Bottleneck Routing Games in Communication Networks

Ron Banner and Ariel Orda
Department of Electrical Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
{banner@tx, ariel@ee}.technion.ac.il

Abstract*—We consider routing games where the performance of each user is dictated by the worst (bottleneck) element it employs. We are given a network, finitely many (selfish) users, each associated with a positive flow demand, and a load-dependent performance function for each network element; the social (i.e., system) objective is to optimize the performance of the worst element in the network (i.e., the network bottleneck). Although we show that such "bottleneck" routing games appear in a variety of practical scenarios, they have not been considered yet. Accordingly, we study their properties, considering two routing scenarios, namely when a user can split its traffic over more than one path (splittable bottleneck game) and when it cannot (unsplittable bottleneck game). First, we prove that, for both splittable and unsplittable bottleneck games, there is a (not necessarily unique) Nash equilibrium. Then, we consider the rate of convergence to a Nash equilibrium in each game. Finally, we investigate the efficiency of the Nash equilibria in both games with respect to the social optimum; specifically, while for both games we show that the price of anarchy is unbounded, we identify for each game conditions under which Nash equilibria are socially optimal.

Keywords- Bottleneck Objectives; Selfish Routing; Nash Equilibrium; Price of Anarchy; Unregulated Traffic.

I. INTRODUCTION

Traditional communication 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. However, it has been recognized that systemwide optimization may be an impractical paradigm for the control of modern networking configurations [2],[3],[17],[18],[21],[24],[27],[31],[34]. Indeed, control decisions in large scale networks are often made by each user independently, according to its own individual interests. Such networks are henceforth called *noncooperative*, and Game Theory [25] provides the systematic framework to study and understand their behavior.

Game theoretic models have been employed in various networking contexts, such as flow control [1],[18],[34], routing [3],[14],[17],[21],[24],[31] and bandwidth allocation [22]. These studies mainly investigated the structure of the network operating points i.e., the Nash equilibria of the respective games. Such equilibria are inherently inefficient [13] and, in general, exhibit suboptimal network performance. As a result, the question of how much worse the quality of a Nash equilibrium is with respect to a centrally enforced optimum has received considerably attention e.g., [12],[20],[30],[31]. In order to quantify this inefficiency, two conceptual measures have been proposed in the literature. The first, termed the *price of anarchy* [26], corresponds to a worst-case analysis and it is the ratio between the *worst* Nash equilibrium and the social optimum. The second, termed the *price of stability* [4] is the ratio between the *best* Nash equilibrium and the optimum, and it quantifies the degradation in performance when the solution is required to be *stable* (i.e., with no agent having an incentive to independently defect out of it once being there).

The above studies focused on the case where the structure of the user performance objective is *additive* i.e., performance is determined by the *sum* of link cost functions. Yet, another fundamental case is that of *bottleneck objectives* (also known as *Max-Min* or *Min-Max* objectives), in which performance is determined by the *worst* component (link). Accordingly, in this study we investigate the case where users route traffic selfishly so as to optimize the performance of *their* bottleneck elements, given the routing strategies of all other users. Such settings give rise to a non-cooperative game, which is henceforth termed *the bottleneck game*.

Bottleneck games emerge in many practical scenarios. One major framework is wireless networks, where each node has a limited battery (i.e., transmission energy) so that the node's lifetime depends on the total flow that emanates from it. In such settings, the social (i.e., system) objective is to maximize the minimum battery lifetime in the network [10],[35], while each user would route traffic so as to maximize the smallest battery lifetime along *its* routing topology (hence, maximizing its connection's lifetime). Bottleneck games also arise when users attempt to enhance their ability to accommodate momentary traffic bursts. In these cases, the users aim at maximizing the smallest residual capacity of the links they employ (while the social objective is to maximize the minimum residual capacity of a link in the network [11]). Traffic engineering is another major framework where bottleneck games are encountered. For example, in view of the limited size of transmission buffers, each of the users is interested in minimizing the utilization of its most utilized buffer in order to avoid deadlocks and reduce packet loss [7]. Similarly, in congested networks it is often desirable to minimize the utilization of the most utilized links so as to move traffic away from congested hot spots to less utilized parts of the network [5],[33]. Other scenarios where bottleneck games appear can be found, for example, in frameworks where users attempt to enhance the ability to survive malicious attacks. Since such attacks are naturally aimed against the links (or nodes) that carry the largest amount of traffic, each user would be interested in minimizing the maximum amount of traffic that a link transfers in its routing topology.

As mentioned, in spite of their fundamental importance, bottleneck games have not been considered in the literature. Accordingly, in this study we consider two classes of bottleneck games. In the first, each user can split its traffic among any number of paths (henceforth, *splittable bottleneck game*), while in the second, each user routes its traffic along a single path (henceforth, *unsplittable bottleneck game*).

A. *Our Results*

First, we prove that, for both splittable and unsplittable bottleneck games, there is a (not necessarily unique) Nash equilibrium; we note that, for the splittable case, a major complication is the inherent discontinuity of the objective functions. Then, we turn to consider the rate of convergence of *best response dynamics* in both games. For unsplittable bottleneck games, we show that convergence to a Nash equilibrium is always achieved within a finite number of steps; moreover, when the number of users is small (i.e., $O(1)$), convergence time is polynomial in the network size. For splittable bottleneck games, we show that the convergence time may be unbounded.

Next, we investigate the efficiency of the Nash equilibrium. Specifically, for each game we compare the network bottleneck (i.e., the performance of the worst element in the network) at Nash equilibrium with that of an optimal flow (i.e., a flow that minimizes the network bottleneck). In particular, for both games, we analyze the price of anarchy and identify conditions under which Nash equilibria are efficient. Considering first the unsplittable case, we show that the price of anarchy may be unbounded[1]. Yet, we show that a *best* Nash equilibrium coincides with the *social optimum*; hence, *the price of stability is 1*. We also show that calculating such a best equilibrium is NP-hard. Consider now the efficiency of Nash equilibria in splittable bottleneck games. Here too, we show that the price of anarchy is unbounded. Yet, we show that, if at a Nash equilibrium the users route their traffic along paths with a minimum number of bottleneck links, the Nash equilibrium is socially optimal. This finding might motivate the employment of pricing mechanisms that penalize the use of paths with an excessive number of bottleneck links.

B. *Organization*

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 bottleneck games. In section 4, we consider the convergence properties of each game. In section 5, we investigate the price of anarchy of both games and identify for each game conditions under which Nash equilibria are socially optimal. Finally, we conclude the paper in section 6. For ease of presentation most proofs appear in an appendix.

[1] Yet, in the special case of link performance functions that are polynomial with a degree p, we show that the price of anarchy is $O\left(M^p\right)$, where M is the number of network links; we also show that this result is tight.

II. MODEL AND TERMINOLOGY

We consider a finite set of users U, which share a communication network that is modeled by a graph $G(V,E)$. Each user $u \in U$ is associated with a positive throughput demand γ_u and a pair of source-destination nodes (s_u,t_u). We denote by $P^{(s_u,t_u)}$ the set of all paths from the source s_u to the destination t_u and by P the set of all the paths in the network; assume that the source and destination nodes of each user are connected i.e., $\left|P^{(s_u,t_u)}\right| \geq 1$ for each $u \in U$.

A user $u \in U$ needs to send γ_u units of flow from s_u to t_u along the paths $P^{(s_u,t_u)}$. We denote by f_p^u the flow of user $u \in U$ on a path $p \in P^{(s_u,t_u)}$. A user u can assign any value to f_p^u, as long as $f_p^u \geq 0$ (nonnegativty constraint) and $\sum_{p\in P^{(s_u,t_u)}} f_p^u = \gamma_u$ (demand constraint); this assignment of traffic to paths shall also be referred to as the *user strategy*. The set of all possible strategies of a user is referred to as the *user strategy space*. The product of all user strategy spaces is termed the *joint strategy space*; each element in the joint strategy space is termed a (flow vector) *profile*; effectively it is a global assignment of traffic to paths that satisfies the demands of all users.

Given a profile $f=\left\{f_p^u\right\}$ and a path $p \in P$, denote by f_p the total flow that is carried over p i.e., $f_p = \sum_{u\in U} f_p^u$; also, denote by f_e^u the total flow that user u transfers through e i.e., $f_e^u = \sum_{p|e\in p} f_p^u$. Finally, for a flow vector f and a link $e \in E$, denote the total flow carried by e as f_e.

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 continuous and increasing in f_e. We define the *network bottleneck* $B(f)$ of a flow f as the performance of the worst link in the network i.e., $B(f) \triangleq \max_{e\in E}\left\{q_e(f_e)\right\}$. Similarly, we define the *bottleneck of a user* $u \in U$ as the performance of the worst link that u employs i.e., $b_u(f) \triangleq \max_{e\in E|f_e^u>0}\left\{q_e(f_e)\right\}$.

Users are selfish i.e., each minimizes its *own* bottleneck. Since the bottleneck of a user $u \in U$ depends on the flow configuration of *all* users, we are faced with a *non-cooperative game* [25]. We consider two types of games. The triple $\left\langle G,U,\{q_e\}\right\rangle$ is termed a *splittable bottleneck game* if each user may route its traffic along more than a single path; it is termed an *unsplittable bottleneck game* if each user can employ only a single path. Similarly, a profile $f = \{ f_p^u | u \in U, p \in P \}$ is said to be an *unsplittable flow vector* if $f_p^u \in \{0, \gamma_u\}$ for each $u \in U$ and $p \in P$; otherwise, *f* is said to be a *splittable flow vector*.

A profile is said to be at *Nash equilibrium* if each user considers its chosen strategy to be the best under the given choices of other users. More formally, considering a splittable (alternatively, unsplittable) bottleneck game, $f = \{ f_p^u \}$ is a *splittable (correspondingly, unsplittable) Nash flow* if, for each user $\tilde{u} \in U$ and splittable (correspondingly, 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 $b_{\tilde{u}}(f) \le b_{\tilde{u}}(g)$.

