

Evolutionary Equilibrium in Bayesian Routing Games: Specialization and Niche Formation

Petra Berenbrink* and Oliver Schulte**

School of Computing Science
Simon Fraser University

Abstract. In this paper we consider Nash Equilibria for the selfish routing model proposed in [12], where a set of n users with tasks of different size try to access m parallel links with different speeds. In this model, a player can use a mixed strategy (where he uses different links with a positive probability); then he is indifferent between the different link choices. This means that the player may well deviate to a different strategy over time. We propose the concept of evolutionary stable strategies (ESS) as a criterion for stable Nash Equilibria, i.e. Equilibria where no player is likely to deviate from his strategy. An ESS is a steady state that can be reached by a user community via evolutionary processes in which more successful strategies spread over time. The concept has been used widely in biology and economics to analyze the dynamics of strategic interactions. We establish that the ESS is *uniquely determined* for a symmetric Bayesian parallel links game (when it exists). Thus evolutionary stability places strong constraints on the assignment of tasks to links.

Keywords: Selfish Routing, Nash Equilibrium, Bayesian Games, Evolutionarily Stable Strategy

1 Introduction

We consider the *selfish routing model* proposed in [12], where users try to access a set of parallel links with different speeds. This scenario gives rise to a strategic interaction between users that combines aspects of both competition, in that users compete for the fastest links, and coordination, in that users want to avoid overloaded links. Koutsoupias and Papadimitriou suggest to study the model in a game-theoretic frame work [12]. They compare the cost of the worst case Nash equilibrium with the cost of an optimal solution; this ratio was called *price of anarchy*. Depending on the cost function that is used to assess the optimal solution, the ratio between Nash equilibria and optimal solutions can vary greatly. For example, the cost of the worst case Nash equilibrium can be similar to the cost of the optimal solution (min-max function considered in [2]), or the cost for every Nash Equilibrium can be far away from that of the optimal solution [1].

* E-mail: petra@cs.sfu.ca

** E-mail: oschulte@cs.sfu.ca

It is an elementary fact that if a player plays a *mixed Nash strategy*, then he is indifferent between the choices that carry positive probability. So, it is not easy to see what keeps the players from deviating to a different strategy with different probabilities. A Nash Equilibrium having a sequence of single-player strategy changes that do not alter their own payoffs but finally lead to a non-equilibrium position, is called **transient** [6]. Games can have several non-stable and transient Nash Equilibria, and it is unlikely that a system will end up in one of these. Hence it might be interesting to answer first the question which Nash Equilibria are stable, and then to compare the cost of stable equilibria to the cost of the optimal solution (see [6]). Several stability models were suggested in the literature [18]. One of the most important models is Maynard Smith’s concept of an *evolutionarily stable strategy*, abbreviated ESS [14]. The criterion proposed by Maynard Smith is as follows: An equilibrium E is evolutionarily stable if there is a threshold $\bar{\epsilon}$ for the size of deviations such that if the fraction of deviating players falls below $\bar{\epsilon}$, then the players following the equilibrium E do better than the deviants.

1.1 New Results

In this paper we study evolutionarily stable equilibria for selfish routing in Koutsoupias and Papadimitriou’s *parallel links model* [12]. We first define a symmetric version of a Bayesian parallel links game where every player is not assigned a task of a fixed size but, instead, is randomly assigned a task drawn from a distribution (Section 2.1). Then we argue that every ESS in this game is a *symmetric* Nash Equilibrium, where every player uses the same strategy.

In Section 3 we show that the symmetric Nash Equilibrium is unique for *link groups*. By *link group uniqueness* we mean the following. Assume that all links with the same speed are grouped together into so-called *link groups*. Then, in every symmetric Nash Equilibrium, the total probability that tasks of a certain size are sent to a link group is unique. This implies that the only flexibility in a symmetric Nash Equilibrium is the probability distribution over links from the same link group, not over different link groups. Then we show that in a symmetric equilibrium two links with different speeds cannot both be used by two or more tasks with different weights. In fact, we show an even stronger result: If link ℓ is used for task w and ℓ' for $w' \neq w$, then at least one of the links will not be optimal for the other link’s task. We also show that tasks with larger weight must be assigned to links with bigger speed.

