit in a non-constructive way) a recent result of Swamy [40] for Stackelberg routing: We bound the PoA under Stackelberg routing by the weighted harmonic mean of the PoA for selfish and altruistic users, whereas Swamy’s bounds give the arithmetic mean. It is known that the harmonic mean is always bounded above by the arithmetic mean. We can also show that the case of Stackelberg routing is in fact the worst case for the bound of Theorem 4.2, in the sense that the right-hand side is maximized. While the bound of Theorem 4.2 will in general not be tight, this nevertheless gives rise to the philosophical interpretation that, conditioned on a given average altruism level $\bar{\beta}$, the scenario in which completely altruistic users or a central authority compensate for completely selfish users is the worst case, while uniform altruism through the population is the best case.

PROPOSITION 4.6. *Conditioned on the mean of ψ being any given $\bar{\beta}$, the quantity $\left(\int_0^1 \psi(t) \frac{1}{\alpha^{(t)}(\mathcal{C})} dt \right)^{-1}$ is maximized when ψ has point mass of $\bar{\beta}$ on 1 and $1 - \bar{\beta}$ on 0. It is minimized when ψ has a point mass of 1 on $\bar{\beta}$.*

Proof. We will show that $\frac{1}{\alpha^{(\beta)}(\mathcal{C})}$ is concave as a function of β . Both results then follow readily from Jensen’s Inequality. To prove concavity, let $p_1, p_2 \geq 0$ satisfy $p_1 + p_2 = 1$. For any cost function $c \in \mathcal{C}$, Definition 3.2 thus gives us

$$\begin{aligned} &\frac{1}{\alpha^{(p_1 \beta_1 + p_2 \beta_2)}(\mathcal{C})} \\ &= \inf_{\lambda} \frac{\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)(p_1 \beta_1 + p_2 \beta_2)c'(r)}{c(r)} \\ &= \inf_{\lambda} \frac{p_1(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_1 c'(r)) + p_2(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_2 c'(r))}{c(r)} \\ &\geq \inf_{\lambda} \frac{p_1(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_1 c'(r)) + p_2(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_2 c'(r))}{c(r)} \\ &= p_1 \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} + p_2 \frac{1}{\alpha^{(\beta_2)}(\mathcal{C})}. \end{aligned}$$

Finally, we take an infimum over all $c \in \mathcal{C}$ on both sides to complete the proof of concavity. \blacksquare

5. CONCLUSIONS

In this paper, we proposed a simple model of altruism and spite in traffic routing, where users’ utilities or perceived costs are linear combinations of their own latency and the increase they cause in other users’ latencies. We proved a $1/\beta$

bound on the price of anarchy even for worst-case networks, latency functions, and commodities, under the assumption that all users are (at least) β -altruistic, and $\beta > 0$. We extended this result to non-uniform altruism distributions for single-commodity flows in parallel link networks. Among others, this result recovers and improves recent bounds on Stackelberg routing by Roughgarden and by Swamy.

Our work suggests many interesting directions for further research. First, the results should be generalized to (more general) network topologies instead of parallel links. Notice that any such result would immediately imply corresponding bounds on Stackelberg routing, so the lower bound of Bonifaci et al. [3] precludes an extension to arbitrary single-commodity flows. However, an extension to *series-parallel graphs* seems plausible at this point.

While we proved the existence of Nash Equilibria for all routing games with non-atomic users, regardless of the distributions of altruism, the proof is non-constructive. The work of Roughgarden [32] implies that finding the *best* Stackelberg strategy is NP-complete. However, it would be interesting whether Stackelberg strategies meeting our bound can always be found efficiently. Alternatively, in light of recent results proving that finding Nash Equilibria is PPAD-complete [6], it may be possible that finding Nash Equilibria for traffic routing games with two (or more) altruism values is also PPAD-complete.

Finally, the study of partially altruistic and spiteful behavior can and should be extended beyond the realm of traffic routing. Several natural games studied in the CS community are known to have unbounded or large PoA under the assumption of selfish agents. It would be interesting to observe if the introduction of finite amounts of altruism (uniformly or not) into such games will also lead to a constant or otherwise reduced PoA. This holds in particular because partial altruism appears to be a more natural model of actual user behavior.

Acknowledgments

We would like to thank Tim Roughgarden and Pierre-Olivier Weill for useful discussions and references, and several anonymous referees for detailed and useful comments.

