CCIT Report #461
January 2004

The Price of Anarchy in Networks with Bottleneck Objectives

Ron Banner and Ariel Orda
Department of Electrical Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
{banner@tx, ariel@ee}.technion.ac.il
Total number of pages (including Appendix): 13

Abstract—**We study the price of anarchy when performance is dictated by the worst (bottleneck) element. We are given a network, finitely many users, each associated with a positive flow demand, and a load-dependent performance function for each network element; the network objective is to route traffic such that the performance of the worst (i.e., bottleneck) element in the network is optimized. In the absence of regulation by some central authority, we assume that each user routes its traffic selfishly i.e., through paths that optimize the *user's* performance, in terms of bottleneck elements. We prove the existence of a Nash equilibrium, considering two routing scenarios, namely when a user can split its traffic over more than one path and when it cannot. Then, we quantify the price of anarchy in both scenarios. Specifically, we show that, when users are allowed to split their traffic, *anarchy comes at no price*. On the other hand, we show that, if each user is limited to a single path, the price of anarchy is *unbounded*. Finally, we turn to consider the case where the network objective is *additive* i.e., minimizing the sum of all link performance functions, while users still optimize the performance of their bottleneck elements. For this case, and for users that can split their traffic, we show that the price of anarchy is at most the number of network links. We then delineate a possible application of this result for the case where both the network and the users consider *additive* objectives.**

Keywords- **bottleneck & additive metrics, Nash equilibrium, price of anarchy/coordination factor, selfish routing, unregulated traffic.**

I. Introduction

Traditional computer networks were designed and operated with systemwide optimization in mind. Accordingly, the actions of the network users were determined so as to optimize the overall network performance. Consequently, users would often find themselves sacrificing some of their own performance for the sake of the entire network. In recent years it has been recognized that systemwide optimization may be an impractical paradigm for the control of modern networking configurations [1],[12],[16],[27]. Indeed, control decisions in large scale networks are often made by each user independently, according to its own individual performance objectives. Such networks are henceforth called *noncooperative*, and Game Theory [17] provides the systematic framework to study and understand their behavior.

Game theoretic models have been employed in the context of flow control [1],[12],[27] routing [16],[20],[21] and bandwidth allocation [14]. These studies mainly investigated, the structure of network operating points i.e., the Nash equilibria of the respective games. Such equilibria are inherently inefficient [11] and, in general, exhibit suboptimal network performance. In order to understand this phenomenon, an investigation of the performance ratio between the worst possible Nash equilibrium and the social (i.e., overall) optimum was initiated in [13]. This ratio, termed the *price of anarchy* (also: *coordination factor*), was first investigated by [10],[13],[15] for routing problems in which a set of users send traffic along a set of parallel links with linear cost functions. A more general framework for general topologies was later considered in [19],[20],[21]; in those studies, the cost of each link was a load-dependent latency function, and each network user chose a minimum-latency path while controlling a negligible fraction of the overall traffic; the network objective was to minimize total latency.

The above studies solely focused on *additive metrics* i.e., the case where performance is determined by the *sum* of link cost functions. Another fundamental case is that of *bottleneck metrics*, in which network performance is determined by the *worst* component (link) in the network. Bottleneck objectives (also known as Max-Min or Min-Max objectives) are of major practical importance. For example, a commonly used objective for traffic engineering is to minimize the utilization of the most utilized link in the network, in order to move away traffic from congested hot spots to less utilized parts of the network [2],[26]. Another example is when the performance goal is to maximize the ability to accommodate momentary traffic bursts by maximizing the minimum residual capacity (or the "common scaling factor" [3] i.e., a common factor that can scale up the traffic demands of all users without violating the capacity constraints). Another commonly used bottleneck objective appears in the context of wireless networks, where each node has a limited battery (i.e., transmission energy). The objective is then to maximize lifetime i.e., the time until the *first* battery in the network drains out [7],[8],[9]. Other scenarios where bottleneck metrics are considered occur, for example, in frameworks that consider fairness objectives, where the performance of the "poorest" user in the network needs to be maximized [4].

In spite of the major importance of bottleneck metrics, they have not been considered in the context of noncooperative networks. Accordingly, in this study we investigate non-cooperative network games with bottleneck objectives. More specifically, we assume that each user routes its traffic selfishly so as to optimize the performance of *its* bottleneck element, given the traffic of all other users. Under this assumption, we can view the network users as *N* independent agents participating in a noncooperative game, and expect the routes chosen by the users to form a Nash equilibrium. Our main goal is to quantify the cost of selfish routing with respect to bottleneck metrics. Throughout the paper, we assume that there are *finitely many agents*, each controlling a *non-negligible amount* of flow. We then consider two routing scenarios, namely, when a user can split its traffic over more than one path (henceforth, *splittable flow*) and when it cannot (henceforth, *unsplittable flow*). We show that, for performance functions that are convex increasing and continuous, there is a (not necessarily unique) Nash equilibrium for both splittable and unsplittable flows. We then study the price(s) of anarchy of splittable and unsplittable flows. It is well known that the price of anarchy is usually more than 1, and, moreover, it is often unbounded. However, we show that, for convex, increasing and continuous performance functions, *anarchy comes at no price* for splittable flows[1]. On the other hand, we show that, for unsplittable flows, the price of anarchy may be arbitrarily large.

Next, we consider the degradation in the *additive* network objective of "total cost" (i.e., minimizing the sum of all link performance functions), when selfish users still consider *bottleneck* objectives. This is the case when the network is interested in *average* performance across the network rather than the performance of the worst component. For example, in the context of wireless networks, users are often interested in maximizing the lifetime of their connections, whereas the network may be interested in the global objective of minimizing the total network power consumption [24]. As before, we study this scenario for splittable and unsplittable flows. For the splittable case, we show that the performance is at most *M*, where *M* is the number of links in the network. For the unsplittable case, we show that the price of anarchy remains unbounded. Therefore, for splittable flows, we conclude that, if there exists a way to enforce users to route their traffic with respect to bottleneck metrics (even if they essentially consider additive metrics), it may be possible to restrict the price of anarchy to *M*. Accordingly, we suggest a possible enforcement mechanism.

The rest of this paper is organized as follows. In section 2, we formulate the model and terminology. In section 3, we establish the existence (and non-uniqueness) of splittable and unsplittable flows at Nash equilibrium. In section 4, we quantify the price of anarchy in the two routing scenarios and for the case where both the users and the network consider bottleneck objectives. In section 5, we quantify the price of anarchy of the corresponding routing scenarios for the case where the network considers an additive objective; we then describe a possible scheme for bounding the price of anarchy when both the network and the users consider additive objectives. Finally, section 6 summarizes the results and discusses directions for further research.

II. Model and Terminology

This section formulates the model and required terminology.

A. The Model

A *network* is represented by a directed graph $G(V,E)$, where V is the set of nodes and E is the set of links. Let $N=|V|$ and $M=|E|$. A (simple) *path* is a finite sequence of distinct nodes $p=(v_0,v_1,\cdots,v_h)$, such that, for $0\le n\le h-1, (v_n,v_{n+1})\in E$. Let $P^{(i,j)}$ denote the collection of all paths from the source *i* to the destination *j* and $P\triangleq\bigcup_{(i,j)\in V\times V}P^{(i,j)}$.

We consider a finite set of users *U*, where each user $u\in U$ is associated with a positive demand γ_u and a pair of source-destination nodes (s_u,t_u). A user $u\in U$ needs to send γ_u units of flow from s_u to t_u. For each $u\in U$, the *connection* of user *u* is the collection of all paths that user *u* employs.

[1] The convexity is needed only to ensure the existence of a Nash equilibrium. Whenever there exists a Nash equilibrium, the price of anarchy remains 1 for performance functions that are just continuous and increasing.

Given are a network $G(V,E)$, and a set of users U. A *splittable flow vector* $f \triangleq \{f_p^u\}$ is a real-valued function $f: P\times U \to \mathbb{R}^+ \cup \{0\}$ that satisfies the flow demands of all users in U i.e., $\sum_{p\in P^{(s_u,t_u)}} f_p^u = \gamma_u$ for each $u\in U$. For a fixed flow vector f and a path $p\in P$, the total flow that is carried over p is $f_p \triangleq \sum_{u\in U} f_p^u$. Next, for a fixed flow vector f, a user $u\in U$ and a link $e\in E$, the total flow that u transfers through e is defined as $f_e^u \triangleq \sum_{p|e\in p} f_p^u$; finally, for a fixed flow vector f and a link $e\in E$, the total flow that is carried by e is $f_e \triangleq \sum_{u\in U} f_e^u$.