In Section 4 we characterize ESS for the symmetric Bayesian parallel links game. We show that every ESS is a Nash Equilibrium, and we show that, to evaluate evolutionary stability, we have to consider only best replies to the current strategy. Then we establish that in an ESS, we not only have link group uniqueness, but also the probability distribution with which links of the same group are chosen by tasks has to be unique. In fact, an ESS requires treating two links with equal speed exactly the same. This result establishes the *uniqueness of ESS*. We show that in an ESS even two links with the same speed cannot both be used by two or more tasks with different weights. This implies that in

an ESS links acquire niches, meaning that there is minimal overlap in the tasks served by different links. We call this *specialization* in the following.

In general, the problem of calculating an ESS is very hard; it is contained in Σ_2^P and is both NP-hard and coNP-hard [5]. We expect that our uniqueness results and the structural properties of ESS for our game will help to develop algorithms that compute an ESS. We also show that, unfortunately, the specialisation condition is necessary for an ESS, but not sufficient. We introduce a sufficient condition called *clustering*—roughly, links must form disjoint niches—and show that every clustered Nash equilibrium is an ESS. Unfortunately, we also show that there exists a game that does not have a clustered ESS, but it has an unclustered ESS, so clustering is not a necessary condition.

1.2 Known results

The Parallel Links Game was introduced by Koutsoupias and Papadimitriou [12], who initiated the study of coordination ratios. In the model of [12], the cost of a collection of strategies is the (expected) maximum load of a link (maximized over all links). The coordination ratio is defined as the ratio between the maximum cost (maximized over all Nash equilibria) divided by the cost of the optimal solution. Koutsoupias and Papadimitriou give bounds on the coordination ratio. These bounds are improved by Mavronicolas and Spirakis [13], and by Czumaj and Vöcking [3] who gave an asymptotically tight bound.

In [8] Gairing et al. introduce a Bayesian version of the selfish routing game. Following Harsanyi’s approach [9], each user can have a set of possible types. Their paper presents a comprehensive collection of results for the Bayesian routing game. Note that their model is more general than ours since they allow different types for different users, whereas our users all have the same type space (possible task assignments). In [7] Fischer and Vöcking adopt an evolutionary approach to a related routing problem (see [16] for a definition).

The concept of evolutionary stability is fundamental in evolutionary game theory, which has many applications in theoretical biology and economics. The seminal presentation of the concept of an ESS is due to Maynard Smith [14]. Since then, the concept has played a central role in evolutionary biology and has been used in thousands of studies. Economists have also applied the concept to analyze many economic and social interactions. Kontogiannis and Spirakis provide an introduction to and motivation for evolutionary analysis from a computer science perspective [11, Sec.3]. Kearns and Suri examine evolutionary stability in graphical games [10].

Sandholm proposes a pricing scheme based on evolutionary stability for minimizing traffic congestion; he notes the potential applicability of his models to computer networks [17, Sec.8]. His approach does not apply the concept of evolutionarily stable strategy. To our knowledge, our combination of congestion game + Bayesian incomplete information + ESS is new in the literature. (Ely and Sandholm remark that “nearly all work in evolutionary game theory has considered games of complete information” [4, p.84].)

2 Basic Models and Concepts

We first introduce Bayesian Parallel Links Games and show some simple observations concerning link load and utilities. We then introduce population games and define evolutionary stable strategies (ESS).

2.1 Bayesian Parallel Links Games

We examine an extension of the original routing game called *Bayesian parallel links game*. Our definition below is a special symmetric case of the definition in [8]. Harsanyi [9] introduced the notion of a Bayesian game to analyze games with *incomplete information* where players are uncertain about some aspect of the game such as what preferences or options the other players have. Bayesian games have found many applications in economics; eventually Harsanyi's work earned him the Nobel Prize. In a Bayesian parallel links game, the uncertainty among the players concerns the task size of the opponents. An agent knows the size of her own message, but not the size of the messages being sent by other users. The Bayesian game of [8] models this uncertainty by a distribution that specifies the probability that w is the task of user i . In our symmetric Bayesian routing model, this distribution is the same for all agents. A natural interpretation of this assumption is that agents are assigned tasks drawn from a common distribution.