6. REFERENCES

- [1] M. Babaioff, R. Kleinberg, and C. Papadimitriou. Congestion games with malicious players. In *Proc. 9th ACM Conf. on Electronic Commerce*, 2007.
- [2] M. Beckmann, C. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [3] V. Bonifaci, T. Harks, and G. Schäfer. The impact of stackelberg routing in general networks. Technical Report COGA Preprint 020-2007, TU Berlin, 2007.
- [4] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proc. 20th Intl. Joint Conf. on Artificial Intelligence*, pages 1207–1214, 2007.
- [5] F. Brandt and G. Weiss. Antisocial agents and vickrey auctions. In *Proc. 8th Workshop on Agent Theories, Architectures and Languages*, pages 335–347, 2001.
- [6] X. Chen and X. Deng. Settling the complexity of two-player nash-equilibrium. In *Proc. 47th IEEE Symp. on Foundations of Computer Science*, pages 261–270, 2006.
- [7] S. Chien and A. Sinclair. Convergence to approximate nash equilibria in congestion games. In *Proc. 18th ACM Symp. on Discrete Algorithms*, 2007.
- [8] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 35th ACM Symp. on Theory of Computing*, pages 521–530, 2003.
- [9] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72:444–467, 2006.
- [10] R. Cominetti, J. Correa, and N. Stier-Moses. Network games with atomic players. In *Proc. 33rd Intl. Colloq. on Automata, Languages and Programming*, pages 525–536, 2006.
- [11] S. Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transportation Science*, 6:73–87, 1972.
- [12] E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114:817–868, 1999.
- [13] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. In *Proc. ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, 2004.
- [14] L. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348:217–225, 2005.
- [15] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 277–285, 2004.
- [16] N. Garg, D. Grosu, and V. Chaudhary. An antisocial strategy for scheduling mechanisms. In *Proc. 19th IEEE Intl. Parallel and Distributed Processing Symp.*, 2005.
- [17] H. Gintis, S. Bowles, R. Boyd, and E. Fehr. *Moral Sentiments and Material Interests : The Foundations of Cooperation in Economic Life*. MIT Press, 2005.
- [18] A. Hayrapetyan, E. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proc. 37th ACM Symp. on Theory of Computing*, pages 89–98, 2006.
- [19] O. Jahn, R. Möhring, A. Schulz, and N. Stier-Moses. System-optimal routing of traffic flows with user constraints in networks with congestion. *Operations Research*, 53:600–616, 2006.
- [20] A. Kaporis and P. Spirakis. The price of optimum in stackelberg games on arbitrary single commodity networks and latency functions. In *Proc. 19th ACM Symp. on Parallel Algorithms and Architectures*, pages 19–28, 2006.
- [21] G. Karakostas and S. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous users. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 268–276, 2004.
- [22] G. Karakostas and S. Kolliopoulos. Stackelberg strategies for selfish routing in general multicommodity networks. Technical Report 2006/08, McMaster University, 2006.
- [23] E. Koutsoupias and C. Papadimitriou. Worst-case

- equilibria. In *Proc. 17th Annual Symp. on Theoretical Aspects of Computer Science*. Springer, 1999.
- [24] J. Ledyard. Public goods: A survey of experimental research. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*, pages 111–194. Princeton University Press, 1997.
- [25] D. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.
- [26] L. Liang and Q. Qi. Cooperative or vindictive: Bidding strategies in sponsored search auctions. In *Proc. 3rd Workshop on Internet and Network Economics (WINE)*, pages 167–178, 2007.
- [27] H. Lin, T. Roughgarden, E. Tardos, and A. Walkover. Braess’s paradox, fibonacci numbers, and exponential inapproximability. In *Proc. 32nd Intl. Colloq. on Automata, Languages and Programming*, pages 497–512, 2005.
- [28] A. Mas-Colell. On a theorem of Schmeidler. *J. of Mathematical Economics*, 13:201–206, 1984.
- [29] A. Pigou. *The Economics of Welfare*. Macmillan, 1920.
- [30] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67:341–364, 2003.
- [31] T. Roughgarden. The maximum latency of selfish routing. In *Proc. 15th ACM Symp. on Discrete Algorithms*, pages 973–974, 2004.
- [32] T. Roughgarden. Stackelberg scheduling strategies. *SIAM J. on Computing*, 33:332–350, 2004.
- [33] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [34] T. Roughgarden. Selfish routing with atomic players. In *Proc. 16th ACM Symp. on Discrete Algorithms*, pages 1184–1185, 2005.
- [35] T. Roughgarden. On the severity of braess’s paradox: Designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72:922–953, 2006.
- [36] T. Roughgarden and E. Tardos. How bad is selfish routing? In *Proc. 41st IEEE Symp. on Foundations of Computer Science*, 2000.
- [37] T. Schelling. *Micromotives and Macrobehavior*. Norton, 1978.
- [38] Y. Sharma and D. Williamson. Stackelberg thresholds in network routing games or the value of altruism. In *Proc. 9th ACM Conf. on Electronic Commerce*, 2007.
- [39] M. Smith. The marginal cost taxation of a transportation network. *Transportation Research, Part B*, 13:237–242, 1979.
- [40] C. Swamy. The effectiveness of stackelberg strategies and tolls for network congestion games. In *Proc. 18th ACM Symp. on Discrete Algorithms*, 2007.
- [41] I. Vetsikas and N. Jennings. Outperforming the competition in multi-unit sealed bid auctions. In *Proc. 6th Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 702–709, 2007.
- [42] J. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.