Note that, in any splittable flow vector, the flow demands can be split along any number of paths. We shall also consider the case where each user is limited to route its flow over a single path. Accordingly, we say that $f = \{f_p^u\}$ is an *unsplittable flow vector* if f is a (splittable) flow vector that satisfies $f_p^u \in \{0,\gamma_u\}$ for each $u\in U$ and $p\in P$.

We associate with each link $e\in E$ a *performance function* $q_e(\cdot)$ that depends on the total flow f_e carried over e. We assume that, for all $e\in E$, $q_e(f_e)$ is convex, continuous and increasing in f_e. Given link performance functions $\{q_e(f_e)\}$, the network objective may be either "bottleneck" or "additive". A network with a "bottleneck" objective aims at minimizing the *network bottleneck* $B(f)$, defined as the worst performance among all links i.e., $B(f) \triangleq \max_{e\in E}\{q_e(f_e)\}$; whereas a network with an "additive" objective aims at minimizing the *network total cost* $C(f)$, defined as the sum of all link performance functions i.e., $C(f) \triangleq \sum_{e\in E} q_e(f_e)$. Finally, for each user $u\in U$, we define the *bottleneck of u* as the worst performance among the links that u employs i.e., $b_u(f) \triangleq \max_{e\in E | f_e^u>0}\{q_e(f_e)\}$; similarly, for each $p\in P$, we define the *bottleneck of path p* as the worst link that belongs to p i.e., $b_p(f) \triangleq \max_{e\in p}\{q_e(f_e)\}$.

The triple $\langle G,U,\{q_e\}\rangle$ is termed a *splittable instance* if all the users in U can split their traffic; it is termed an *unsplittable instance* if all the users of U are limited to a single path.

B. Nash Equilibrium

We assume that each user $u\in U$ selects an assignment of traffic to paths from $P^{(s_u,t_u)}$ so as to meet its traffic demand γ_u. This selection shall also be referred to as the *user strategy*. The collection of all possible strategies of a user is referred to as the *user strategy space*. Finally, the product of all user strategy spaces is called the *joint strategy space*; each element in the joint strategy space is termed a (strategy) *profile*, and it is actually a flow vector that satisfies the demands of all users.

Users are assumed to behave selfishly, i.e., ship their demands through connections that have the lowest possible bottlenecks. This setting gives rise to a *non-cooperative game* [17]. A flow f is said to be at Nash equilibrium if each user considers its chosen traffic assignment to be the best under the given choice of other users. This is formalized as follows.

Definition 1: Considering a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, a flow vector $f=\{f_p^u\}$ is *at Nash equilibrium* if, for each user $\tilde{u}\in U$ and each feasible (splittable) flow vector $g=\{g_p^u\}$ that satisfies $g_p^u = f_p^u$ for each $u\in U\setminus\{\tilde{u}\}$, it holds that $b_{\tilde{u}}(f)\le b_{\tilde{u}}(g)$. f is then termed *a splittable Nash flow*.

The Nash equilibrium for unsplittable flow is defined similarly, as follows.

Definition 2: Considering an unsplittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, a flow vector $f=\{f_p^u\}$ is *at Nash equilibrium* if, for each user $\tilde{u}\in U$ and each feasible (unsplittable) flow vector $g=\{g_p^u\}$ that satisfies $g_p^u = f_p^u$ for each $u\in U\setminus\{\tilde{u}\}$, it holds that $b_{\tilde{u}}(f)\le b_{\tilde{u}}(g)$. f is then termed *an unsplittable Nash flow*.

Let $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ be a splittable instance for which f is a Nash flow, f_A^* is an optimal flow with respect to the (additive) network total cost objective $C(\cdot)$ and f_B^* is an optimal flow with respect to the network bottleneck objective $B(\cdot)$. *The prices of anarchy of splittable flows with respect to bottleneck* and *additive network objectives* are $\rho_{B\times S} \triangleq \frac{B(f)}{B(f_B^*)}$ and $\rho_{A\times S} \triangleq \frac{C(f)}{C(f_A^*)}$, respectively. Similarly, let $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ be an unsplittable instance for which $\widehat{f}$ is a Nash flow, $\widehat{f_A^*}$ is an optimal flow with respect to the (additive) network total cost objective $C(\cdot)$

and $\widehat{f_B^*}$ is an optimal flow with respect to the network bottleneck objective $B(\cdot)$. *The prices of anarchy of unsplittable flows with respect to bottleneck* and *additive network objectives* are $\rho_{B\times US} \triangleq \frac{B(\widehat{f})}{B(\widehat{f_B^*})}$ and $\rho_{A\times US} \triangleq \frac{C(\widehat{f})}{C(\widehat{f_A^*})}$, respectively.

Remark 1: It is easy to show that the optimal flows $f_A^*, f_B^*, \widehat{f_A^*}, \widehat{f_B^*}$ can be formulated as solutions of convex programs over the variables $\{f_p^u\}$. While the formulations of $f_A^*, \widehat{f_A^*}$ require the minimization of $\sum_{e\in E} q_e(f_e)$, the formulation of $f_B^*, \widehat{f_B^*}$ require the minimization of $\max_{e\in E}\{q_e(f_e)\}$, where in all formulations $f_e = \sum_{u\in U}\sum_{p|e\in p} f_p^u$ for each $e \in E$. Moreover, while the formulations of the splittable flows (namely, f_A^*, f_B^*) allow each variable f_p^u to take any value in the range 0 to γ_u, the formulations of the unsplittable flows (namely, $\widehat{f_A^*}, \widehat{f_B^*}$) restrict f_p^u to be either 0 or γ_u for each $u \in U$ and path $p \in P$.

Remark 2: It is easy to show that the results of this study on bottleneck objectives hold also for link performance functions $\{q_e(f_e)\}$ that are *concave* continuous and *decreasing* in the flows $\{f_e\}$, and the goal (of the users and the network) is to *maximize* the bottleneck values.

III. Existence and Non-Uniqueness of Flows at Nash Equilibrium

In this section, we establish the existence and non-uniqueness of splittable and unsplittable Nash flows.

A. Existence of Nash Equilibrium

We begin with the splittable case.

Theorem 1: Given a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, there exists a splittable flow vector *f* that is at Nash equilibrium.

Proof (sketch) It was established in [18] that, if the allowed strategies for an *n*-person game are limited to a convex, closed and bounded set in some Euclidean space, and if the payoff function of each player is continuous and convex, then the game has a Nash equilibrium. In the Appendix, we show that these conditions are satisfied in our case. ■

Next, the following theorem establishes the existence of a Nash equilibrium for *unsplittable* flows.

Theorem 2 Given an unsplittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, there exists an unsplittable flow vector *f* that is at Nash equilibrium.

Proof: First, note that, since *f* is an unsplittable flow vector, each user selects one out of at most $|P|$ possible paths; hence, each user has at most $|P|$ strategies. Therefore, since there are $|U|$ users, the joint strategy space of the game is finite. Assume by way of contradiction that there is no Nash equilibrium. Hence, for each profile in the joint strategy space, there is at least one user that can improve its bottleneck. Let $\langle \theta_1,\theta_2,\cdots\rangle$ be a sequence of profiles, such that, for each two consecutive profiles θ_i,θ_{i+1}, it holds that exactly one user in θ_{i+1} reroutes its traffic and improves its bottleneck with respect to θ_i. Since the joint strategy space is finite, it follows that, after a finite number of transitions between successive profiles in $\langle \theta_1,\theta_2,\cdots\rangle$, we must encounter some profile θ for the second time, and, without loss of generality, let it be θ_1 and let $\theta_{n+1}=\theta_1$. Consider then the sequence of profiles $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$. Let $\widetilde{U} \subseteq U$ be the collection of users whose bottleneck is not constant over all profiles of $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$. Let $u \in \widetilde{U}$ be a user that produces the largest (worst) bottleneck among all users of $\widetilde{U}$ and over all profiles of $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$. Let $\Theta=\{\theta_k\}$ be the collection of all profiles in $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$ for which *u* achieves the worst bottleneck. Since $u \in \widetilde{U}$, it follows that there exists at least one profile $\theta_k \in \Theta$ such that $\theta_{k-1} \notin \Theta$; let θ_k be such a profile.