Remark 1: In our model, a higher bottleneck value means poorer performance. Clearly, all the results of this study can be adapted to the case where the link performance functions $\{q_e(f_e)\}$ are continuous and *decreasing* in the flows $\{f_e\}$, and the goal (of the users and the network) is to *maximize* the (corresponding) bottleneck values.

III. Existence and Non-Uniqueness of Nash Equilibria

In this section we establish the existence and non-uniqueness of a Nash equilibrium in splittable and unsplittable bottleneck games. We begin with the splittable case.

In general, there are several standard techniques to establish the existence of a Nash equilibrium in infinite games (e.g., Kakutani's fixed-point theorem [19] and its generalization [16], Rosen's theorem [29] and Debreu's theorem [13]). However, these techniques cannot be employed for splittable bottleneck games, since the utility function of each user (i.e., $b_u(\cdot)$) is *discontinuous*, as shown by the following example.

Specifically, we now show that, since the *max* operation in the user utility function $b_u(f) = \max_{e \in E | f_e^u > 0} \{ q_e(f_e) \}$ is defined over a *strategy-dependent* set, it is *discontinuous* in the user's strategy space. Indeed, for a single user that needs to transfer one unit of flow from *s* to *t* in the network of Fig. 1, both $g = \langle f_{e_1} = \varepsilon, f_{e_2} = 1 - \varepsilon \rangle$ and $h = \langle f_{e_1} = 0, f_{e_2} = 1 \rangle$ are feasible flow vectors that satisfy $\|g - h\| = \sqrt{2} \cdot \varepsilon$. However, while $\|g - h\| \to 0$ for $\varepsilon \to 0$, it holds that $b_u(g) - b_u(h) > 1$ for all $\varepsilon > 0$.

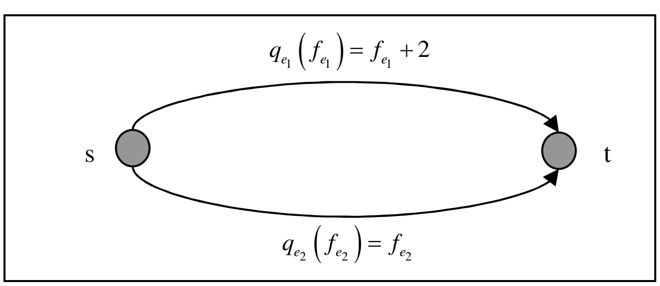

Fig. 1**:** Discontinuous user objectives in splittable bottleneck games

Consequently, we proceed to construct an existence proof that does not rely on the continuity of the cost functions. Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$, denote for each user $u \in U$ the set of strategies available to *u* by F_u and let $F \triangleq \times_{u \in U} F_u$. In addition, let F_{-u} denote the set of strategies of all users other than *u* i.e., $F_{-u} \triangleq \times_{w \in U \setminus u} F_w$. Finally, for each user $u \in U$, let the pair $(f_u, f_{-u}) \in F_u \times F_{-u}$ denote a profile in *F* with a strategy of f_u for the user *u* and a strategy profile f_{-u} for all other users.

Definition 1: A splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$ is *compact* if each F_u is a nonempty compact subset of a topological vector space and each user bottleneck function $b_u(\cdot)$ is bounded in *F*.

Lemma 1: Splittable bottleneck games are compact.

The proof Appears in Appendix A.1.

Definition 2: A user *u* can *secure* a bottleneck $b \in \mathbb{R}$ at a profile $f \in F$ if there exists a reply $\overline{f_u} \in F_u$, such that $b_u(\overline{f_u}, \widetilde{f_{-u}}) \le b$ for all $\widetilde{f_{-u}}$ in some neighborhood of f_{-u}.

In other words, a bottleneck can be *secured* by a user *u* at $f \in F$ if *u* has a strategy that guarantees that bottleneck even if the other users slightly deviate from *f*.

For the next definition, we introduce the following notation. Given a profile $f \in F$, denote the vector of all user bottlenecks by $\overline{b(f)}$ i.e., $\overline{b(f)} \triangleq \left(b_{u_1}(f), b_{u_2}(f), \cdots b_{u_{|U|}}(f) \right)$. The *set of the user bottleneck vectors* is the subset of $F \times \mathbb{R}^{|U|}$ given by $\left\{ \left(f, \overline{b(f)} \right) | f \in F \right\}$.

Definition 3: Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$, the pair $(f^*, b^*) \in F \times \mathbb{R}^{|U|}$ is in the *closure* of the user bottleneck vector set if there exists a sequence of profiles $\{f^n\} \subseteq F$ such that $(f^*, b^*) = \lim_{n \to \infty} \left(f^n, \overline{b(f^n)} \right)$.

Considering Definition 3, we show in the proof of lemma 2 that $f^* \in F$ and therefore, $b_u(f^*)$ is well-defined for all *u*; however, it is important to note that, since the user bottleneck function is discontinuous, it is possible that $b^* \ne \overline{b(f^*)}$.

Definition 4: A splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$ is *better-reply secure* if, whenever (f^*, b^*) is in the closure of the user bottleneck vector set, and f^* is not an equilibrium, some user $u \in U$ can secure a bottleneck strictly below b_u^* at f^*.

In other words, a game is *better-reply secure* if, for every non-equilibrium profile $f^* \in F$ and sequence $\{f^n\} \subseteq F$ such that $(f^*, b^*) = \lim_{n \to \infty} \left(f^n, \overline{b(f^n)} \right)$, there exists a user $u \in U$ that can attain a bottleneck below b_u^* even if the others slightly deviate from f^*.

Lemma 2: Splittable bottleneck games are better-reply secure.

The proof Appears in Appendix A.1.

Definition 5: A function $f : F \rightarrow \mathbb{R}$ is said to be *quasi-convex* on F if, for a each $\alpha \in \mathbb{R}$, the set $\left\{ x \in F \middle| f(x) \le \alpha \right\}$ is convex.

Definition 6: A splittable bottleneck game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ is *quasi-convex* if, for each $u \in U$, the set F_u is convex and, for every $f_{-u} \in F_{-u}$, the payoff function $b_u(\ \cdot\ , f_{-u})$ is quasi-convex on F_u.

Lemma 3: Splittable bottleneck games are quasi-convex.

The proof Appears in Appendix A.1.

We are finally ready to prove the existence of a Nash equilibrium in a splittable bottleneck game.

Theorem 1: A splittable bottleneck game admits a Nash equilibrium.

Proof: In [28] it is shown that, if a game *G* is better reply secure, quasi-convex and compact, then it possesses a Nash equilibrium. This, together with Lemmas 1- 3, establish the Theorem. ■

For *unsplittable* bottleneck games, the existence of a Nash equilibrium follows directly from Theorem 3 (see Section V.A.2).

Corollary 1: An unsplittable bottleneck game admits a Nash equilibrium.

Finally, we now show, by way of an example, that a Nash equilibrium may not be unique both in splittable and unsplittable bottleneck games.

Consider the network presented in Figure 2. 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, for both splittable and unsplittable bottleneck games, the flow vectors $\langle f_{p_1} = 1, f_{p_2} = 0 \rangle$ and $\langle f_{p_1} = 0, f_{p_2} = 1 \rangle$ minimize the user bottleneck. For the case of a *single-user game*, each of these is a Nash flow.

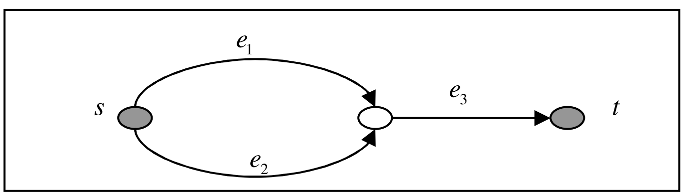

Fig. 2**:** Non-uniqueness of flow at Nash Equilibrium

IV. CONVERGENCE TO NASH EQUILIBRIUM

In this section we address the convergence properties of *best response dynamics* in splittable and unsplittable bottleneck games. We shall focus on the *Elementary Stepwise System* (ESS) [24], in which players update their actions sequentially, and at each update a player uses the best response action given the actions of the other players[1]. We begin with the unsplittable case.

Theorem 2: In an unsplittable bottleneck game with $|U|$ users, starting out of any flow configuration, the ESS scheme converges to a Nash equilibrium within at most $2^{|U|^2} \cdot |E|^{|U|}$ steps.

The proof appears in Appendix A.2.

In particular, for a number of users that is constant in the size of the network, the number of steps required to reach a Nash equilibrium is polynomial. Hence, since the best response action of each user (i.e., the establishment of a path with an optimal bottleneck) can be computed in polynomial time [23], for $O(1)$ users the ESS scheme provides a polynomial-time algorithm for computing a Nash equilibrium.

For splittable bottleneck games, it is shown in Appendix A.2 that convergence time is unbounded even for two users.

V. (IN)EFFICIENCY OF NASH EQUILIBRIA

In this section, we investigate the degradation in network performance due to the selfish behavior of users. For both games we show that the price of anarchy is unbounded. Yet, for each game we identify conditions under which Nash equilibria are socially optimal. We begin with the unsplittable case.

A. *(In)Efficiency of unsplittable Nash equilibria*

We begin with a simple example that shows that, in general, unsplittable bottleneck games have unbounded price of anarchy.

Example 1: Consider the network presented in Fig. 3, 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 2γ 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* $B(f)/B(f^*) = e^{\frac{\gamma}{3}}$; *as γ can be arbitrarily large, this ratio is unbounded.*

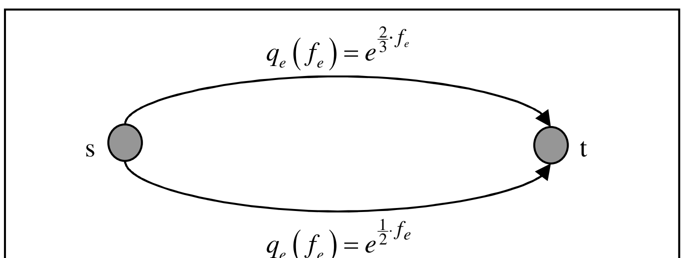

Fig. 3: Unbounded price of anarchy for unsplittable bottleneck games

[1] We note that, if all players are allowed to move simultaneously, the system might oscillate and never reach a Nash equilibrium.