A game is *symmetric* if (1) all players have the same set of strategy options, and (2) all that matters is what strategies are chosen, and how often they are chosen, but not by which player. Our Bayesian version of the game is symmetric, whereas the parallel links game is symmetric for uniform users only. The formal definition of a symmetric Bayesian routing model is as follows.

Definition 1. A *symmetric Bayesian routing model* is a tuple $\langle N, W, \mu, L, P \rangle$ where

1. $N = [n]$ is the set of users
2. $W = \{w_1, w_2, \dots, w_k\}$ is a finite set of task weights, and $\mu : W \rightarrow (0, 1]$ is a probability distribution over the weights W . The distribution μ is used to assign weights i.u.r. to players $1, \dots, n$.
3. $L = [m]$ is the set of links. For $i \in [m]$, link i has speed c_i .
4. An allocation s is a vector in $[0, 1]^m$ such that $\sum_{j=1}^m s(j) = 1$.
5. A mixed strategy p is a total function $W \rightarrow [0, 1]^m$, i.e., a strategy assigns an allocation to every task weight in W . Then P is the set of all strategies. A strategy profile p_1, \dots, p_n assigns a strategy $p_i \in P$ to each player i .

Now fix a routing model with strategy set P . In the following we use $p_i(j, w)$ for the probability that, in strategy $p_i \in P$, user i assigns a task with weight w to link j . As usual, (p_i, p_{-i}) denotes a strategy profile where user i follows strategy p_i and the other players' strategies are given by the vector p_{-i} of length $n - 1$. Similarly, (w_i, w_{-i}) denotes a weight vector where user i is assigned task size w_i and the other players' weights are given by the vector w_{-i} of length $n - 1$. The concept of a mixed strategy in a symmetric Bayesian routing model may be

interpreted as follows. Each player chooses a strategy before the game is played. Then tasks w_1, w_2, \dots, w_n are assigned i.u.r. to players 1 through n according to the distribution μ . Each player learns their own task but not that of the others. Next for each player i we "execute" the strategy p_i given task w_i , such that task w_i is sent to link j with probability $p_i(j, w)$. Thus, strategies have a natural interpretation as programs that take as input a task w and output a link for w or a probability distribution over links for w .

Like Koutsoupias and Papadimitriou [12], we assume that the latency of a link depends linearly on the load of a link. Thus we have the following definition of the load on a link.

Definition 2. Let $B = \langle N, W, \mu, L, P \rangle$ be a symmetric Bayesian routing model. Let p_1, \dots, p_n be a vector of mixed strategies.

1. We define

$$\text{load}_i(p_1, \dots, p_n, w_1, \dots, w_n) \equiv \frac{1}{c_i} \sum_{j=1}^n w_j \cdot p_j(i, w_j)$$

to be the **conditional expected load** on link i for fixed w_1, w_2, \dots, w_n .

2. We define

$$\text{load}_i(p_1, \dots, p_n) \equiv \sum_{w_1, \dots, w_n \in W^n} \text{load}_i(p_1, \dots, p_n, w_1, \dots, w_n) \cdot \prod_{j=1}^n \mu(w_j).$$

to be the **expected load** on link i .

The next observation shows that the load function is additive in the sense that the total load on link i due to n users is just the sum of the loads due to the individual users. The proof can be found in the full version of this paper.

Observation 1 Let $B = \langle N, W, \mu, L, P \rangle$ be a symmetric Bayesian routing model. Let p_1, \dots, p_n be a vector of mixed strategies; then for any user j we have $\text{load}_i(p_1, \dots, p_n) = \text{load}_i(p_{-j}) + \text{load}_i(p_j)$. Therefore $\text{load}_i(p_1, \dots, p_n) = \sum_{j=1}^n \text{load}_i(p_j)$.

A symmetric Bayesian routing game is a symmetric Bayesian routing model with utility functions for the players. In a symmetric game, the payoff of each player depends only on what strategies are chosen, and not on which players choose particular strategies. Hence for a fixed profile of opponents strategy p yields the same payoff no matter which player i follows strategy p . This allows us to simplify our notation and write $p_j(w)$ or $p_\ell(w)$ for the probability that strategy p uses a link j or ℓ for a task with size w .