In the transition from the profile θ_{k-1} to the profile θ_k, there exists exactly one user $u' \in U$ that reroutes its traffic in order to improve its bottleneck. Since u' improves its bottleneck, it follows that the bottleneck of *u*' is not constant in $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$; hence, $u' \in \widetilde{U}$. Next, since the bottleneck of *u* has a smaller value in θ_{k-1} than in θ_k, it follows that *u*' transfers a positive amount of flow through some bottleneck of *u* in the profile θ_k. Therefore, by definition, the bottleneck of *u*' in the profile θ_k is equal to that of *u*. However, since the bottleneck of *u* in the profile θ_k is maximal with respect to all users in $\widetilde{U}$ and all profiles of $\langle \theta_1,\theta_2,\cdots\theta_n,\theta_1\rangle$, the bottleneck of the user $u' \in U$ is not improved in the transition from θ_{k-1} to θ_k. This contradicts the way by which user u' was selected, hence establishing the theorem. ■

B. Non-Uniqueness of Flows at Nash Equilibrium

We now show, by way of an example, that the Nash Equilibrium is not unique, for both splittable and unsplittable flows. Consider the network presented in Figure 1. Suppose that there exists a single user that needs to transfer one unit of flow from *s* to *t*, and assume that each link $e \in E$ is assigned with a performance function $q_e(f_e) = f_e$. Let $p_1 = (e_1, e_3)$ and $p_2 = (e_2, e_3)$. One can see that both $\langle f_{p_1} = 1, f_{p_2} = 0 \rangle$ and $\langle f_{p_1} = 0, f_{p_2} = 1 \rangle$ are optimal flow vectors with respect to both splittable and unsplittable flows. As optimal and Nash flows coincide in the case of a single user, we have identified two different Nash flows.

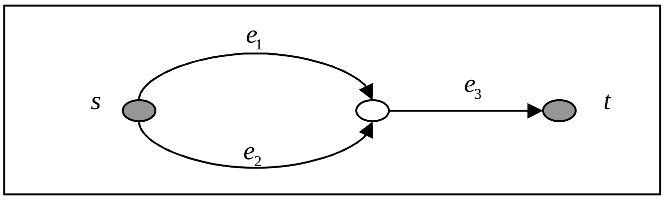

Fig. 1: Non-uniqueness of flow at Nash Equilibrium

IV. The Price of Anarchy for Bottleneck Network Objectives

In this section, we investigate the price of anarchy with respect to *bottleneck network objectives* for splittable and unsplittable flows. Specifically, we show that, when users are allowed to split their traffic over any number of paths, the price of anarchy $\rho_{B\times S}$ equals to 1 i.e., *the lack of regulation for the splittable case incurs no cost*. On the other hand, we show that, when such splitting is not possible, then the price of anarchy can be unbounded i.e., $\rho_{B\times US} = \infty$.

A. No Price of Anarchy with Splittable flows

This subsection establishes that the (network) bottleneck obtained by selfish users is optimal for splittable flows. We begin by introducing the following definition.

Definition 3: Given is a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ for which f is a Nash flow. Denote by $\Gamma(f)$ the collection of all users that ship traffic through a network bottleneck i.e., $\Gamma(f) = \{u \in U \,|\, b_u(f) = B(f)\}$.

The following lemma states that, if *f* is a splittable Nash flow, then deleting users that are not in $\Gamma(f)$ and zeroing their respective flows, results in a flow that is at Nash equilibrium (with respect to the new game that corresponds to the reduced set of users). Moreover, the new flow keeps the bottleneck of the network unchanged.

Lemma 1: Given a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ for which f is a Nash flow, consider the function $g{:}P\times\Gamma(f)\rightarrow\mathbb{R}^{+}\cup\{0\}$ that satisfies $f_p^u = g_p^u$ for each $u\in\Gamma(f)$ and $p \in P$. Then, *g* is a splittable flow vector that satisfies the demands of all users in $\Gamma(f)$, such that the following three properties hold:

(i) *g* is at Nash equilibrium;
(ii) the (network) bottleneck of *g* equals to that of *f*, i.e., $B(f)=B(g)$;
(iii) the bottleneck of each user $u\in\Gamma(f)$ is equal to that of the network i.e., $b_u(g)=B(g)$ for each $u\in\Gamma(f)$.

The proof appears in the Appendix. While the details are rather tedious, the basic idea is quite straightforward, namely, the removal of users not in $\Gamma(f)$ cannot change the flow on any of the bottleneck links in the network.

Remark 3: As indicated in the Appendix, Lemma 1 applies only to splittable flows, and its proof makes a crucial use of the fractional nature of such flows.

We are now ready to prove a rather surprising result, which essentially states that, for bottleneck (network) objectives, anarchy comes for free when users are allowed to split their traffic.

*Remark 4***:** In the following theorem we assume that the traffic between each source-destination pair belongs to at most one user. However, it is easy to see that, for the splittable case, this assumption incurs no loss of generality, since all the users that ship traffic between the same source-destination pair can be regarded as a single "aggregated" user with a traffic demand that is equal to the sum of demands of all the corresponding users. On the other hand, we note that, in the unsplittable case, such an assumption does incur loss of generality.

Theorem 3: Given a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, the price of anarchy with respect to bottleneck network objectives is 1, i.e., $\rho_{B\times S} = 1$.

*Proof***:** Consider the Nash flow *f* for the instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ and the corresponding set of users $\Gamma(f)$. Let $g{:}P\times\Gamma(f)\rightarrow\mathbb{R}^{+}\cup\{0\}$ satisfy $g_p^u=f_p^u$ for each $u\in\Gamma(f)$ and $p \in P$. From Lemma 1 it follows that *g* is a splittable flow vector that satisfies the demands of all users in $\Gamma(f)$, such that:

(i) g is a Nash flow;

(ii) $B(f)=B(g)$ i.e., the (network) bottleneck of g equals to that of f;

(iii) $b_u(g)=B(g)$ for each $u\in\Gamma(f)$ i.e., the bottleneck of each user is equal to that of the network.

We now show that $B(g)$ is the minimum (network) bottleneck among the bottlenecks of the (splittable) flow vectors that satisfy the demands of the users in $\Gamma(f)$.

First, we note that, since $b_u(g)=B(g)$ for each $u\in\Gamma(f)$, it holds by definition that there exists a $p\in P^{(s_u,t_u)}$ such that $g_p>0$ and $b_p(g)=B(g)$ for each $u\in\Gamma(f)$.

Assume, by way of contradiction, that $B(g)$ is not minimum. Hence, there exists a feasible flow vector h that satisfies the demands of all users in $\Gamma(f)$ such that $B(h)<B(g)$. Consider a user $u\in\Gamma(f)$ and its corresponding source-destination pair $s_u,t_u\in V$. Denote by $E^{(s_u,t_u)}$ the collection of all network bottlenecks with respect to g that belong to a path in $P^{(s_u,t_u)}$ i.e., $E^{(s_u,t_u)}=\left\{e\,\middle|\,e\in p, p\in P^{(s,t)}, q_e(g_e)=B(g)\right\}$. In addition, denote by $\overline{P^{(s_u,t_u)}}$ the collection of all paths in $P\setminus P^{(s_u,t_u)}$ that intersect with paths in $P^{(s_u,t_u)}$ on at least one network bottleneck from $E^{(s_u,t_u)}$ i.e., $\overline{P^{(s_u,t_u)}}\triangleq\left\{p\,\middle|\,\begin{array}{l}p\in P\setminus P^{(s_u,t_u)},\exists p'\in P^{(s_u,t_u)}\text{ s.t. for at least one}\\ \text{common link, } e\in p\cap p',\text{ it holds that } e\in E^{(s_u,t_u)}\end{array}\right\}$. Since g is a Nash flow, it follows that *every* path in $P^{(s_u,t_u)}$ traverses through at least one of the bottlenecks from $E^{(s_u,t_u)}$; indeed, otherwise u could have decreased its bottleneck value (which equals $B(g)$) by rerouting flow into paths whose bottlenecks are lower. Therefore, in order to reduce the value of *all* network bottlenecks in $E^{(s_u,t_u)}$, it follows that either the flow demand of user u or the total traffic that is carried over the paths from $\overline{P^{(s_u,t_u)}}$ must decrease. Since $b_p(g)=B(g)>B(h)$ for each $p\in P^{(s_u,t_u)}$, it holds that $q_e(h_e)<q_e(g_e)$ for each $e\in E^{(s_u,t_u)}$. Therefore, it holds for the flow vector h that either the flow demand of user u or the total traffic that is carried by the paths from $\overline{P^{(s_u,t_u)}}$ is smaller than in g. Since the flow demand requirement of user u is γ_u both in g and in h, it follows that the total flow of the paths in $\overline{P^{(s_u,t_u)}}$ is reduced i.e., $\sum_{p\in\overline{P^{(s_u,t_u)}}}h_p<\sum_{p\in\overline{P^{(s_u,t_u)}}}g_p$.