In Appendix A.3 we show that, when the link performance functions are polynomial with a degree *p*, the price of anarchy is $O\left(|E|^p\right)$; we also show that the latter result is tight.

While we have shown that in unsplittable bottleneck games the price of anarchy is unbounded, we now show that the price of stability is 1. Thus, although in the worst case unsplittable Nash equilibria can be very inefficient, every unsplittable bottleneck game has at least one Nash equilibrium that optimizes the network bottleneck. Hence, finding such an equilibrium provides an *optimal* solution that is also *stable*, with no user having an incentive to discard it once adopted by all other users.

Theorem 3: Given a flow vector f^* for the unsplittable bottleneck game $\langle G, U, \{q_e\}\rangle$, there exists a Nash flow *f* such that $B(f) \le B(f^*)$; hence, the price of stability is 1.

The proof appears in Appendix A.3.

The existence of "good" Nash equilibria is of major importance in the design of efficient protocols. Indeed, in many networking applications, a collective solution is proposed to all users. Therefore, it is in the interest of the protocol designer to seek the best solution that selfish users can agree upon i.e., the best Nash equilibrium. Unfortunately, while Theorem 3 guarantees that, in an unsplittable bottleneck game, there exists a Nash equilibrium that is socially optimal, the following theorem shows that computing it is NP-hard.

Theorem 4: Given an unsplittable bottleneck game $\langle G(V,E), U, \{q_e\}\rangle$ and a value *B*, it is NP-hard to determine if the game has a Nash equilibrium with a bottleneck of at most *B*.

The proof of the theorem appears in Appendix A.3, and is based on a reduction from the *Disjoint Connecting Paths Problem* [15].

B. (In)Efficiency of splittable Nash equilibria

We now turn to consider the efficiency of Nash equilibria in *splittable* bottleneck games. Again, we first quantify the price of anarchy and then provide conditions under which Nash equilibria are optimal.

1) The price of anarchy of splittable bottleneck games

The following example shows that the price of anarchy of splittable bottleneck games is unbounded.

Example 2: Consider the network depicted in Fig. 4. Assume that each link $e \in E$ *has a performance function* $q_e(f_e) = 2^{f_e}$. *There are two users A and B, each with a flow demand of* γ *units. The source-destination pairs of A and B are* (s_1,t_1) *and* (s_2,t_2)*, correspondingly. Let* $p_1 = (s_1, s_2, u, t_1)$, $p_2 = (s_1, v, t_2, t_1)$, $p_3 = (s_2, u, v, t_2)$ *and* $p_4 = (s_2, v, t_2)$. *It is easy to see that the path flow* $\left\langle f_{p_1} = \frac{\gamma}{2}, f_{p_2} = \frac{\gamma}{2}, f_{p_3} = \gamma, f_{p_4} = 0 \right\rangle$ *is at Nash equilibrium with a (network) bottleneck of* $2^{\frac{3\cdot\gamma}{2}}$*, while the path flow* $\langle f_{p_1} = \gamma, f_{p_2} = 0, f_{p_3} = 0, f_{p_4} = \gamma\rangle$ *is an optimal solution with a bottleneck of* 2^{γ}. *Therefore, the ratio between the bottlenecks is* $2^{\frac{\gamma}{2}}$*; as* γ *can be arbitrarily large, the price of anarchy is unbounded.*

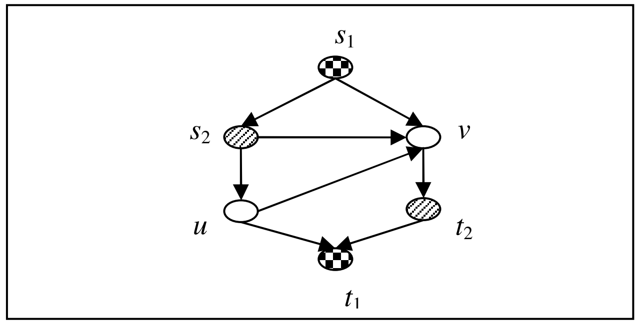

Fig. 4: Unbounded price of anarchy for splittable bottleneck games

2) When are splittable Nash equilibria optimal?

Note that in the (inefficient) Nash equilibrium of Example 2, user B is not routing along paths with a minimum number of bottlenecks; indeed, user B could shift its flow demand from the path p_3 (that contains two bottlenecks) to the path p_4 (that contains only one bottleneck) without affecting its performance. Thus, B unnecessarily ships traffic through an excessive number of network bottlenecks and, as a result, affects user A that shares one of these bottlenecks. On the other hand, in the optimal solution $\langle f_{p_1} = \gamma, f_{p_2} = 0, f_{p_3} = 0, f_{p_4} = \gamma\rangle$, users A and B are at Nash equilibrium and each routes its traffic along paths with a minimum number of bottlenecks.

In the following, we generalize the above observation to all splittable bottleneck games. More specifically, we prove that the (network) bottleneck of a splittable Nash equilibrium is optimal if the users route their traffic along paths that consist of a minimum number of bottlenecks. To that end, we need the following definition.

Definition 7: Given a Nash flow *f* for the splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$ and a path $p \in P$, denote the number of network bottlenecks over *p* by $N_f(p)$ i.e., $N_f(p) \triangleq \left|\{e \in p \mid q_e(f) = B(f)\}\right|$. Then, *f* is said to *satisfy the efficiency condition* if all users route their traffic along paths with a minimum number of bottlenecks i.e., for each $u \in U$ and $p_1, p_2 \in P^{(s_u, t_u)}$ with $f_{p_1}^u > 0$ it holds that $N_f(p_1) \le N_f(p_2)$.

In the reminder of this section, we show that every splittable Nash flow that satisfies the efficiency conditions is a social optimum. To that end, we introduce the notion of *induced flow*, which is defined as follows.

Definition 8: Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$ and a feasible flow vector *f*, denote by $\Gamma(f)$ the set of users that ship traffic through a network bottleneck i.e., $\Gamma(f) = \{u \in U \mid b_u(f) = B(f)\}$. Then, the flow vector $g: P \times \Gamma(f) \to \mathbb{R}^+ \cup \{0\}$ is *induced by f iff* for each $p \in P$ it holds that $g_p^u = \begin{cases} f_p^u & u \in \Gamma(f) \\ 0 & \text{else} \end{cases}$.

In other words, given a flow vector *f*, the induced flow *g* is attained by deleting and zeroing the flows of all users that are not shipping traffic through any network bottleneck. The following lemma claims that an induced flow of a splittable Nash flow that satisfies the efficiency condition is at Nash equilibrium (with respect to the new game that corresponds to the reduced set of users); moreover, the new flow satisfies the efficiency condition and keeps the bottleneck of each user $u \in \Gamma(f)$ equal to that of the network.

Lemma 4: Given a splittable bottleneck game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ for which *f* is a Nash flow that satisfies the efficiency condition, the induced flow *g* satisfies the demands of all users in $\Gamma(f)$, such that the following two properties hold:

(i) *g* is a Nash flow that satisfies the efficiency condition;

(ii) the bottleneck of each user $u \in \Gamma(f)$ in the flow vector *g* equals to that of the network i.e., $b_u(g)=B(g)$ for each $u \in \Gamma(f)$.

The proof of Lemma 4 appears in Appendix A.4. While the details are rather tedious, the basic idea is quite straightforward, namely, the removal of users that are not in $\Gamma(f)$ cannot change the flow on any of the bottleneck links in the network.

Lemma 5: Given is a splittable bottleneck game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, for which *f* is a Nash flow that satisfies the efficiency condition; let *g* be the flow that is induced by *f*. Then, *g* has the smallest network bottleneck among all splittable flow vectors that satisfy the demands of the users in $\Gamma(f)$ i.e., $B(g) \leq B(h)$ for each flow vector *h* that has $\sum_{p \in P^{(s_u,t_u)}} h_p^u = \gamma_u$ for each $u \in \Gamma(f)$.

Proof: For simplicity, we assume that the traffic demand between each source-destination pair belongs to at most one user. It is easy to see[1] that this assumption incurs no loss of generality for splittable flows.

Let *f* and *g* satisfy the hypothesis of the theorem. Assume, by way of contradiction, that the bottleneck $B(g)$ is not the minimum among the (network) bottlenecks of all (splittable) flow vectors that satisfy the demands of the users in $\Gamma(f)$. Hence, there exists a feasible flow vector *h* that satisfies the demands of all users in $\Gamma(f)$ such that $B(h) < B(g)$. Denote by $\overline{E}$ the set of all network bottlenecks with respect to *g*. In addition, for each $e \in \overline{E}$, denote by $P(e)$ the set of all paths that traverse through *e*.

Since *g* is induced by the Nash flow *f* it follows according to Lemma 4 that *g* is at Nash equilibrium. Therefore, *every* path $p \in P^{(s_u,t_u)}$, where $u \in \Gamma(f)$, must traverse through at least one network bottleneck from $\overline{E}$; indeed, otherwise a user $u \in \Gamma(f)$ could have decreased its bottleneck value (which, according to Lemma 4 equals to the network bottleneck $B(g)$) by rerouting flow into paths whose bottlenecks are lower. Therefore, $P^{(s_u,t_u)} \subseteq \bigcup_{e \in \overline{E}} P(e)$ for each $u \in \Gamma(f)$.