Definition 3. A **symmetric Bayesian routing game** $B = \langle N, W, \mu, L, u, P \rangle$ is a routing model $\langle N, W, \mu, L, P \rangle$ with a utility function u . We write $u(p; p_1, \dots, p_{n-1})$ to denote the payoff of following strategy p when the other players' strategies are given by p_1, \dots, p_{n-1} . Then the payoff is defined as

$$u(p; p_1, \dots, p_{n-1}) = - \sum_{w \in W} \sum_{\ell \in L} (w/c_\ell + \text{load}_\ell(p_1, \dots, p_{n-1})) \cdot p_\ell(w) \cdot \mu(w).$$

Fix a symmetric Bayesian routing game B with n players. A strategy profile (p_1, \dots, p_n) is a Nash equilibrium if no player can improve their payoff unilaterally by changing their strategy p_i . To simplify notation for games in which several players follow the same strategy, we write p^k for a vector (p, p, \dots, p) with p repeated k times. Then the mixed strategy p is a **best reply** to p^{n-1} if for all mixed strategies p' we have $u(p; p^{n-1}) \geq u(p'; p^{n-1})$. If all players in a symmetric game follow the same strategy, then p^n is the resulting mixed strategy profile. The strategy profile p^n is a (symmetric) **Nash equilibrium** if p is a best reply to p^{n-1} . Hence, a symmetric Nash equilibrium for a symmetric Bayesian routing game with n players is a Nash equilibrium (p, p, \dots, p) in which each player follows the same strategy. It follows from Nash's existence proof [15] that a symmetric game, such as a symmetric Bayesian Routing Game, has a symmetric Nash equilibrium.

2.2 Population Games and Evolutionary Stability for the Parallel Links Game

We give a brief introduction to population games and evolutionary stability. A more extended introduction from a computer science point of view is provided in [11, Sec.3]. The standard population game model considers a very large **population** A of agents [19, 14]. The agents play a symmetric game like our symmetric Bayesian routing game. Every agent in the population follows a strategy p *fixed before* the game is played. A *match* is a particular instance of the base game that results when we match n i.u.r. (independent and uniformly at random) chosen agents together to play the base game. Since strategies occur with a certain frequency in the population, there is a fixed probability with which a task with a given size is assigned to a link. *Hence, with a population A we can associate a mixed strategy that we denote by p_A .*

Consider now the expected payoff that an agent using strategy p receives in a match given a fixed population A . This is equivalent to playing strategy p against $n - 1$ opponents whose choices are determined by *the same distribution*, the population distribution p_A . In other words, the expected payoff is given by $u(p; (p_A)^{n-1})$, namely the payoff of using strategy p when the other $n - 1$ players follow mixed strategy p_A . A population is *in equilibrium* if no agent benefits from changing her strategy unilaterally given the state of the population. Formally, a population A with associated mixed strategy p_A is in equilibrium if every mixed strategy p that occurs with frequency greater than zero in the population is a best reply to $(p_A)^{n-1}$. It is easy to see that this is the case if and only if the symmetric strategy profile (p, p, \dots, p) is a Nash equilibrium. So *population equilibria correspond exactly to symmetric Nash equilibria*. While restricting attention to symmetric Nash equilibria may seem like an artificial restriction for non-population models, in large population models symmetric Nash equilibria characterize the natural equilibrium concept for a population.

The main idea in evolutionary game theory is Maynard Smith's concept of *stability against mutations*. Intuitively, a population is evolutionarily stable if a small group of mutants cannot invade the population. Consider a population A

that encounters a group A^M of mutants. Then the mixed population is $A \cup A^M$. Suppose that in this mixed population the proportion of mutants is ε . The distribution for the mixed population is the probabilistic mixture $(1 - \varepsilon)p_A + \varepsilon p_{A^M}$. We may view a mutation A^M as successful if the average payoff for invaders *in the mixed population* is at least as great as the average payoff for nonmutants in the mixed population. The expected payoff for a strategy p in the mixed population $A \cup A^M$ is given by $u(p; [(1 - \varepsilon)p_A + \varepsilon p_{A^M}]^{n-1})$. So the average payoff for the non-mutants is $u(p_A; [(1 - \varepsilon)p_A + \varepsilon p_{A^M}]^{n-1})$ and for the mutants it is $u(p_{A^M}; [(1 - \varepsilon)p_A + \varepsilon p_{A^M}]^{n-1})$.