As already mentioned, since g is a Nash flow, every path in $P^{(s_u,t_u)}$ traverses at least one link from $E^{(s_u,t_u)}$. Thus, since $\overline{P^{(s_u,t_u)}}$ is the collection of all paths from $P\setminus P^{(s_u,t_u)}$ that intersect with paths from $P^{(s_u,t_u)}$ on at least one network bottleneck, it follows that $\overline{P^{(s_u,t_u)}}$ is the collection of all paths in $P\setminus P^{(s_u,t_u)}$ that intersect with paths from $P^{(s_u,t_u)}$ i.e., $\overline{P^{(s_u,t_u)}}=\left\{p\,\middle|\,p\in P\setminus P^{(s_u,t_u)}\,\middle|\,\exists p'\in P^{(s_u,t_u)}\text{ s.t. } p'\cap p\neq\phi\right\}$. Therefore, $\overline{P}\triangleq\bigcup_{u\in\Gamma(f)}\overline{P^{(s_u,t_u)}}$ is the collection of all the paths that connect the source-destination pair of some user and intersect with paths that connect the source-destination pair of some other user i.e., $\overline{P}=\left\{p\in P\,\middle|\,\exists p'\in P^{(s_u,t_u)}\text{ s.t. } p'\cap p\neq\phi\text{ and } p\notin P^{(s_u,t_u)}\right\}$. Thus, since we have shown that $\sum_{p\in\overline{P^{(s_u,t_u)}}}h_p<\sum_{p\in\overline{P^{(s_u,t_u)}}}g_p$ for each $u\in\Gamma(f)$, it follows that $\sum_{p\in\overline{P}}h_p=\sum_{p\in\bigcup_{u\in\Gamma(f)}\overline{P^{(s_u,t_u)}}}h_p=\sum_{u\in\Gamma(f)}\sum_{p\in\overline{P^{(s_u,t_u)}}}h_p<\sum_{u\in\Gamma(f)}\sum_{p\in\overline{P^{(s_u,t_u)}}}g_p=$
$=\sum_{p\in\bigcup_{u\in\Gamma(f)}\overline{P^{(s_u,t_u)}}}g_p=\sum_{p\in\overline{P}}g_p$.

Finally, note that, for each $u\in\Gamma(f)$, the total traffic from s_u to t_u equals γ_u both in g and in h i.e., $\sum_{p\in P^{(s_u,t_u)}}h_p=\sum_{p\in P^{(s_u,t_u)}}g_p$ for each $u\in\Gamma(f)$. Hence, the total traffic that is carried over the network paths must be the same in both cases i.e., $\sum_{p\in P}h_p=\sum_{p\in P}g_p$. Therefore, since we established that $\sum_{p\in\overline{P}}h_p<\sum_{p\in\overline{P}}g_p$, it follows that $\sum_{p\in P/\overline{P}}h_p>\sum_{p\in P/\overline{P}}g_p$; that is, the total traffic over the paths that do not intersect with paths of other users is larger in h than in g. Therefore, there exists at least one user $u\in\Gamma(f)$ that satisfies

$$\sum_{p\in(P/\overline{P})\cap P^{(s_u,t_u)}}h_p>\sum_{p\in(P/\overline{P})\cap P^{(s_u,t_u)}}g_p \tag{1}$$

i.e., the total traffic of user u, carried over the paths that do not intersect with paths of other users, is increased. Since the paths $(P/\overline{P})\cap P^{(s_u,t_u)}$ do not intersect with paths of other users, it follows that the traffic that traverses over each link $e\in p$, where $p\in(P/\overline{P})\cap P^{(s_u,t_u)}$, is only that of user u. Therefore, since each path $p\in P^{(s_u,t_u)}$ must traverse at least one link that is a bottleneck with respect to g, it follows from (1) that there exists a link $e \in p$, where $p \in \left(P / \overline{P}\right) \cap P^{(s_u, t_u)}$, that was a bottleneck with respect to g and satisfies $h_e > g_e$. Therefore, since $q_e(\cdot)$ is increasing, it follows that $q_e(h_e) > q_e(g_e) = B(g)$. Obviously, this contradicts the assumption that $B(h) < B(g)$. Hence, $B(g)$ has the minimum bottleneck over all flow vectors that satisfy the demands of the users in $\Gamma(f)$.

We point out that, while we have shown that g achieves the minimum value when considering only the users in $\Gamma(f)$, we still need to show that f achieves the minimum value when considering all users in U. Denote by f^* an optimal flow, i.e., one that has the smallest bottleneck with respect to all flow vectors that satisfy the demands of the users of U. Also, let H be the set of all splittable flow vectors that satisfy the demands of the users in $\Gamma(f)$. We transform the vector f^* into a flow that belongs to the set H by zeroing the flow of all users in $U \setminus \Gamma(f)$ i.e., by zeroing all the flows in $\left\{ f_p^{*u} \middle| p \in P, u \in U \setminus \Gamma(f) \right\}$; denote the resulting flow vector by $\widetilde{f^*}$ and note that, by construction, $f_e^* \geq \widetilde{f_e^*}$ for each $e \in E$. Since $q_e(\cdot)$ is increasing, it follows that $q_e\left(f_e^*\right) \geq q_e\left(\widetilde{f_e^*}\right)$ for each $e \in E$. Therefore, by definition,

$$B\left(f^*\right) \geq B\left(\widetilde{f^*}\right). \tag{2}$$

However, since $\widetilde{f^*} \in \mathrm{H}$, it follows that $B\left(\widetilde{f^*}\right) \geq \min\left\{B(h) \middle| h \in \mathrm{H}\right\}$. Therefore, from (2), it follows that $B\left(f^*\right) \geq \min\left\{B(h) \middle| h \in \mathrm{H}\right\}$. However, we have shown that $B(g)$ has the minimum bottleneck over all flow vectors that satisfy the demands of the users in $\Gamma(f)$ i.e., $B(g) = \min\left\{B(h) \middle| h \in \mathrm{H}\right\}$. Therefore, $B\left(f^*\right) \geq B(g)$. Finally, since we established (item (ii) in the beginning of the proof) that $B(g) = B(f)$, it follows that $B\left(f^*\right) \geq B(f)$; therefore, since f^* is optimal, $B\left(f^*\right) = B(f)$ i.e., $\rho_{B \times S} = \dfrac{B(f)}{B\left(f^*\right)} = 1$. Thus, the Theorem is established. ■

Remark 5: The Braess paradox [6] shows that addition of links to a noncooperative network can negatively impact the performance of both the network and each of the users. However, it follows from Theorem 3 that, for bottleneck network objectives and splittable flows, this paradox cannot occur.

B. Unbounded Price of Anarchy with Unsplittable Flows

We now consider the case where each user is restricted to route over a *single path*. It turns out that, with such a restriction, the result of the previous subsection no longer holds. Indeed, the following example shows that an unsplittable Nash flow may have an arbitrarily larger bottleneck than that of an optimal flow.