Since we assume that $B(h) < B(g)$, it follows that $q_e(h_e) < q_e(g_e)$ for each $e \in \overline{E}$. Therefore, since $q_e(\cdot)$ is increasing, it holds that $h_e < g_e$ for each $e \in \overline{E}$. Therefore, for each $e \in \overline{E}$, the total traffic that traverses the paths $P(e)$ is smaller in *h* than in *g* i.e., $\sum_{p \in P(e)} h_p < \sum_{p \in P(e)} g_p$ for each $e \in \overline{E}$. Therefore,

$$\sum_{e \in \overline{E}} \sum_{p \in P(e)} h_p < \sum_{e \in \overline{E}} \sum_{p \in P(e)} g_p \quad (1)$$

Recall that, by Definition 7, the number of links that belong to the set $\overline{E}$ and a path $p \in P$ is $N_g(p)$ i.e., $N_g(p)=\left|p \cap \overline{E}\right|$. Thus, $\sum_{e \in \overline{E}} \sum_{p \in P(e)} h_p = \sum_{p \in \bigcup_{e \in \overline{E}} P(e)} h_p \cdot N_g(p)$ and $\sum_{e \in \overline{E}} \sum_{p \in P(e)} g_p = \sum_{p \in \bigcup_{e \in \overline{E}} P(e)} g_p \cdot N_g(p)$. Therefore, from (1), $\sum_{p \in \bigcup_{e \in \overline{E}} P(e)} h_p \cdot N_g(p) < \sum_{p \in \bigcup_{e \in \overline{E}} P(e)} g_p \cdot N_g(p)$. Thus, if we define $\Pi(u) \triangleq \left(\bigcup_{e \in \overline{E}} P(e)\right) \cap P^{(s_u,t_u)}$ for each $u \in \Gamma(f)$, then

$$\sum_{p \in \bigcup_{e \in \overline{E}} P(e)} (g_p - h_p) \cdot N_g(p) > 0 \Rightarrow \sum_{u \in \Gamma(f)} \sum_{p \in \Pi(u)} (g_p - h_p) \cdot N_g(p) > 0 \Rightarrow$$

$$\Rightarrow \sum_{u \in \Gamma(f)} \left(\sum_{p \in \Pi(u) | g_p > 0} (g_p - h_p) \cdot N_g(p) + \sum_{p \in \Pi(u) | g_p = 0} (g_p - h_p) \cdot N_g(p) \right) > 0 \Rightarrow$$

$$\Rightarrow \sum_{u \in \Gamma(f)} \left(\sum_{p \in \Pi(u) | g_p > 0} (g_p - h_p) \cdot N_g(p) - \sum_{p \in \Pi(u) | g_p = 0} h_p \cdot N_g(p) \right) > 0.$$

Since the given flow vector g is induced by a Nash flow that satisfies the efficiency condition, it holds (by Lemma 4) that the efficiency condition is satisfied for *g* as well. Hence, by definition $N_g(p_1) \leq N_g(p_2)$ for each $u \in \Gamma(f)$ and $p_1, p_2 \in P^{(s_u,t_u)}$ with $g_{p_1}^u > 0$; in particular, $N_g(p_1) = N_g(p_2) \triangleq N^{(s_u,t_u)}$ for each $p_1, p_2 \in P^{(s_u,t_u)}$ with $g_{p_1}^u > 0$ and $g_{p_2}^u > 0$. Hence, $\sum_{u \in \Gamma(f)} \left(\sum_{p \in \Pi(u) | g_p > 0} (g_p - h_p) \cdot N^{(s_u,t_u)} - \sum_{p \in \Pi(u) | g_p = 0} h_p \cdot N_g(p) \right) > 0$; on the other hand, the efficiency condition guarantees that $N^{(s_u,t_u)} \leq N_g(p)$ for each path $p \in \Pi(u)$ with $g_p = 0$. Thus, it holds that $\sum_{u \in \Gamma(f)} \left(\sum_{p \in \Pi(u) | g_p > 0} (g_p - h_p) \cdot N^{(s_u,t_u)} - \sum_{p \in \Pi(u) | g_p = 0} h_p \cdot N^{(s_u,t_u)} \right) > 0$; consequently, $\sum_{u \in \Gamma(f)} N^{(s_u,t_u)} \cdot \left(\sum_{p \in \Pi(u) | g_p > 0} (g_p - h_p) + \sum_{p \in \Pi(u) | g_p = 0} (g_p - h_p) \right) > 0$; therefore, $\sum_{u \in \Gamma(f)} N^{(s_u,t_u)} \cdot \sum_{p \in \Pi(u)} (g_p - h_p) > 0$.

[1] By adding fictitious sources and destinations and zero-cost fictitious links.

Finally, since we have shown that $P^{(s_u,t_u)} \subseteq \bigcup_{e \in \bar{E}} P(e)$, it holds that $\Pi(u) = \left(\bigcup_{e \in \bar{E}} P(e)\right) \cap P^{(s_u,t_u)} = P^{(s_u,t_u)}$ for each $u \in \Gamma(f)$. Therefore, it is satisfied for each user $u \in \Gamma(f)$ that

$$\sum_{u \in \Gamma(f)} N^{(s_u,t_u)} \cdot \sum_{p \in P^{(s_u,t_u)}} (g_p - h_p) > 0 \cdot \qquad (2)$$

Yet, as $\sum_{p \in P^{(s_u,t_u)}} g_p = \sum_{p \in P^{(s_u,t_u)}} h_p = \gamma_u$ it holds that $\sum_{p \in P^{(s_u,t_u)}} (g_p - h_p) = 0$ for each $u \in \Gamma(f)$; hence, $\sum_{u \in \Gamma(f)} N^{(s_u,t_u)} \cdot \sum_{p \in P^{(s_u,t_u)}} (g_p - h_p) = 0$. Obviously, the latter contradicts (2). ■

The following lemma establishes a useful relation between any splittable flow vector *f* and the flow that is induced by *f*.

Lemma 6: Given a flow vector *f* for the splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$, denote by g the flow that is induced by *f*. If g has the smallest network bottleneck among the flow vectors that satisfy the demands of the users in $\Gamma(f)$, then *f* has the smallest network bottleneck among the flow vectors that satisfy the demands of the users in *U*.

The proof of Lemma 6 appears in Appendix A.4. We are finally ready to introduce the main result of this section, which follows directly from Lemmas 5 and 6.

Corollary 2: Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\} \rangle$, a feasible flow vector f^* and a splittable Nash flow *f* that satisfies the efficiency condition, it holds that $B(f) \le B(f^*)$; hence, a splittable Nash equilibrium that satisfies the efficiency condition is socially optimal.

Remark 4: The Braess paradox [9] shows that addition of links to a noncooperative network can negatively impact the performance of both the network and each of the users. However, for the case that $q_e(0) = 0$ for each $e \in E$, it holds that the bottleneck of the network is also a bottleneck of some user in *U*. Hence, for this case, it follows from Theorem 5 that the efficiency condition is sufficient to guarantee that the paradox does not show up; indeed, at least the bottleneck of the "poorest" user (which equals the bottleneck of the network) is minimized.

Remark 5: In [6] we use the results of this section to show that, when the efficiency condition is satisfied, the *average* performance across all links at a splittable Nash equilibrium is at most $|E|$ times larger than the minimum value (obtained by a flow that minimizes the average link performance); we also provide an example that shows that this result is tight. On the other hand, we show that, without the efficiency condition, the average link performance of a splittable Nash flow may be arbitrarily larger than the optimum.

VI. Conclusion

Although bottleneck objectives emerge in many practical scenarios, the behavior of selfish users optimizing such objectives has not been considered before. Accordingly, in this study we investigated bottleneck games in noncooperative networks, both for the case where users can split their traffic between several paths, and for the case where they cannot. We have shown that a bottleneck game always admits a Nash equilibrium; moreover, best response dynamics in unsplittable games converge to a Nash equilibrium in finite time. Yet, we have shown that a Nash equilibrium (both in splittable and unsplittable bottleneck games) can be very inefficient. In order to cope with this inefficiency, we investigated for each game "reasonable" conditions under which Nash equilibria are socially optimal; these conditions suggest network design rules for ensuring that selfish behavior in bottleneck games result in a desirable outcome. Specifically, for unsplittable bottleneck games we have shown that a best Nash equilibrium is a social optimum, and for splittable bottleneck games we have shown that a Nash equilibrium is optimal if all users route their traffic along paths with a minimum number of bottlenecks. Accordingly, a major direction for future research is the establishment of pricing mechanisms that would steer users to choose particular Nash equilibria, or to incorporate additional (reasonable) criteria in the routing decision, thus optimizing the overall network performance.

References

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

[2] E. Altman and L. Wynter, *Equilibrium, games, and pricing in transportation and telecommunication networks*, Networks and Spatial Economics, vol.4, no. 1, 2003.

[3] E. Altman, T. Basar, T. Jimenez and N. Shimkin , *Competitive routing in networks with polynomial cost*, IEEE Transactions on Automatic Control, Vol 47, pp. 92-96, Jan. 2002.

[4] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden, *The Price of Stability for Network Design with Fair Cost Allocation*, In Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, October 2004.

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

[6] R. Banner and A. Orda, *Bottleneck Games in Noncooperative Networks*, CCIT Report No. 538, Department of Electrical Engineering, Technion, Haifa, Israel, 2005. Available from: http://www.ee.technion.ac.il/people/ron/Bottleneck.pdf

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

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

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

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

[11] H. S. Chang, B. W. Kim, C. G. Lee, S. L. Min, Y. Choi, H. S. Yang, D. N. Kim and C. S. Kim., *FSA-Based Link-Assignment and Routing in Low-Earth Orbit Satellite Networks* IEEE Transactions on Vehicular Technology, 1997.

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

[13] G. Debreu, *A Social Equilibrium Existence Theorem*, In Proceedings of the National Academy of Sciences, 38, 886-893, 1952.

[14] E. Even-Dar and Y. Mansour, *Fast Convergence of Selfish Rerouting*, In Proceedings of Sixteenth ACM-SIAM Symposium on Discrete Algorithms (SODA), 2005.

[15] M. R. Garey and D. S. Johnson, *Computers and Intractability*, W.H. Freeman and Co., 1979.

[16] I., L., Glicksberg, *A further Generalization of the Kakutani Fixed Point Theorem*, Proceedings of the American Mathematical Society, 3, 170-174, 1952.

[17] H. Kameda, *How harmful the paradox can be in Braess/Cohen-Kelly-Jeffries networks*, In Proceedings of IEEE INFOCOM 2002, New York, pp. 437-445, June 2002.

[18] 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.

[19] S. Kakutani, *A generalization of Brouwer's Fixed Point Theorem*, Duke Mathematical Journal, 8: 457-458, 1941.

[20] 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.