Definition 4 (ESS). *Let B be a symmetric Bayesian routing game with n players. A mixed strategy p^* is an **evolutionarily stable strategy (ESS)** \iff there is an $\bar{\varepsilon} > 0$ such that for all $0 < \varepsilon < \bar{\varepsilon}$ and mixed strategies $p \neq p^*$ we have $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) > u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*

3 Link Group Uniqueness of Symmetric Nash Equilibria

This section investigates the structure of symmetric Nash equilibria and establishes that symmetric equilibria are uniquely determined up to the distribution of tasks within link groups. A *link group* \mathcal{L} in a symmetric Bayesian routing game B is a maximal set of links with the same speed, that is, $c_\ell = c_{\ell'}$ for all $\ell, \ell' \in \mathcal{L}$. Then, for any mixed strategy p , the probability that p sends task w to a link in link group \mathcal{L} is given by $p_{\mathcal{L}}(w) \equiv \sum_{\ell \in \mathcal{L}} p_\ell(w)$. The main result of this section is that in any symmetric Bayesian routing game the aggregate distribution over groups of links with the same speed is *uniquely determined* for symmetric Nash equilibria. In other words, the probabilities $p_{\mathcal{L}}$ are uniquely determined in a symmetric Nash equilibrium; if p^n and $(p')^n$ are Nash equilibria in a routing game B , then for every link group \mathcal{L} and every task weight w we have $p_{\mathcal{L}}(w) = p'_{\mathcal{L}}(w)$.

In the following we say that link ℓ is **optimal** for task w given p^{n-1} iff ℓ minimizes the function $w/c_\ell + (n-1) \cdot \text{load}_\ell(p)$. In this case we write $w \in \text{opt}_\ell(p)$. A mixed strategy p **uses** link ℓ for task w if $p_\ell(w) > 0$; we write $w \in \text{support}_\ell(p)$. The next proposition asserts that a best reply p' to a strategy profile p^{n-1} uses a link for a task only if the link is optimal for the task given p^{n-1} . This is a variant of the standard characterization of Nash equilibrium according to which all pure strategies in the support of an equilibrium strategy are best replies. The proof can be done similar to the proof of the standard Nash characterization and is omitted.

Proposition 1. *Let B be a symmetric Bayesian routing game with n players, and let p be a mixed strategy. A strategy p' is a best reply to $p^{n-1} \iff$ for all tasks w , links ℓ , if $p'_\ell(w) > 0$, then ℓ is an optimal link for w given p^{n-1} . That is, $\text{support}_\ell(p') \subseteq \text{opt}_\ell(p)$.*

The next Lemma gives a clear picture of what a symmetric Nash equilibrium looks like. Intuitively, this picture is the following. (1) Tasks with bigger weights are placed on faster links. (2) Faster links have a bigger load. (3-5) For every link

ℓ there is an “interval” of task weights $\{w_i, \dots, w_l\}$ such that ℓ is optimal for all and only these weights. (6) Any pair of links with different speeds are optimal for at most one common task weight.

Lemma 1. *Let B be a symmetric Bayesian routing game with n players and a symmetric Nash equilibrium p^n . Fix any two links ℓ and ℓ' .*

1. *If $c_\ell > c_{\ell'}$, strategy p uses ℓ for task w , and p uses ℓ' for w' , then $w \geq w'$.*
2. *If $c_\ell > c_{\ell'}$, then $\text{load}_\ell(p^n) > \text{load}_{\ell'}(p^n)$, or $\text{load}_\ell(p^n) = \text{load}_{\ell'}(p^n) = 0$. And if $c_\ell = c_{\ell'}$, then $\text{load}_\ell(p^n) = \text{load}_{\ell'}(p^n)$.*
3. *If $c_\ell > c_{\ell'}$, then there cannot exist tasks $w > w'$ with $w, w' \in \text{support}_\ell(p)$ and $w, w' \in \text{support}_{\ell'}(p)$.*
4. *If $w_1 \geq w_2 \geq w_3$ and $w_1 \in \text{opt}_\ell(p)$ and $w_3 \in \text{opt}_{\ell'}(p)$, then $w_2 \in \text{opt}_\ell(p)$.*
5. *If $c_{\ell_1} \geq c_{\ell_2} \geq c_{\ell_3}$ and $w \in \text{opt}_{\ell_1}(p)$ and $w \in \text{opt}_{\ell_3}(p)$, then $\text{opt}_{\ell_2}(p) = \{w\}$.*
6. *If $c_\ell > c_{\ell'}$, then $|\text{opt}_\ell(p) \cap \text{opt}_{\ell'}(p)| \leq 1$.*