Consider the network presented in Fig. 2, and suppose there are two users, each with the same source *s* and destination *t*. For each $\gamma > 0$, the first user (user *A*) has to transfer a demand of γ units and the other user (user *B*) has to transfer a demand of $2 \cdot \gamma$ units. Finally, assume that both users must transfer their demands *unsplittably*. In the optimal solution, *A* is assigned to the upper link and *B* is assigned to the lower link; the corresponding bottleneck is $\max\left\{e^{\frac{2}{3}\cdot\gamma}, e^{\frac{1}{2}\cdot 2\cdot\gamma}\right\} = e^{\gamma}$. On the other hand, a profile where *A* chooses the lower link and *B* chooses the upper link is an unsplittable Nash flow with a bottleneck of $\max\left\{e^{\frac{2}{3}\cdot 2\gamma}, e^{\frac{1}{2}\cdot\gamma}\right\} = e^{\frac{4\gamma}{3}}$. Hence, the price of anarchy is $\rho_{B \times US} = \dfrac{B(f)}{B\left(f^*\right)} = \dfrac{e^{\frac{4\gamma}{3}}}{e^{\gamma}} = e^{\frac{\gamma}{3}}$; as γ can be arbitrarily large, $\rho_{B \times US}$ is unbounded.

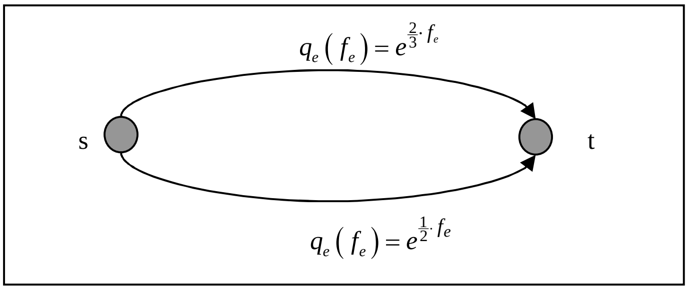

Fig. 2: Unbounded price of anarchy for unsplittable flows

V. The Price of Anarchy for Additive Network Objectives

In this section, we investigate the price of anarchy with respect to additive network objectives (i.e., minimization of total cost). Specifically, we show that, if all users splittably route their traffic with respect to *bottleneck* metrics, the resulting total cost is at most *M* times the optimal cost i.e., the price of anarchy $\rho_{A \times S}$ is at most *M*. On the other hand, we show that the price of anarchy for unsplittable flows remains unbounded also for additive objectives i.e., $\rho_{A\times US} = \infty$. Then, we turn to indicate that the bound of *M* on the price of anarchy when splittable flows are considered suggests a potential mechanism for bounding the price of anarchy in the case where both the *users* and the *network* consider *additive* objectives.

Before proceeding we mention that, for unsplittable flows, the price of anarchy remains unbounded. Indeed, consider again the example presented in Subsection IV.B. Since users still consider bottleneck objectives, the original Nash flow remains to be such. Also, it is easy to see that the original optimal network flow remains to be such in the (network) additive case, hence, $\rho_{A\times US} = \frac{C(f)}{C(f^*)} = \frac{e^{\frac{4\cdot\gamma}{3}} + e^{\frac{\gamma}{2}}}{e^{\frac{2\cdot\gamma}{3}} + e^{\gamma}} \xrightarrow[\gamma\to\infty]{} \infty$.

A. The Price of Anarchy is at Most M with Splittable flows

In this subsection we show that, when users route their traffic with respect to bottleneck metrics and are allowed to split their traffic, the resulting total cost is at most *M* times the optimal cost i.e., $\rho_{A\times S} \le M$.

Theorem 4: Given a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, the price of anarchy with respect to additive network objectives is at most *M*, i.e., $\rho_{A\times S} \le M$.

Proof: Let f *and* f^* denote a Nash flow and an optimal flow for the splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, correspondingly. It follows from Theorem 3 that the bottleneck $B(f) = \max_{e\in E}\{q_e(f_e)\}$ is optimal with respect to all flow vectors that consider the demands of the users in *U*. Hence, $B(f) \le B(f^*)$. Therefore, by definition, $\max_{e\in E}\{q_e(f_e)\} \le \max_{e\in E}\{q_e(f_e^*)\}$. Hence, it follows that

$$C(f) = \sum_{e\in E} q_e(f_e) \le M \cdot \max_{e\in E}\{q_e(f_e)\} \le M \cdot \max_{e\in E}\{q_e(f_e^*)\}. \quad (3)$$

Thus, since $\max_{e\in E}\{q_e(f_e^*)\} \le \sum_{e\in E} q_e(f_e^*)$, it follows from (3) that $C(f) \le M \cdot \max_{e\in E}\{q_e(f_e^*)\} \le M \cdot \sum_{e\in E} q_e(f_e^*) = M \cdot C(f^*)$.

Hence, $\rho_{A\times S} = \frac{C(f)}{C(f^*)} \le M$, and the theorem is established.

■

B. Possible Application

The above worst case ratio *M* with respect to additive network objectives (for splittable users) potentially gives rise to a scheme for bounding the price of anarchy in the case where both the network *and* the users consider *additive objectives*, e.g., delay. To that end, we first formalize the corresponding traffic model. For concreteness, we focus on delay minimization.

Assume that we are given a network $G(V,E)$, a *link latency function* $l_e(\cdot)$ for each $e\in E$, and a set of users *U* where each $u\in U$ intends to send γ_u units of flow from s_u to t_u. Following the traffic model of [20], we define the average latency of a flow *f* (often termed as *total latency*) as $\sum_{e\in E} l_e(f_e)\cdot f_e$ In addition, we assume that each user routes its traffic so as to minimize the latency it experiences i.e., each user $u\in U$ minimizes the value $\sum_{e\in E} l_e(f_e)\cdot f_e^u$.

Recently, it has been established [22] that, when users are allowed to split their traffic, the price of anarchy in the above model is at most $|U|$ [1]. Obviously, when $|U|$ is large, this price may be prohibitive.

In the previous subsection we showed that, if all users splittably route their traffic with respect to *bottleneck* metrics, the price of anarchy with respect to the (additive) total cost objective is at most *M*. Therefore, if the performance function of each link $e\in E$ is set to $q_e(f_e) = l_e(f_e)\cdot f_e$ and users are driven to route their traffic according to bottleneck metrics, the price of anarchy with respect to total latency is bounded by *M*. Thus, for a large number of users, specifically $|U| > M$, it is tempting to make users consider *bottleneck* metrics (i.e., even though their original objective is additive).

Accordingly, we are currently working on a mechanism for driving selfish users (that aim at minimizing their delays) route their traffic with respect to bottleneck metrics. The basic idea is that the network advertises to the users *solely* the condition of the worst links (i.e., the network bottlenecks) and avoids advertising the condition of the rest of the links (i.e., the network image from the user's perspective is $q_e(f_e) = 0$ if $l_e(f_e)\cdot f_e \ne \max_{e'\in E}\{l_{e'}(f_{e'})\cdot f_{e'}\}$ and $q_e(f_e) = l_e(f_e)\cdot f_e$ otherwise).

VI. Conclusion

We studied the behavior of self-optimizing users with bottleneck objectives. We evaluated the worst-possible degradation of network performance (due to such behav-

[1] Furthermore, in [20] it has been established that, for $|U|\to\infty$ and $\gamma_u \to 0$, the price of anarchy is unbounded.

ior) with respect to two main classes of metrics namely, bottleneck and additive. Our results are summarized in Table 1.

Network Objective \ Flow	Splittable	Unsplittable
Bottleneck	1	∞
Additive	*M*[1]	∞

Table 1: The price of anarchy for splittable and unsplittable flows with respect to additive and bottleneck network objectives.

Our results clearly show that the combination of lack of regulation and inability to split traffic may incur a prohibitive cost. Therefore, in the absence of some global regulation, the employment of splittable flows is much more desired. In particular, for bottleneck network objectives, the lack of regulation incurs no cost. This finding suggest an important network design rule, favoring *multipath routing* schemes over the more traditional single-path routing [25].

There are several directions for future research that stem from this study. First, it remains an open question whether the upper bound of *M* on the price of anarchy of splittable flows with respect to additive objectives is tight. At a more general level, one combination that has not been considered yet[2], is that of a network with bottleneck objectives and users with additive objectives. Furthermore, it is of interest to investigate a "hybrid" case, where some users consider bottleneck objectives while others consider additive objectives. Finally, in view of the prohibitive value of the price of anarchy for a large population of users that (together with the network) consider additive objectives [20],[22], it is of interest to find schemes for driving "additive" self optimizing users consider bottleneck objectives. As mentioned, we are currently working on such a scheme.