[21] R. La and V. Anantharam, *Optimal Routing Control: Repeated Game Approach*, IEEE Transactions on Automatic Control, Vol. 47, No. 3, pp. 437 -450, March 2002.

[22] A. A. Lazar, A. Orda and D. E. pendarakis, *Virtual path bandwidth allocation in multi-user networks*, IEEE/ACM Transactions on Networking, vol. 5, No. 6 pages 861-871, December 1997.

[23] Q. Ma and P. Steenkiste, *On path selection for traffic with bandwidth guarantees*, in IEEE International Conference on Network Protocols, 1997.

[24] A. Orda, R. Rom and N. Shimkin, *Competitive routing in multiuser communication networks*, IEEE/ACM Transactions on Networking vol. 1, pages 510-521, October 1993.

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

[26] C. Papadimitriou, *Algorithms, games and the Internet,* In proceedings of the 33rd Annual ACM Symposium on the Theory of Computing, Heraklion, Crete, Greece, July 2001.

[27] K. Park , M. Sitharam and S. Chen, *Quality of service provision in noncooperative networks: heterogenous preferences, multi-dimensional QoS vectors, and burstiness*, In proceedings of the first international conference on Information and computation economies, pp.111-127, October 1998, Charleston, South Carolina.

[28] P. Reny, *On the Existence of Pure and Mixed Strategy Nash Equilibrium in Discontinuous Games*, Econometrica, Vol 67, pp. 1029-1056, September 1999.

[29] J. B. Rosen, *Existence and Uniqueness of Equilibrium Points for Concave N-Person Games*, Econometria, Vol. 33, pp. 510-534, July 1965

[30] 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.

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

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

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

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

[35] G. Zussman and A. Segall, *Energy Efficient Routing in Ad Hoc Disaster Recovery Networks*, Proc. IEEE Infocom 2003, San Francisco, April 2003

APPENDICES

Appendix A.1: Existence of Nash Equilibrium in Splittable Bottleneck Games

In this Appendix we provide the proofs of Lemmas 1, 2 and 3 that are used to establish the existence of a Nash equilibrium in splittable bottleneck games.

Lemma 1: Splittable bottleneck games are compact.

Proof: Assume that we are given a splittable bottleneck game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$. Note that, since we assume that the source-destination nodes of all users are connected, each user has at least one strategy for routing its flow demand; hence, $F_u \neq \phi$ for each $u \in U$.

Next we show that, for each $u \in U$, the set F_u is a compact subset of $\mathbb{R}^{|P|}$. To that end, note that, for each $u \in U$, F_u is the space of solutions of the following system (A1)-(A3):

$$\sum_{p \in P^{(s_u,t_u)}} f_p^u = \gamma_u \qquad \text{(A1)}$$

$$f_p^u \geq 0 \qquad \forall p \in P^{(s_u,t_u)} \qquad \text{(A2)}$$

$$f_p^u = 0 \qquad \forall p \in P \setminus P^{(s_u,t_u)} \qquad \text{(A3)}$$

In order to prove that the set $F_u \subseteq \mathbb{R}^{|P|}$ is compact, it is sufficient to show that F_u is bounded and closed [32]. We first show that F_u is bounded. To that end, note that, for each strategy $f_u \in F_u$ in the solution space of (A1)-(A3), it holds that $0 \leq f_p^u \leq \gamma_u$ for each $p \in P$. Therefore, for each two strategies $h_u, g_u \in F_u$, it holds that $\|h_u - g_u\| \leq \sqrt{|P|} \cdot \gamma_u$. We turn to show that F_u is closed. To that end, note that the solution space of (A1)-(A3) is the surface of some polyhedron in $\mathbb{R}^{|P|}$, which is closed in $\mathbb{R}^{|P|}$ [32]. Therefore, the set $F_u \subseteq \mathbb{R}^{|P|}$ is compact.

It remains to be shown that the bottleneck function $b_u(\cdot)$ is bounded on F for each $u \in U$. To that end, note that $b_u(f) = \max_{e \in E | f_e^u > 0} \{q_e(f)\} \leq \max_{e \in E} \{q_e(f)\}$ for each $f \in F$. Moreover, since 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, it follows that the function $\max_{e \in E}\{q_e(\cdot)\}$ is continuous[1]. Therefore, since F is compact (as a finite Cartesian product of compact metric spaces) $\max_{e \in E}\{q_e(\cdot)\}$

[1] Although we have shown that the function $b_u(f) = \max_{e \in E | f_e^u > 0}\{q_e(f_e)\}$ is, in general, discontinuous in the strategy space, here we consider the function $\max_{e \in E}\{q_e(f_e)\}$, in which the *max* operation is defined over the (*strategy-independent*) set *E*.

attains the maximum over F [32]; hence, there exists an $M>0$ such that $b_u(f) \le \max_{e\in E}\{q_e(f)\} < M$ for each $u \in U$.■

Lemma 2: Splittable bottleneck games are better-reply secure.

Proof: Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$, let $(f^*, b^*) \in F \times \mathbb{R}^{|U|}$ be a pair in the closure of the user bottleneck vectors set. By definition, there exists a sequence $\{f^n\} \subseteq F$ such that $(f^*, b^*) = \lim_{n\to\infty}\left(f^n, \overline{b(f^n)}\right)$. Note that, since we established in Lemma 1 that F is compact, we have that $f^* \in F$, and $b_u(f^*)$ is well-defined for all u.

We first show that, for each $u \in U$, it holds that $b_u(f^*) \le b_u^*$. To that end, note that, since $f^* = \lim_{n\to\infty} f^n$, there must exist an integer $N>0$, such that, if in the profile f^* some user $u \in U$ transfers a positive flow through some link e (i.e., $f_e^{*u} > 0$), it must also does so through each of the profiles $\{f^n\}_{N+1}^{\infty}$ (i.e., $f_e^{n,u} > 0$ for each $n > N$); hence, for each $u \in U$ and $n > N$, the set of links that are employed by u in f^* is *contained* in the set of links that are employed by u in f^n. Therefore, $b_u^* = \lim_{n\to\infty} b_u(f^n) = \lim_{n\to\infty} \max_{e\in E | f_e^{n,u}>0}\{q_e(f_e^n)\} \ge \lim_{n\to\infty} \max_{e\in E | f_e^{*u}>0}\{q_e(f_e^n)\}$ for each $u \in U$. Moreover, since $q_e(\cdot)$ is continuous for each $e \in E$, it holds that $\lim_{n\to\infty} q_e(f_e^n) = q_e(f_e^*)$ for each $e \in E$. Thus, $b_u^* \ge \lim_{n\to\infty} \max_{e\in E | f_e^{*u}>0}\{q_e(f_e^n)\} = \max_{e\in E | f_e^{*u}>0}\{q_e(f_e^*)\} = b_u(f^*)$ for each $u \in U$.

Following the definition of *better-reply security*, assume that f^* is not an equilibrium point. Hence, there exists a user $u \in U$ with a better reply g_u, such that $b_u(f^*) > b_u(g_u, f_{-u}^*)$. Therefore, since we have just shown that $b_u(f^*) \le b_u^*$ for each $u \in U$, it follows that

$$b_u(g_u, f_{-u}^*) < b_u^*. \qquad \text{(A4)}$$

Literally, user u in profile f^* has a reply g_u such that $b_u(f_{-u}^*, g_u) < b_u^*$. We now show that the user u can *secure* a bottleneck below b_u^* at f^*. To that end, it is sufficient to show that, for the reply g_u, there exists some neighborhood of f_{-u}^*, say $N(f_{-u}^*)$, such that $b_u(g_u, \widetilde{f_{-u}}) < b_u^*$ for each $\widetilde{f_{-u}} \in N(f_{-u}^*)$.

Define $\varepsilon \triangleq \dfrac{b_u^* - b_u(g_u, f_{-u}^*)}{2} > 0$. Since $q_e(\cdot)$ is continuous for each $e \in E$, there exists a neighborhood $N(f_{-u}^*)$, such that, for each $\widetilde{f_{-u}} \in N(f_{-u}^*)$ and $e \in E$, it holds that $q_e(g_u, \widetilde{f_{-u}}) - q_e(g_u, f_{-u}^*) < \varepsilon$; hence, $b_u(g_u, \widetilde{f_{-u}}) - b_u(g_u, f_{-u}^*) = \max_{e\in E | g_e^u>0}\{q_e(g_u, \widetilde{f_{-u}})\} - \max_{e\in E | g_e^u>0}\{q_e(g_u, f_{-u}^*)\} < \varepsilon$ for each $\widetilde{f_{-u}} \in N(f_{-u}^*)$. Hence, it holds that $b_u(g_u, \widetilde{f_{-u}}) < b_u(g_u, f_{-u}^*) + \varepsilon = \dfrac{b_u^* + b_u(g_u, f_{-u}^*)}{2}$ for each $\widetilde{f_{-u}} \in N(f_{-u}^*)$. Therefore, from (A4), we obtain that $b_u(g_u, \widetilde{f_{-u}}) < \dfrac{b_u^* + b_u(g_u, f_{-u}^*)}{2} < \dfrac{b_u^* + b_u^*}{2} = b_u^*$ for each $\widetilde{f_{-u}} \in N(f_{-u}^*)$. Thus, by definition, user u can secure a bottleneck strictly below b_u^* at f^*. Hence, the game is better reply secure. ■

Lemma 3: Splittable bottleneck games are quasi-convex.

Proof: Given is a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$. We first show that, for each $u \in U$, the set of feasible flow vectors F_u is convex. Although the proof is quite immediate, we provide the details for completeness. For each value $\lambda \in [0,1]$ and pair of strategies (feasible flow vectors) $h_u, g_u \in F_u$, it holds that $\sum_{p\in P^{(s_u,t_u)}} \lambda\cdot g_p^u + (1-\lambda)\cdot h_p^u = \lambda \cdot \sum_{p\in P^{(s_u,t_u)}} g_p^u + (1-\lambda)\cdot \sum_{p\in P^{(s_u,t_u)}} h_p^u = \lambda\cdot\gamma_u + (1-\lambda)\cdot\gamma_u = \gamma_u$. Moreover, since $h_p^u, g_p^u \ge 0$ for each $p \in P^{(s_u,t_u)}$ it holds that $\lambda\cdot h_p^u + (1-\lambda)\cdot g_p^u \ge 0$ for each $p \in P^{(s_u,t_u)}$. Therefore, $\lambda\cdot h_u + (1-\lambda)\cdot g_u$ is also a feasible flow vector i.e., $\lambda\cdot h_u + (1-\lambda)\cdot g_u \in F_u$; Thus, by definition, F_u is convex [8].