Proof. The proof can be found in the full version.

We note that Lemma 1 holds for Nash equilibria in general, not just symmetric ones. Specifically, let p' be a Nash equilibrium for a symmetric Bayesian routing game, and fix any player i such that $p' = (p'_i, p'_{-i})$. Then Lemma 1 holds if we replace a mixed strategy p with p'_i , and p^{n-1} with p'_{-i} , and p^n with (p'_i, p'_{-i}) .

We extend our notation for links to link groups \mathcal{L} such that $c_{\mathcal{L}}$ denotes the speed of all links in group \mathcal{L} . We also define

$$\text{load}_{\mathcal{L}}(p^n) \equiv \sum_{\ell \in \mathcal{L}} \text{load}_\ell(p^n).$$

The next theorem is the main result of this section. It states that for a user population in equilibrium (corresponding to a symmetric Nash equilibrium), the distribution of tasks to *link groups* is uniquely determined. Thus the only way in which population equilibria can differ is by how tasks are allocated within a link group. This result is the first key step for establishing the uniqueness of an ESS for a symmetric Bayesian routing game.

Theorem 1 (Link Group Uniqueness). *Let B be a symmetric Bayesian routing game with n players and two symmetric Nash equilibria p^n and (p') . Then we have $p_{\mathcal{L}}(w) = p'_{\mathcal{L}}(w)$ and $\text{load}_{\mathcal{L}}(p^n) = \text{load}_{\mathcal{L}}((p')^n)$ for all task sizes w and link groups \mathcal{L} of B .*

The proof is in the full version.

The next section investigates the structure of evolutionarily stable equilibria and proves that an ESS is uniquely determined when it exists.

4 Characterization of Evolutionary Stability

In this section we prove a necessary and sufficient condition for a mixed strategy p^* to be an ESS. In the following let $p_\ell(W) = \sum_{w \in W} p_\ell(w) \cdot \mu(w)$. The next proposition shows that for sufficiently small sizes of mutations, only best replies

to the incumbent distribution p^* have the potential to do better than the incumbent. The proposition also implies that an ESS corresponds to a symmetric Nash equilibrium (Corollary 1).

Proposition 2. *Let B be a symmetric Bayesian routing game with n players, and let p^* be a mixed strategy. Then there is a threshold $\bar{\varepsilon}$ such that for all ε with $0 < \varepsilon < \bar{\varepsilon}$, for all mixed strategies p :*

1. *If $u(p^*; (p^*)^{n-1}) > u(p; (p^*)^{n-1})$, then $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) > u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*
2. *If $u(p^*; (p^*)^{n-1}) < u(p; (p^*)^{n-1})$, then $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) < u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*

Proof. The proof requires only standard techniques from evolutionary game theory [19] and is omitted for reasons of space. Intuitively, the result holds because we can choose our threshold $\bar{\varepsilon}$ small enough (as a function of B and p^*) so that any difference in the case in which the mutant and incumbent face 0 mutants outweighs the differences in their payoffs when they face one or more mutants.

Corollary 1. *Let B be a symmetric Bayesian routing game with n players, and let p^* be an ESS. Then $(p^*)^n$ is also a Nash Equilibrium.*

Proof. (sketch) If p^* is not a best reply to $(p^*)^{n-1}$, then there is a mutant p such that $u(p; (p^*)^{n-1}) > u(p^*; (p^*)^{n-1})$. Proposition 2(2) then implies that p is a successful mutation no matter how low we choose the positive threshold $\bar{\varepsilon}$.