Acknowledgement: We would like to thank Amir Ronen for helpful discussions.

REFERENCES

[1] E. Altman, *Flow Control using the theory of zero-sum Markov games*, IEEE Transactions on Automatic Control 39(4) (April 1994), 814-818.

[2] D. Awduche and J. Malcolm, J. Agogbua, M. O'Dell and J. McManus, *Requirements for Traffic Engineering Over MPLS*, IETF RFC 2702, September 1999.

[3] S. Banerjee and J. Yoo, *Minimizing Maximum Logical Link Congestion in Packet Switched Optical Networks*, ICC, Monréal, Québec, Canada, June 1997.

[4] D. Bertsekas and R. Gallager, *Data networks*, Prentice-Hall, 1992.

[5] S. Boyd and L. Vandenberghe, *Convex Optimization*, available from http://www.stanford.edu/~boyd/bv_cvxbook.pdf.

[6] D. Braess, *Uber ein paradoxon der verkehrsplanung*, Unternehmensforschung, 12:258-268, 1968.

[7] J. H. Chang and L. Tassiulas, *Routing for Maximum System Lifetime in Wireless Ad-Hoc Networks*, In proceedings of the 37th Annual Allerton Conf. On Commun., Control, and Comp., September 1999.

[8] J. H. Chang and L. Tassiulas, *Energy Conserving Routing in Wireless Ad-Hoc Networks*, In proceedings of IEEE Infocom, Tel-Aviv, Israel, March 2000.

[9] J. H. Chang and L. Tassiulas, *Fast Approximate Algorithm for Maximum Lifetime Routing in Wireless Ad-Hoc Networks*, In proceedings of IFIP-TC6 Networking, Lecture Notes in Computer Science, vol. 1815, Springer, Berlin, May 2000.

[10] A. Czumaj and B. V¨ocking, *Tight bounds for worst-case equilibria*, In proceedings of the 13th Annual Symposium on Discrete Algorithms, San Francisco CA, January 2002.

[11] P. Dubey, *Inefficiency of Nash Equilibria*, Mathematics of Operations Research, 11(1): 1-8, February 1986.

[12] Y. A. Korilis and A. A. Lazar, *On the existence of equilibria in noncooperative optimal flow control*, Journal of the ACM 42(3), May 1995.

[13] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, In proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 1999.

[14] A. A. Lazar, A. Orda and D. E. pendarakis, *Virtual path bandwidth allocation in multi-user networks*, In proceedings of Infocom 1995, Boston, MA, April 1995.

[15] M. Mavronicolas and P. Spirakis, *The price of selfish routing*, In proceedings of the 33rd Annual ACM Symposium on the Theory of Computing, Heraklion, Crete, Greece, July 2001.

[16] A. Orda, R. Rom and N. Shimkin, *Competitive routing in multiuser communication networks*, IEEE/ACM Transaction on Networking 1(5), October 1993.

[17] G. Owen. *Game Theory*, Academic Press, Third Edition, 1995.

[18] J. B. Rosen, *Existence and uniqueness of equilibrium points for concave N-person games*, Econometria, Vol. 33, pp. 510-534, July 1965.

[19] T. Roughgarden, *Designing networks for selfish users is hard*, In proceedings of the 42nd Symposium on Foundations of Computer Science, Las Vegas, NV, October 2001.

[20] T. Roughgarden and E. Tardos, *How bad is selfish routing?*, Journal of the ACM, vol. 49, no. 2, pp. 236-259, March 2002.

[21] T. Roughgarden, *The price of anarchy is independent of network topology*, In proceedings of the 34th Annual ACM Symposium on Theory of Computing Montréal, Québec, Canada, May 2002.

[22] T. Roughgarden, Private Communication, 2004.

[23] W. Rudin, *Principles of mathematical analysis*, Third Edition, McGraw Hill, 1976.

[24] S. Singh, M. Woo, and C. Raghavendra, *Power-aware routing in mobile ad hoc networks*, In proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking, Dallas, Texas, October 1998.

[25] C. Villamizar, *OSPF Optimised Multipath (OSPF-OMP)*, Internet Draft <draft-ietf-ospf-omp-02.txt> 1998.

[1] More precisely, we have shown that the worst price of anarchy for splittable flows with respect to additive network objectives is *upper-bounded* by *M*, but this is not a tight bound necessarily.

[2] Neither by this study nor by [20],[22].

[26] Y. Wang and Z. Wang, *Explicit routing algorithms for Internet Traffic Engineering*, In proceedings of ICCN'99, Boston, October 1999.

[27] Z. Zhang and D. Doulieris, *Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment*, IEEE Transactions on Computers vol.40, No.8, August 1992.

Appendix

This Appendix contains the proofs of Theorem 1 and Lemma 1. In addition, we indicate why Lemma 1 does not holds for unsplittable flows.

A.1 Proof of Theorem 1

Theorem 1: Given a splittable instance $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, there exists a splittable flow vector *f* that is at Nash equilibrium.

Proof It was established in [18] that, if the allowed strategies for an *n*-person game are limited to a convex, closed and bounded set in some Euclidean space, and if the payoff function of each player is continuous and convex, then the game has a Nash equilibrium. We now show that these conditions are satisfied in our case. Let Ω be the set of all splittable flow vectors that satisfy the demands of the users in *U*. Since for each $k \geq 1$ it holds that the function $\max(\cdot)$ is continuous in the Euclidean space $\mathbb{R}^k$, and since for each $e \in E$ we assume that $q_e(\cdot)$ is continuous in f_e (hence also continuous in Ω), it follows that the user's payoff function, namely, $b_u(\cdot) = \max_{e\in E | f_e^u>0} \{q_e(\cdot)\}$, is continuous in Ω. Next, since we assume that, for each splittable flow vector $f \in \Omega$ and link $e \in E$, the performance function $q_e(f_e)$ is convex in f_e, it follows that, for each pair of splittable flow vectors $h,g \in \Omega$, link $e \in E$ and $\lambda \in [0,1]$, it holds that $q_e(\lambda \cdot h_e + (1-\lambda)\cdot g_e) \leq \lambda \cdot q_e(h_e) + (1-\lambda)\cdot q_e(g_e)$. Thus, it follows that $b_u(\lambda \cdot h + (1-\lambda)\cdot g) = \max_{e\in E | f_e^u>0} \{q_e(\lambda \cdot h_e + (1-\lambda)\cdot g_e)\} \leq$

$\leq \max_{e\in E | f_e^u>0} \{\lambda \cdot q_e(h_e) + (1-\lambda)\cdot q_e(g_e)\}$. Hence, it holds that

$$b_u(\lambda \cdot h + (1-\lambda)\cdot g) \leq \max_{e\in E | f_e^u>0} \{\lambda \cdot q_e(h_e) + (1-\lambda)\cdot q_e(g_e)\} \leq$$

$$\leq \lambda \cdot \max_{e\in E | f_e^u>0} \{q_e(h_e)\} + (1-\lambda)\cdot \max_{e\in E | f_e^u>0} \{q_e(g_e)\} =$$

$$= \lambda \cdot b_u(h) + (1-\lambda)\cdot b_u(g).$$

Therefore, $b_u(f)$ is convex in *f* for each user $u \in U$ and splittable flow vector $f \in \Omega$.

Clearly, the set of feasible (splittable) flow vectors Ω is convex. It is also easy to show that it is bounded and closed. For completeness we provide the details.