We turn to show that, for each $u \in U$ and for every $f_{-u} \in F_{-u}$, the payoff function $b_u(\ \cdot\ , f_{-u})$ is quasi-convex on F_u. By definition, we have to show that, for each $u \in U$, $\alpha \in \mathbb{R}$, and $f_{-u} \in F_{-u}$, the set $\{f_u \in F_u | b_u(f_u\ , f_{-u}) \le \alpha\}$ is convex. Accordingly, given a pair of feasible flow vectors $h_u, g_u \in F_u$ with $b_u(h_u\ , f_{-u}) \le \alpha$ and $b_u(g_u\ , f_{-u}) \le \alpha$, it is sufficient to show that for each value $\lambda \in [0,1]$, $\lambda\cdot h_u + (1-\lambda)\cdot g_u$ is also a feasible flow vector with $b_u(\lambda\cdot h_u + (1-\lambda)\cdot g_u\ , f_{-u}) \le \alpha$; yet, since we have just proven that the set of all feasible flow vectors F_u is convex, it follows that $\lambda\cdot h_u + (1-\lambda)\cdot g_u \in F_u$ i.e., $\lambda\cdot h_u + (1-\lambda)\cdot g_u$ is a feasible flow vector. Thus, we only left to show that $b_u(\lambda\cdot h_u + (1-\lambda)\cdot g_u\ , f_{-u}) \le \alpha$ for each $\lambda \in [0,1]$. Denote by f_e^{-u} the flow that f_{-u} induces on a link e. Then, given a value $\lambda \in [0,1]$, we have to show by definition that $q_e(\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u\ + f_e^{-u}) \le \alpha$ for each link $e \in E$ with $\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u > 0$.

Given is a link $e \in E$ with $\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u > 0$. We only have to show that $q_e(\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u\ + f_e^{-u}) \le \alpha$. Without loss of generality assume that $h_e^u \le g_e^u$. Therefore, it follows that (i) $\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u \le g_e^u$ and (ii) $g_e^u > 0$. Considering (i), since $q_e(\cdot)$ is monotonically increasing, it holds that (iii) $q_e(\lambda\cdot h_e^u + (1-\lambda)\cdot g_e^u + f_e^{-u}) \le q_e(g_e^u + f_e^{-u})$. Finally, since we establish in (ii) that $g_e^u > 0$ and we assume that $b_u(g_u\ , f_{-u}) \le \alpha$ it follows by definition (of $b_u(\cdot)$) that

$q_e\left(g_e^u + f_e^{-u}\right) \le \alpha$. This together with (iii) establish that $q_e\left(\lambda \cdot h_e^u + (1-\lambda)\cdot g_e^u + f_e^{-u}\right) \le \alpha$. Thus, the lemma is established. ■

Appendix A.2: Convergence to Nash Equilibrium

In this Appendix we provide the proof of Theorem 2, which claims that, for unsplittable bottleneck games, the ESS scheme converges to a Nash equilibrium within a finite number of steps. Then, for splittable bottleneck games, we show, by way of an example, that convergence time can be unbounded.

Theorem 2: In unsplittable bottleneck games, starting with any flow configuration, the ESS scheme converges to a Nash equilibrium within at most $2^{|U|^2} \cdot |E|^{|U|}$ steps.

Proof: Given a flow configuration *f* and a link $e \in E$, denote by $U(e) \subseteq U$ the collection of all users that ship their traffic through *e*. Note that since we consider unsplittable flows it holds that $q_e(f_e) = q_e\left(\sum_{u\in U(e)} \gamma_u\right)$ for each link $e \in E$; hence, the number of values that $q_e(f_e)$ can take is at most the number of different subsets $U(e)$. Thus, since the latter equals $2^{|U|}$, it holds that in any flow configuration the performance on each link can take at most $2^{|U|}$ different values. Therefore, since (by definition) the bottleneck of each user is equal to the performance of some link in *E*, the bottleneck values of each user in *U* in all possible flow configurations are included in a set of at most $2^{|U|} \cdot |E|$ different values.

Given a profile *f* denote by *B*(*f*) the vector of all user bottlenecks i.e., $B(f) \triangleq \left(b_{u_1}(f), b_{u_2}(f), \ldots, b_{u_{|U|}}(f)\right)$. Obviously, since the bottleneck of each user can take at most $2^{|U|} \cdot |E|$ different values, the number of different user bottleneck vectors is at most $\left(2^{|U|} \cdot |E|\right)^{|U|} = 2^{|U|^2} \cdot |E|^{|U|}$.

Let $\langle f_1, f_2, \cdots \rangle$ be the sequence of profiles obtained throughout the execution of the ESS scheme. By definition of the scheme, for each two consecutive profiles f_i, f_{i+1}, it holds that exactly one user in f_{i+1} reroutes its traffic and improves its bottleneck with respect to f_i. Since there are at most $2^{|U|^2} \cdot |E|^{|U|}$ user bottleneck vectors, it follows that, after at most $2^{|U|^2} \cdot |E|^{|U|}$ transitions between successive profiles in $\langle f_1, f_2, \cdots \rangle$, we must encounter a pair of profiles f_i, f_j where $i < j \le 2^{|U|^2} \cdot |E|^{|U|}$ such that $B(f_i) = B(f_j)$. Consider then the sequence of profiles $\langle f_i, f_{i+1}, \cdots f_j \rangle$. Assume, by way of contradiction, that the EES scheme does not converge to a Nash equilibrium in $2^{|U|^2} \cdot |E|^{|U|}$ steps. Hence, for each profile in $\langle f_i, f_{i+1}, \cdots f_j \rangle$, there exists at least one user that can improve its bottleneck.

Let $\widetilde{U} \subseteq U$ be the set of users whose bottleneck is not constant over all profiles of $\langle f_i, f_{i+1}, \cdots f_j \rangle$. In addition, let $u \in \widetilde{U}$ be a user that produces the worst (i.e., largest) bottleneck among all users of $\widetilde{U}$ and over all profiles of $\langle f_i, f_{i+1}, \cdots f_j \rangle$. Finally, let $\Theta = \{f_k\}$ be the set of all profiles in $\langle f_i, f_{i+1}, \cdots f_j \rangle$ for which *u* achieves the worst bottleneck. Since $B(f_i) = B(f_j)$ it holds that *u* has the same bottleneck in profiles f_i and f_j. Therefore, since the bottleneck of *u* is not constant over all profiles in $\langle f_i, f_{i+1}, \cdots f_j \rangle$, it follows that there must exists at least one transition from a profile $f_{k-1} \notin \Theta$ to a profile $f_k \in \Theta$ (which increases the bottleneck of *u*) and at least one transition from a profile $f_{l-1} \in \Theta$ to a profile $f_l \notin \Theta$ (which reduces the bottleneck of *u*). In order to achieve a contradiction it is sufficient to show that the transition from a profile $f_{k-1} \notin \Theta$ to a profile $f_k \in \Theta$ is unattainable.

Assume, by way of contradiction, that there exists a profile $f_k \in \Theta$ such that $f_{k-1} \notin \Theta$. In the transition from the profile f_{k-1} to the profile f_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 $\langle f_i, f_{i+1}, \cdots f_j \rangle$; hence, $u' \in \widetilde{U}$. Next, since the bottleneck of *u* has a smaller value in f_{k-1} than in f_k, it follows that *u*' transfers a positive amount of flow through some bottleneck of *u* in the profile f_k. Therefore, by definition, the bottleneck of *u*' in the profile f_k is equal to that of *u*. However, since the bottleneck of *u* in the profile f_k is maximal with respect to all users in $\widetilde{U}$ and all profiles of $\langle f_i, f_{i+1}, \cdots f_j \rangle$, the bottleneck of the user $u' \in U$ is not improved in the transition from f_{k-1} to f_k. This contradicts the way by which user u' was selected; hence, a transition from a profile $f_{k-1} \notin \Theta$ to a profile $f_k \in \Theta$ is unattainable. Thus, the Theorem is established. ■

On the other hand we now show, by way of an example, that, for splittable bottleneck games, convergence time may be unbounded even for two users. To that end, consider the network depicted in Fig. A.1. Assume that each link $e \in E$ has a performance function $q_e(f_e) = f_e$. Assume that there are two users, A and B, each with a demand of two units. The source-destination pairs of A and B are (s_1,t_1) and (s_2,t_2), correspondingly. Suppose that the initial flow configuration is $\left\langle f_{p_1} = 1, f_{p_2} = 1, f_{p_3} = 0, f_{p_4} = 2 \right\rangle$. It is easy to see that, at the k[th] step of the ESS scheme, we have $f_{p_2}^k + f_{p_3}^k = 2 - \sum_{i=0}^{k} (-1/2)^i$, for $k \in [0, \infty)$. Moreover, note that the (unique) Nash flow is $\left\langle f_{p_1} = 4/3, f_{p_2} = 2/3, f_{p_3} = 2/3, f_{p_4} = 4/3 \right\rangle$; hence, a flow is at Nash equilibrium in this game only if $f_{p_2} + f_{p_3} = 4/3$. However, $f_{p_2}^k + f_{p_3}^k = 2 - \sum_{i=0}^{k} (-1/2)^i \ne 4/3$ for all finite values of *k*.