The next lemma 2 provides a necessary and sufficient condition for when a best reply is a successful mutation, which is key for our analysis of evolutionarily stable strategies in a given network game.

Lemma 2. *Let B be a symmetric Bayesian routing game with n players. Let $(p^*)^n$ be a Nash equilibrium, and consider any best reply p to $(p^*)^{n-1}$. Define $\Delta_\ell(p, p^*) = \sum_{\ell \in L} [\text{load}_\ell(p^*) - \text{load}_\ell(p)] \cdot [p_\ell^*(W) - p_\ell(W)]$. Then for all ε with $0 < \varepsilon < 1$, if*

1. *$\Delta_\ell(p, p^*) < 0$, then $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) < u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*
2. *$\Delta_\ell(p, p^*) = 0$, then $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) = u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*
3. *$\Delta_\ell(p, p^*) > 0$, then $u(p^*; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1}) > u(p; [\varepsilon p + (1 - \varepsilon)p^*]^{n-1})$.*

The proof of Lemma 2 can be found in the full version. It shows that a best reply p to $(p^*)^{n-1}$ that has a negative quantity $\sum_{\ell \in L} [\text{load}_\ell(p^*) - \text{load}_\ell(p)] \cdot [p_\ell^*(W) - p_\ell(W)]$ is successful in the strong sense that it yields a better payoff than the incumbent strategy p^* no matter what its size. It will be convenient to say that a mutation p **defeats** the incumbent strategy p^* if $\sum_{\ell \in L} [\text{load}_\ell(p^*) - \text{load}_\ell(p)] \cdot [p_\ell^*(W) - p_\ell(W)] < 0$. Similarly, we say that a mutation p **equals** an incumbent p^* if $\sum_{\ell \in L} [\text{load}_\ell(p^*) - \text{load}_\ell(p)] \cdot [p_\ell^*(W) - p_\ell(W)] = 0$. In this terminology our results so far yield the following characterization of evolutionary stability. The proof directly follows from Proposition 2 and Lemma 2.

Corollary 2. *Let B be a symmetric Bayesian routing game with n players. A mixed strategy p^* is an ESS for $B \iff$ the strategy profile $(p^*)^n$ is a Nash equilibrium, and no best reply $p \neq p^*$ to $(p^*)^{n-1}$ defeats or equals p^* .*

Lemma 1 clarified the structure of user populations in equilibrium. The next section applies the criterion from Corollary 2 to establish additional properties of populations in an evolutionarily stable equilibrium. In fact, these properties imply that an evolutionarily stable equilibrium is unique when it exists.

5 Uniqueness and Structure of Evolutionary Stable Strategies

We analyze the structure of evolutionary equilibria and show the uniqueness of ESS. For the first point, our focus is on the allocation of tasks to links that are consistent with evolutionary stability. Such results tell us how the structure of the network shapes evolutionary dynamics. They can be helpful for the development of algorithms calculating an ESS for a given system. The next theorem shows that in an evolutionary equilibrium there is minimal overlap in the tasks served by different links, in that two distinct links (even with the same speed) may not be used by tasks with different weights. In fact the result is stronger in that if link ℓ is used for task w and ℓ' for $w' \neq w$, then at least one of the links must not be optimal for the other link's task. This specialization result can be regarded as a stronger version of Lemma 1(6), where ℓ and ℓ' can have the same speed. Unfortunately, the specialization condition of the Theorem is necessary but not sufficient, as Observation 2 will show.

The idea of the proof of the next theorem is that if two distinct links ℓ and ℓ' are used with a probability > 0 by users with different tasks, it is possible to create a “better” mutant distribution. The mutant distribution increases the load on one of the two links, say ℓ (by putting the task with the bigger weight with a larger probability onto ℓ , and, in turn, by putting the smaller task with smaller probability onto ℓ), but uses the link overall with a smaller probability. Note that this strategy is only possible if we have different task weights.

Theorem 2 (Specialization). *Let B be a symmetric Bayesian routing game with mixed strategy p^* . Assume $w \neq w'$, $\ell \neq \ell'$ with $c_\ell \geq c_{\ell'}$, and suppose the following conditions are fulfilled.*

1. $w \in \text{support}_\ell(p^*)$ and $w' \in \text{support}_{\ell'}(p^*)$,
2. $w, w' \in \text{opt}_\ell(p^*)$, and $w, w' \in \text{opt}_{\ell'}(p^*)$.