First, for each value $\lambda \in [0,1]$, pair of feasible flow vectors $h,g \in \Omega$ and user $u \in U$, it holds that $\lambda \cdot \sum_{p\in P^{(s_u,t_u)}} g_p^u + (1-\lambda)\cdot \sum_{p\in P^{(s_u,t_u)}} h_p^u = \lambda \cdot \gamma_u + (1+\lambda)\cdot \gamma_u = \gamma_u$. Therefore, $\lambda \cdot h + (1-\lambda)\cdot g$ is also a feasible splittable flow vector i.e., $\lambda \cdot h + (1-\lambda)\cdot g \in \Omega$. Thus, by definition, Ω is a convex set [5]. We turn to show that Ω is bounded. To that end, consider the Euclidean metric $d(x_1,x_2) \triangleq \sqrt{\sum_{i=1}^{|P|\cdot|U|} \left(x_1^i - x_2^i\right)^2}$, which is associated with the Euclidean space $\mathbb{R}^{|P|\cdot|U|}$ (recall that *U* is the set of all users and *P* is the collection of all paths in the network). Note that, for each $f \in \Omega$, it holds that $0 \leq f_p^u \leq \gamma_u$ for each link $e \in E$ and user $u \in U$. Therefore, for each two points (flow vectors) $h,g \in \Omega$, it holds that

$$d(h,g) = \sqrt{\sum_{p\in P, u\in U} \left(h_p^u - g_p^u\right)^2} \leq \sqrt{\sum_{p\in P, u\in U} (\gamma_u)^2} =$$

$= \sqrt{|P|\cdot|U|} \cdot \max_{u\in U}\{\gamma_u\}$. Hence, Ω is bounded. Finally, we show that Ω is a closed set. To that end, note that Ω is the space of solutions to the following algebraic equations.

$$\sum_{p\in P^{(s_u,t_u)}} f_p^u = \gamma_u \qquad ,\forall u \in U \quad ,\forall v \in V \setminus \{s_u,t_u\}$$

$$f_p^u \geq 0 \qquad ,\forall p \in P, \ \forall u \in U$$

The space of solutions Ω is a polyhedron in $\mathbb{R}^{|P|\cdot|U|}$. Moreover, since each of the above algebraic equations considers the sign of *equality*, it follows that Ω also consists of the polyhedron's surface. Therefore, the complement of Ω (denoted by Ω^c) is an open set [23]. Thus, since any set is open *iff* its complement is closed [4], it follows that $(\Omega^c)^c = \Omega$ is a closed set. Thus, we conclude that Ω is closed, convex and bounded. Hence, the Theorem is established. ■

A.2 Proof of Lemma 1

Lemma 1: Given are a network $G(V,E)$, a set of users *U*, for each $e\in E$ a performance function $q_e(\cdot)$, and a splittable Nash flow $f:P\times U\rightarrow\mathbb{R}^+\cup\{0\}$. Consider the function $g:P\times\Gamma(f)\rightarrow\mathbb{R}^+\cup\{0\}$ that satisfies $g_p^u=f_p^u$ for each $u\in\Gamma(f)$ and $p\in P$. Then, *g* is a splittable flow vector that satisfies the demands of all users in $\Gamma(f)$, such that the following three properties hold:

(i) *g* is at Nash equilibrium;

(ii) the (network) bottleneck of *g* equals to that of *f*, i.e., $B(f)=B(g)$;

(iii) the bottleneck of each user $u\in\Gamma(f)$ is equal to that of the network i.e., $b_u(g)=B(g)$ for each $u\in\Gamma(f)$.

Proof: First, note that, since $g_p^u=f_p^u$ for each $u\in\Gamma(f)$ and $p\in P$, it holds that g is a splittable flow vector that satisfies the flow demands of all users in $\Gamma(f)$ i.e., $\sum_{p\in P^{(s_u,t_u)}} g_p^u=\sum_{p\in P^{(s_u,t_u)}} f_p^u=\gamma_u$ for each $u\in\Gamma(f)$.

Next, we show that the bottleneck of *g* equals to that of *f* i.e., $B(f)=B(g)$. To that end, consider the given Nash flow $f:P\times U\rightarrow\mathbb{R}^+\cup\{0\}$ that satisfies the flow demand of all users in *U*. Denote by $\tilde{E}$ the set of all links that are the bottlenecks of the network with respect to *f* i.e., $\tilde{E}\triangleq\{e\in E\,|\,q_e(f_e)=B(f)\}$. Recall that $\Gamma(f)\subseteq U$ is the collection of all users that ship traffic through one or more bottleneck from $\tilde{E}$; hence, by definition, all users of *U* that ship positive traffic through a link $e\in\tilde{E}$ are in $\Gamma(f)$. Therefore, it holds that $f_e=\sum_{u\in U}\sum_{p|e\in P} f_p^u=\sum_{u\in\Gamma(f)}\sum_{p|e\in P} f_p^u$ for each $e\in\tilde{E}$; hence, since $g_p^u=f_p^u$ for each $u\in\Gamma(f)$ and $p\in P$, it holds that $f_e=\sum_{u\in\Gamma(f)}\sum_{p\in P} f_p^u=\sum_{u\in\Gamma(f)}\sum_{p\in P} g_p^u=g_e$ for each $e\in\tilde{E}$. Therefore,

$$q_e(g_e)=q_e(f_e) \text{ for each } e\in\tilde{E}; \qquad \text{(A1)}$$

hence, since $q_e(f_e)=B(f)$ for each $e\in\tilde{E}$, it follows that $q_e(g_e)=B(f)$; hence, by definition, $B(g)=\max_{e\in E}\{q_e(g_e)\}\geq B(f)$. However, since, by construction, $g_e\leq f_e$ for each $e\in E$, and since the performance functions $\{q_e(\cdot)\}$ are increasing, it follows that $q_e(g)\leq q_e(f)$ for each $e\in E$; therefore $B(g)\leq B(f)$; thus, we conclude that

$$B(g)=B(f). \qquad \text{(A2)}$$

We turn to prove that the bottleneck of each user of *g* is equal to that of the network i.e., $b_u(g)=B(g)$ for each $u\in\Gamma(f)$. To that end, consider a user $u\in\Gamma(f)$. Since $u\in\Gamma(f)$, it follows by definition that *u* must ship positive traffic through at least one bottleneck from $\tilde{E}$ in the flow vector *f*. Therefore, since $g_p^u=f_p^u$ for each $u\in\Gamma(f)$ and $p\in P$, it follows that user *u* must ship positive traffic through at least one link from $\tilde{E}$ also in the flow vector *g*; let $e\in\tilde{E}$ be a link that carries a positive traffic of user *u* in the flow vector *g*. Since we have shown in (A1) that $q_e(g)=q_e(f)$ for each $e\in\tilde{E}$, and since, by definition, $q_e(f)=B(f)$ for each $e\in\tilde{E}$, it follows that the bottleneck of user *u* in the flow vector *g* is at least $B(f)$ i.e., $b_u(g)=\max_{e\in E|g_e^u>0}\{q_e(g_e)\}\geq B(f)$. Therefore, from (A2), $b_u(g)\geq B(g)$. Finally, since the bottleneck of each user is at most that of the network i.e., $b_u(g)\leq B(g)$ for each $u\in\Gamma(f)$, it follows that

$$b_u(g)=B(g) \text{ for each } u\in\Gamma(f). \qquad \text{(A3)}$$

We turn to show that *g* is at Nash equilibrium. By way of contradiction, assume otherwise. Hence, there exists at least one user $u\in\Gamma(f)$ in *g* that can improve its bottleneck. Let $\tilde{u}\in\Gamma(f)$ be such a user. Since we established in (A3) that $b_u(g)=B(g)$ for each $u\in\Gamma(f)$, it follows that, in the flow vector *g*, user $\tilde{u}$ can decrease its bottleneck below $B(g)$. More formally, there exists a flow vector $\tilde{g}$ that satisfies $\widetilde{g_p^u}=g_p^u$ for each $u\in\Gamma(f)\setminus\{\tilde{u}\}$, such that $b_{\tilde{u}}(\tilde{g})<B(g)$. Consider this flow vector (i.e., the flow vector $\tilde{g}$). Since $b_{\tilde{u}}(\tilde{g})<B(g)$, it follows by definition that $q_e(\widetilde{g_e})<B(g)$ for each link $e\in E$ that carries positive traffic of user $\tilde{u}$, i.e., a link that satisfies $\widetilde{g_e^{\tilde{u}}}>0$; hence, since $B(g)=B(f)$, it follows that $q_e(\widetilde{g_e})<B(f)$ for each link $e\in E$ that satisfies $\widetilde{g_e^{\tilde{u}}}>0$. Let $\tilde{\tilde{E}}\subseteq\tilde{E}$ be the set of all network bottlenecks of *f* that satisfy $\widetilde{g_e^{\tilde{u}}}>0$. Since $q_e(\widetilde{g_e})<B(f)$ for each $e\in\tilde{\tilde{E}}$, it follows by definition that

$q_e\left(\widetilde{g_e}\right) < q_e\left(f_e\right)$ for each $e \in \tilde{\tilde{E}}$; hence, since $q_e(\cdot)$ is an increasing function, it follows that

$$\widetilde{g_e^{\tilde{u}}} < f_e^{\tilde{u}} \text{ for each } e \in \tilde{\tilde{E}}. \quad \text{(A4)}$$

We employ (A4) in order to establish that g is a Nash flow, as follows.