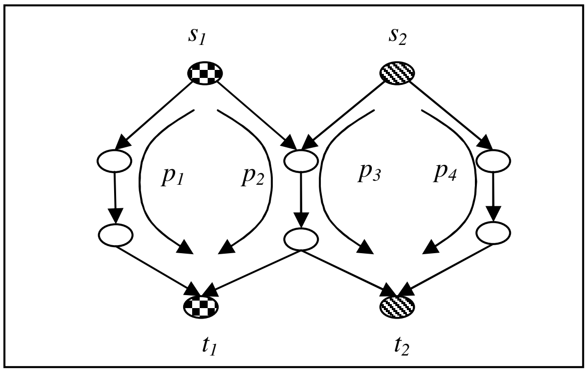

Fig. A.1: Unbounded convergence time for splittable games

Appendix A.3: (In)Efficiency of Unsplittable Nash Equilibria

In this Appendix, we show that the price of anarchy is $O\left(|E|^p\right)$ when the link performance functions are polynomial with a degree p; we also show by way of an example, that the latter result is tight. Then, we present the proof of Theorem 3, which shows that every unsplittable bottleneck game has at least one Nash equilibrium that optimizes the network bottleneck. Finally, we present the proof of Theorem 4, which shows that finding such "good" Nash equilibria is NP-hard. We begin with the following theorem that bounds the price of anarchy for polynomial link performance functions.

Theorem A.1: Given an unsplittable bottleneck game $\langle G(V,E),U,\{q_e\}\rangle$, where $q_e(f_e)=a\cdot(f_e)^p$ for each $e\in E$, the price of anarchy is $O\left(|E|^p\right)$.

Proof: Suppose f and f^* are an unsplittable Nash flow and an unsplittable optimal flow, respectively. It is easy to see that, since the performance functions of all links are identical, the worst possible network bottleneck can be obtained when the traffic of all users traverses some common link i.e., when the total flow through a link is $\sum_{u\in U}\gamma_u$. In that case, the network bottleneck is $a\cdot\left(\sum_{u\in U}\gamma_u\right)^p$; hence,

$$B(f)\le a\cdot\left(\sum_{u\in U}\gamma_u\right)^p. \qquad \text{(A5)}$$

On the other hand, in every (feasible) unsplittable flow vector, there exists at least one link that transfers at least $\frac{1}{|E|}\cdot\sum_{u\in U}\gamma_u$ flow units. Thus, we conclude that, the network bottleneck is at least $a\cdot\left(\frac{1}{|E|}\cdot\sum_{u\in U}\gamma_u\right)^p$; hence,

$$B(f^*)\ge a\cdot\left(\frac{1}{|E|}\cdot\sum_{u\in U}\gamma_u\right)^p. \qquad \text{(A6)}$$

Combing (A5) and (A6) yields the following upper bound on the price of anarchy:

$$\frac{B(f)}{B(f^*)}\le\frac{a\cdot\left(\sum_{u\in U}\gamma_u\right)^p}{a\cdot\left(\frac{1}{|E|}\cdot\sum_{u\in U}\gamma_u\right)^p}=|E|^p. \qquad ■$$

We proceed to show that the result of Theorem A.1 is tight. Specifically, we provide an example where the ratio between the Nash flow and the optimal flow is $\Omega\left(|E|^p\right)$.

Example A.2: Consider the network presented in Fig. A.2, and suppose that there are $|U|$ users, each with a demand of γ flow units from source s to destination t. Assume that each link $e\in E$ has a performance function $q_e(f_e)=(f_e)^p$, $p\ge 0$. It is easy to see that, when all users transfer their traffic through the path (s, v_1, v_2, v_3,..., v_k, t), they are at a (inefficient) Nash equilibrium with a network bottleneck of $(|U|\cdot\gamma)^p$. On the other hand, note that an optimal solution allocates a single path from $\Pi\triangleq\{(s,v_1,t),(s,v_2,v_3,t),(s,v_4,v_5,t),\cdots,(s,v_{k-1},v_k,t)\}$ to a set of at most $\left\lceil\frac{|U|}{|\Pi|}\right\rceil$ users from U. Obviously, the network bottleneck in that case is at most $\left(\left\lceil\frac{|U|}{|\Pi|}\right\rceil\cdot\gamma\right)^p$. Hence, the price of anarchy is at least $(|U|\cdot\gamma)^p\Big/\left(\left\lceil\frac{|U|}{|\Pi|}\right\rceil\cdot\gamma\right)^p=\Omega\left(|\Pi|^p\right)$. Finally, note that $|\Pi|=\frac{k+1}{2}$ and $k=\Omega(|E|)$; hence, the price of anarchy is $\Omega\left(|E|^p\right)$.

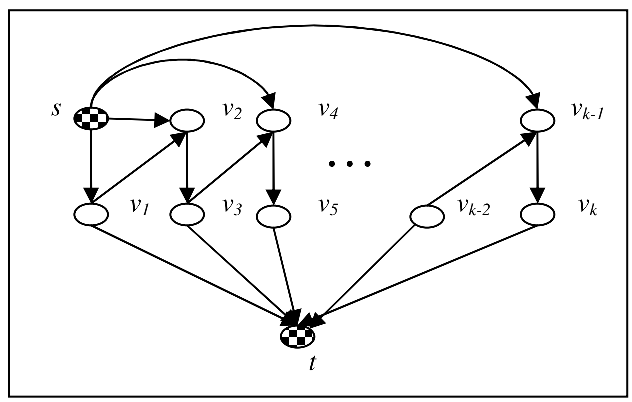

Fig. A.2*:* A lower bound of $\Omega\left(|E|^p\right)$ on the price of anarchy of unsplittable bottleneck games with polynomial performance functions

We turn to present the proof of Theorem 3, which establishes that the price of stability is 1 for unsplittable bottleneck games.

Theorem 3: Given a flow vector f^* for the unsplittable bottleneck game $\langle G,U,\{q_e\}\rangle$, there exists a Nash flow *f* such that $B(f)\le B\left(f^*\right)$; hence, the price of stability is 1.

Proof: Note that, in an unsplittable game, 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 i.e., the game has *finitely* many profiles (feasible unsplittable flow vectors).

Given an unsplittable bottleneck game $\langle G,U,\{q_e\}\rangle$, consider its (finite) set of unsplittable flow vectors Θ. For each $f\in\Theta$, consider the set of all user bottlenecks $\{b_u(f)|u\in U\}$. Let $A(f)$ be a vector that contains all the elements of $\{b_u(f)|u\in U\}$ *sorted* in a non-increasing order and let $A(f)_i$ denote the value of the *i*-th element in $A(f)$. Finally, denote by $F\subseteq\mathbb{R}^{|U|}$ the set of all vectors $A(f)$ that correspond to a feasible flow vector for the game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$ i.e., $F\triangleq\{A(f)|f \text{ is a feasible flow vector for the given game}\}$.

We define the relation < over $\mathbb{R}^{|U|}$ as follows. For each $A(x),A(y)\in\mathbb{R}^{|U|}$ we say that $A(x)<A(y)$ *iff* there exists an $1<i\le|U|$ such that $A(x)_j=A(y)_j$ for each $j\in[1,i-1]$ and $A(x)_i<A(y)_i$. Moreover, for each $A(x)\in F$, we say that $A(x)=\min F$ if there is no $A(y)\in F$ such that $A(y)<A(x)$. Obviously, since we have shown that there are finitely many feasible flow vectors in the game, *F* is finite as well. Therefore, there exists an $A(f)\in F$ for which there is no $A(g)\in F$ such that $A(g)<A(f)$; hence, there exists an $A(f)\in F$ such that $A(f)=\min F$.

Let $A(f)=\min F$. We now show that *f* is at Nash equilibrium. By definition, for each $1<i\le|U|$, the bottleneck $A(f)_i$ of user *i* is minimized with respect to all vectors $A(g)\in F$ that satisfy $A(f)_j=A(g)_j$ for each $j\in[1,i-1]$. Thus, user *i* can improve its bottleneck *only* if the bottleneck $A(f)_j$ of at least one user $j\in[1,i-1]$ is modified. Therefore, for each $1<i\le|U|$, given the (fixed) flow of the users 1,2,…, *i*-1, user *i* cannot improve its bottleneck unless it modifies the bottleneck $A(f)_j$ of at least one user $j\in[1,i-1]$. We distinguish between two cases in which user *i* can modify the bottleneck of users in $[1,i-1]$. In the first case, user *i* increases the bottleneck $A(f)_j$ of at least one user $j\in[1,i-1]$; obviously, this can be done only if *i* ships a positive amount of traffic through one of the bottlenecks of *j*; hence, since $A(f)_j\ge A(f)_i$ for each $j\in[1,i-1]$, such a strategy can only worsen the bottleneck of *i*. In the second case, user *i* only decreases (and never increases) the bottlenecks of users in $[1,i-1]$. However, it is easy to see that such a strategy is unattainable since, by definition, it implies that there exists a feasible flow vector *g* that satisfies $A(g)<A(f)$; obviously, this contradicts the selection of *f* i.e., the fact that $A(f)=\min F$. Thus, we conclude that, for each $1<i\le|U|$, user *i* cannot improve its bottleneck; hence *f* is an unsplittable Nash flow.

Finally, in order to prove the theorem, we have to show that *f* produces the optimal network bottleneck among all unsplittable flows. To that end, let *g* be an unsplittable flow vector for the game $\langle G(V,E),U,\{q_e\}\rangle$. We have to show that $B(f)\le B(g)$.

Since $A(f)=\min F$, it holds by definition that the worst user bottleneck is optimized i.e., $\max_{u\in U}\{b_u(f)\}\le\max_{u\in U}\{b_u(g)\}$.

Hence, $\max_{e\in E|f_e>0}\{q_e(f_e)\}=\max_{u\in U}\left\{\max_{e\in E|f_e^u>0}\{q_e(f_e)\}\right\}=\max_{u\in U}\{b_u(f)\}\le$

$\le\max_{u\in U}\{b_u(g)\}=\max_{u\in U}\left\{\max_{e\in E|g_e^u>0}\{q_e(g_e)\}\right\}=\max_{e\in E|g_e>0}\{q_e(g_e)\}\le\max_{e\in E}\{q_e(g_e)\}$.