For each $\lambda \in [0,1]$, consider the vector $f(\lambda) = \left\{ f(\lambda)_p^u \right\}$, where $f(\lambda)_p^u = \begin{cases} f_p^u & u \in U \setminus \{\tilde{u}\} \\ \lambda \cdot \widetilde{g_p^{\tilde{u}}} + (1-\lambda) \cdot f_p^{\tilde{u}} & \text{otherwise} \end{cases}$ for each $p \in P$ and $u \in U$. It is easy to see that $f(\lambda)$ is a feasible flow vector that satisfies the demands of all users in U for each $\lambda \in [0,1]$. We now show that there exists a λ, $0 < \lambda < 1$, such that $b_{\tilde{u}}\left(f(\lambda)\right) < b_{\tilde{u}}\left(f\right)$. Since by definition $f(\lambda)_p^u = f_p^u$ for each user $u \in U \setminus \{\tilde{u}\}$ and $p \in P$, we actually show that user $\tilde{u}$ in the flow vector f can decrease its bottleneck by rerouting its traffic; obviously, this contradicts the fact that f is at Nash equilibrium and proves that g is a Nash flow. Note that, since $\tilde{u} \in \Gamma(f)$, it follows by definition that $b_{\tilde{u}}(f) = B(f)$; hence, in order to show that $b_{\tilde{u}}\left(f(\lambda)\right) < b_{\tilde{u}}(f)$, it is sufficient to show that there exists some $0 < \lambda < 1$ such that $b_{\tilde{u}}\left(f(\lambda)\right) < B(f)$.

To that end, consider the set $E \setminus \tilde{E}$ i.e., the set of all links that are not the bottleneck of the network with respect to f. It follows by definition that $q_e(f_e) < B(f)$ for each $E \setminus \tilde{E}$. Therefore, since the performance function $q_e(\cdot)$ is continuous for each $e \in E$, it follows by construction of $f(\lambda)$ that there exists some $\varepsilon > 0$ such that

$$q_e\left(f_e(\varepsilon)\right) < B(f) \text{ for each } e \in E \setminus \tilde{E}. \quad \text{(A5)}$$

Therefore, in order to establish that $b_{\tilde{u}}\left(f(\varepsilon)\right) < B(f)$, we only need to show that $q_e\left(f_e(\varepsilon)\right) < B(f)$ for each link $e \in \tilde{E}$ that satisfies $f(\varepsilon)_e^{\tilde{u}} > 0$. To that end, we distinguish between two cases, namely the links $\tilde{\tilde{E}} \subseteq \tilde{E}$ and the links $\tilde{E} \setminus \tilde{\tilde{E}}$. Consider the links $\tilde{\tilde{E}} \subseteq \tilde{E}$. Since we have proven in (A4) that $\widetilde{g_e^{\tilde{u}}} < f_e^{\tilde{u}}$ for each $e \in \tilde{\tilde{E}}$, it follows that $f(\varepsilon)_e^{\tilde{u}} = \varepsilon \cdot \widetilde{g_e^{\tilde{u}}} + (1-\varepsilon) \cdot f_e^{\tilde{u}} < \varepsilon \cdot \widetilde{f_e^{\tilde{u}}} + (1-\varepsilon) \cdot f_e^{\tilde{u}} = f_e^{\tilde{u}}$ for each $e \in \tilde{\tilde{E}}$. Hence, since it holds by construction that $f(\varepsilon)_e^u = f_e^u$ for each $u \in U \setminus \{\tilde{u}\}$ and $e \in E$, it follows that

$$f(\varepsilon)_e = f(\varepsilon)_e^{\tilde{u}} + \sum_{u \in U \setminus \{\tilde{u}\}} f(\varepsilon)_e^u < f_e^{\tilde{u}} + \sum_{u \in U \setminus \{\tilde{u}\}} f(\varepsilon)_e^u = f_e^{\tilde{u}} + \sum_{u \in U \setminus \{\tilde{u}\}} f_e^u = f_e$$

for each $e \in \tilde{\tilde{E}}$. Thus, since $q_e(\cdot)$ is increasing, it follows that $q_e\left(f(\varepsilon)_e\right) < q_e\left(f_e\right)$ for each $e \in \tilde{\tilde{E}}$. Therefore, since by definition $q_e\left(f_e\right) = B(f)$ for each $e \in \tilde{\tilde{E}}$, it follows that $q_e\left(f(\varepsilon)_e\right) < B(f)$ for each $e \in \tilde{\tilde{E}}$; in particular,

$$q_e\left(f(\varepsilon)_e\right) < B(f) \text{ for each } e \in \tilde{\tilde{E}} \text{ that satisfies } f(\varepsilon)_e^{\tilde{u}} > 0. \quad \text{(A6)}$$

Consider now the links $\tilde{E} \setminus \tilde{\tilde{E}}$. We just need to show that $q_e\left(f(\varepsilon)\right) < B(f)$ for each $e \in \tilde{E} \setminus \tilde{\tilde{E}}$ that satisfies $f(\varepsilon)_e^{\tilde{u}} > 0$. To that end, note that, by definition, the links of $\tilde{E} \setminus \tilde{\tilde{E}}$ are the links in $\tilde{E}$ that satisfy $\widetilde{g_e^{\tilde{u}}} = 0$. Hence, since $\varepsilon > 0$, it follows that, if $f_e^{\tilde{u}} > 0$, then it holds that $f(\varepsilon)_e^{\tilde{u}} = \varepsilon \cdot \widetilde{g_e^{\tilde{u}}} + (1-\varepsilon) \cdot f_e^{\tilde{u}} = (1-\varepsilon) \cdot f_e^{\tilde{u}} < f_e^{\tilde{u}}$ for each $e \in \tilde{E} \setminus \tilde{\tilde{E}}$. Thus, following the same arguments as in the previous case (i.e., the case of links from $\tilde{\tilde{E}}$) it is easy to see that

$$q_e\left(f(\varepsilon)_e\right) < B(f) \text{ for each } e \in \tilde{E} \setminus \tilde{\tilde{E}} \text{ that satisfies } f(\varepsilon)_e^{\tilde{u}} > 0. \quad \text{(A7)}$$

Therefore, from (A5)-(A7), we conclude that $q_e\left(f(\varepsilon)_e\right) < B(f)$ for each link $e \in E$ that satisfies $f(\varepsilon)_e^{\tilde{u}} > 0$. Hence, by definition, $b_{\tilde{u}}\left(f(\varepsilon)\right) < B(f)$; finally, as explained, since $b_{\tilde{u}}(f) = B(f)$ for the user $\tilde{u} \in \Gamma(f)$, it follows that $b_{\tilde{u}}\left(f(\lambda)\right) < b_{\tilde{u}}(f)$. As this conclusion contradicts the fact that f is at Nash equilibrium, we deduce that g is a Nash flow. Thus, the Lemma is established. ■

A.3 *Why Lemma 1 does not hold for unsplittable flows?*

The proof of the lemma breaks at the point that the flow vector $f(\varepsilon)$ is introduced. Indeed, for the user $\tilde{u}$, the vector $f(\varepsilon)$ is a superposition of two different unsplittable flow vectors namely, $f(\varepsilon)_p^{\tilde{u}} = \varepsilon \cdot \widetilde{g_p^{\tilde{u}}} + (1-\varepsilon) \cdot f_p^{\tilde{u}}$ for each $p \in P$, Obviously, if f and g ship the traffic of user $\tilde{u}$ on two different paths then, in the flow vector $f(\varepsilon)$, the traffic of $\tilde{u}$ is *split* among the corresponding paths.