Moreover, since the performance functions are increasing, it holds that $q_e(0)\le q_e(g_e)$ for each $e\in E$; hence, $\max_{e\in E|f_e=0}\{q_e(f_e)\}=\max_{e\in E|f_e=0}\{q_e(0)\}\le\max_{e\in E}\{q_e(g_e)\}$. Thus, since both $\max_{e\in E|f_e>0}\{q_e(f_e)\}\le\max_{e\in E}\{q_e(g_e)\}$ and $\max_{e\in E|f_e=0}\{q_e(f_e)\}\le\max_{e\in E}\{q_e(g_e)\}$ hold, we conclude that $\max_{e\in E}\{q_e(f_e)\}\le\max_{e\in E}\{q_e(g_e)\}$ i.e., $B(f)\le B(g)$. ■

Finally, we present the proof of Theorem 4, which establishes that it is intractable to compute a best Nash equilibrium.

Theorem 4: Given an unsplittable bottleneck game $\langle G(V,E),U,\{q_e\}\rangle$ and a value *B*, it is NP-hard to determine if the game has a Nash equilibrium with a bottleneck of at most *B*.

Proof: The reduction is from the *Disjoint Connecting Paths Problem* [15]. Given a network *G*(*V*,*E*) and *k* distinct source-destination nodes $(s_1,t_1),(s_2,t_2),\ldots,(s_k,t_k)$, the goal of the Disjoint Connecting Paths Problem is to find *k* mutually link-disjoint paths connecting the *k* node pairs. Given an instance $\langle G(V,E),\{(s_1,t_1),(s_2,t_2),\ldots,(s_k,t_k)\}\rangle$ of the Disjoint Connecting Paths Problem, associate with each link a performance function $q_e(f_e)=f_e$ and assign to each source-destination pair a user with a demand of *B* flow units. We show that there exists an unsplittable Nash flow with a network bottleneck of at most *B* *iff* there exist *k* mutually disjoint paths connecting the *k* node pairs.

⇒: Assume that there exists an unsplittable Nash flow with a network bottleneck of at most *B*. Then, each of the *k* users transfers its traffic unsplittably without intersecting with the routing paths of the other users. Hence, there exist *k* mutually disjoint paths that connect the node pairs $(s_1,t_1),(s_2,t_2),\ldots,(s_k,t_k)$.

⇐: Assume there exist *k* mutually disjoint paths connecting the *k* node pairs. Assign *B* flow units to each path. It is easy to see that the resulting flow is an unsplittable Nash flow with a network bottleneck of at most B. ■

Appendix A.4: (In)Efficiency of Splittable Nash Equilibria

In this appendix, we present the proofs of Lemmas 4 and 6 that are used in Section V.B to show that every splittable Nash flow that satisfies the efficiency condition is a social optimum. To that end, we first prove the following lemma.

Lemma A.1: Given a flow vector *f* for the splittable bottleneck game $\langle G(V,E),U,\{q_e(\cdot)\}\rangle$, denote by $\tilde{E}$ the set of all links that are the bottlenecks of the network with respect to *f* i.e., $\widetilde{E} \triangleq \left\{ e \in E \middle| q_e(f_e) = B(f) \right\}$; let g be the flow that is induced by f. Then, the following two properties hold:

(i) The network bottleneck of g equals that of f i.e., $B(f) = B(g)$.

(ii) For each $e \in \widetilde{E}$ it holds that $q_e(g) = B(g)$.

Proof: Let f, g and $\widetilde{E}$ satisfy the hypothesis of the Theorem. We first show that $B(f) = B(g)$. To that end, recall that $\Gamma(f) \subseteq U$ is the set of all users that ship traffic through one or more bottlenecks from $\widetilde{E}$ in the flow vector f; hence, by definition, the traffic that is carried over the links of $\widetilde{E}$ belongs only to the users in $\Gamma(f)$. Therefore, since $g_p^u = f_p^u$ for each $u \in \Gamma(f)$ and $p \in P$, it holds that $f_e = g_e$ for each $e \in \widetilde{E}$. Therefore,

$$q_e(g_e) = q_e(f_e) = B(f) \text{ for each } e \in \widetilde{E}; \tag{A7}$$

Hence, $B(g) = \max_{e \in E}\{q_e(g_e)\} \geq \max_{e \in \widetilde{E}}\{q_e(g_e)\} = \max_{e \in \widetilde{E}}\{q_e(f_e)\} = B(f)$. Finally, since g is induced by f it holds that $g_e \leq f_e$ for each $e \in E$; hence, since $q_e(\cdot)$ is increasing, it holds that $q_e(g_e) \leq q_e(f_e)$ for each $e \in E$; therefore, $B(g) \leq B(f)$. Thus, we conclude that

$$B(g) = B(f). \tag{A8}$$

Consequently, item (i) of the Lemma is established by equation (A8) and item (ii) follows from the combination of equations (A7) and (A8). Hence, the lemma is established. ■

We can now prove Lemmas 4 and 6.

*Lemma 4***:** Given a splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$ for which f is a Nash flow that satisfies the efficiency condition, the induced flow vector g satisfies the demands of all users in $\Gamma(f)$, such that the following two properties hold:

(i) g is a Nash flow that satisfies the efficiency conditions;

(ii) the bottleneck of each user $u \in \Gamma(f)$ in flow vector g equals 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 prove that, in g, the bottleneck of each user of $\Gamma(f)$ is equal to that of the network i.e., $b_u(g) = B(g)$ for each $u \in \Gamma(f)$. To that end, denote by $\widetilde{E}$ the set of all links that are the bottlenecks of the network with respect to f i.e., $\widetilde{E} \triangleq \left\{ e \in E \middle| q_e(f_e) = B(f) \right\}$. 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 network bottleneck $e' \in \widetilde{E}$ 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 also ship (positive) traffic through the link e' in the flow vector g. Finally, since g is induced by f, it holds according to Lemma A.1 that $q_e(g) = B(g)$ for each $e \in \widetilde{E}$; in particular it holds that $q_{e'}(g) = B(g)$; hence, by definition $b_u(g) = B(g)$.

We turn to prove that the efficiency condition is satisfied for g. To that end, first note that $\widetilde{E}$ is also the set of network bottlenecks with respect to g; indeed, g is established by zeroing the flows of all users that are not shipping traffic through any network bottleneck in f; thus, since it follows that the performance of the links in $\widetilde{E}$ remain unchanged in g and the performance of all other links is not worsen, $\widetilde{E}$ is the set of network bottlenecks also with respect to g. Next, since we assume that the efficiency condition is satisfied for f, it holds by definition that $N_f(p_1) \leq N_f(p_2)$ for each $u \in U$ and $p_1, p_2 \in P^{(s_u,t_u)}$ with $f_{p_1}^u > 0$. Therefore, since $g_p^u = f_p^u$ for each $u \in \Gamma(f)$ and $p \in P$, it holds that $N_g(p_1) \leq N_g(p_2)$ for each $u \in \Gamma(f)$ and $p_1, p_2 \in P^{(s_u,t_u)}$ with $g_{p_1}^u > 0$; hence, by definition, the efficiency condition is satisfied for g.

We turn to show that g is at Nash equilibrium. To that end, recall that the flow vector g is constructed out of the given Nash flow f by removing the flows of all users that are not in $\Gamma(f)$. By definition, the users that are not in $\Gamma(f)$ do not transfer (positive) flow through a network bottleneck (i.e., a link in $\widetilde{E}$). Hence, their removal can improve *only* the performance of the non-bottleneck links (i.e., links that before the removal had already better performance then the links in $\widetilde{E}$). Thus, since the performance of the bottleneck links $\widetilde{E}$ remain unchanged after the removal of $U \setminus \Gamma(f)$, each user in $\Gamma(f)$ can improve its bottleneck in the flow vector g only if it can reduce its flow from each link $e \in \widetilde{E}$ that it employs, without increasing its flow over any other (unemployed) link $e' \in \widetilde{E}$. However, since before the removal (i.e., in flow vector f) all the users are at Nash equilibrium, it follows that there is no user $u \in \Gamma(f)$ in flow vector f that can reduce its flow from every network bottleneck $e \in \widetilde{E}$ that it employs without increasing its flow over other network bottleneck $e' \in \widetilde{E}$; indeed, otherwise a user $u \in \Gamma(f)$ could have decreased its bottleneck value. Therefore, since all the strategies of the users in $\Gamma(f)$ remain unchanged in g (i.e., $g_p^u = f_p^u$ for each $u \in \Gamma(f)$ and $p \in P$), the latter holds also for flow vector g (i.e., there is no user that can reduce its flow from each employed link $e \in \widetilde{E}$, without employing other link $e' \in \widetilde{E}$). Therefore, there is no user in g that can improve its bottleneck. Thus, g is a Nash flow and the Lemma is established. ■

Lemma 6: Given a flow vector f for the splittable bottleneck game $\langle G(V,E), U, \{q_e(\cdot)\}\rangle$, denote by g the flow that is induced by f. If g has the smallest network bottleneck among the flow vectors that satisfy the demands of the users in $\Gamma(f)$, then f has the smallest network bottleneck among the flow vectors that satisfy the demands of the users in U.

Proof: Let f^* be an optimal flow i.e., a flow that has the smallest (network) bottleneck with respect to all flow vectors that satisfy the demands of the users in U. Denote by H 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,

$$B\left(f^*\right)\geq B\left(\widetilde{f^*}\right). \tag{A9}$$

Yet, 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 (A9), it follows that

$$B\left(f^*\right)\geq\min\left\{B(h)\,\middle|\,h\in\mathrm{H}\right\}. \tag{A10}$$

Let f and g satisfy the hypothesis of the Lemma and assume that g has the smallest network bottleneck among the flow vectors that satisfy the demands of the users in $\Gamma(f)$ i.e., $B(g)=\min\left\{B(h)\,\middle|\,h\in\mathrm{H}\right\}$. From (A10) it holds that $B(f^*)\geq B(g)$. Therefore, since it follows from Lemma A.1[1] that $B(g)=B(f)$, it holds that $B(f^*)\geq B(f)$; hence, since f^* is optimal, $B(f^*)=B(f)$. Thus, the lemma is established. ■

---

[1] Obviously, the conditions of Lemma A.1 are satisfied since g is induced by f.