Then there is a mutation p that defeats p^ , and hence p^* is not evolutionarily stable.*

The next observation gives a counterexample showing that the specialization condition from Theorem 2 is necessary but unfortunately not sufficient for an ESS.

Observation 2 *There exists a symmetric Bayesian routing game B with a strategy p such that p meets the specialization condition of Theorem 2 for any $w \neq w'$ and $\ell \neq \ell'$, but p is not an ESS.*

Proof. Assume three resources ℓ_1, ℓ_2, ℓ_3 with speeds $c_{\ell_1} = 6, c_{\ell_2} = 4,$ and $c_{\ell_3} = 2$. We have two users and two task sizes $w = 21$ and $w' = 1$. We define $\mu(21) = 2/3$ and $\mu(1) = 1/3$. The strategy p with $p_{\ell_1}(21) = 19/21, p_{\ell_2}(21) = 2/21, p_{\ell_2}(1) = 1/3,$ and $p_{\ell_3}(1) = 2/3$ defines a symmetric Nash equilibrium fulfilling the conditions of Theorem 2. But strategy p' with $p'_{\ell_1}(21) = 19/21 - 0.001, p'_{\ell_2}(21) = 2/21 + 0.001, p'_{\ell_2}(1) = 1/3 - 0.008,$ and $p'_{\ell_3}(1) = 2/3 + 0.008$ constitutes a successful mutation.

Theorem 2 is the last result required to establish the uniqueness of an ESS for symmetric Bayesian routing games.

Theorem 3 (Uniqueness). *Let B be a symmetric Bayesian routing game with ESS p^* .*

1. *Fix any two links $\ell \neq \ell'$ with the same speed, i.e. $c_\ell = c_{\ell'}$. Then for all task weights w we have $p_\ell^*(w) = p_{\ell'}^*(w)$ and $|\text{support}_\ell(p^*)| \leq 1$.*
2. *The ESS p^* is the unique ESS for B .*

Now we give a structural condition that is sufficient for an ESS. It can be used to construct an ESS in a wide variety of models where the ESS exists. Theorem 2 shows that an ESS requires links to "specialize" in tasks where distinct links do not share two distinct tasks. A stronger condition is to require that if a link is optimal for two distinct tasks, then no other link is optimal for either of the tasks. We call such a distribution clustered.

Definition 5. *In an n -player symmetric Bayesian routing game B , a symmetric strategy profile p^n is **clustered** if for any two distinct tasks w, w' and any link ℓ , if ℓ is optimal for both w and w' , then no other link is optimal for w or w' .*

The next theorem establishes that every clustered symmetric Nash equilibrium is an ESS. The proof can be found in the full version.

Theorem 4 (Clustering). *Every clustered Nash equilibrium is an ESS, but not vice versa. More precisely:*

1. *Let B be a symmetric Bayesian routing game. If $(p^*)^n$ is a clustered Nash equilibrium in B , then p^* is an ESS in B .*
2. *There is a symmetric Bayesian routing game B that has a non-clustered ESS and no clustered ESS.*

Acknowledgements: We thank Funda Ergun for helpful discussions and Tom Friedetzky for the example in Observation 2. We also thank the anonymous reviewer who simplified the proof of Theorem 1.

References

1. P. Berenbrink, L. A. Goldberg, P. Goldberg, and R. Martin. Utilitarian Resource Assignment. *Journal of Discrete Algorithms* Vol. 4 Issue 4, (Dec. 2006), 567-587.
2. George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. *In Proc. of the 37th Annual Symposium on Theory of Computing (STOC)*, pp 67-73, 2005.
3. A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. *In Proc. 13th Annual Symposium on Discrete Algorithms (SODA)*, pp. 413-420, 2002
4. J. C. Ely and W. H. Sandholm Evolution in Bayesian games I: Theory *Games and Economic Behavior*, Volume 53, Issue 1, October 2005, pp 83-109
5. K. Etessami and A. Lochbihler. The Computational Complexity of Evolutionarily Stable Strategies. *Electronic Colloquium on Computational Complexity*, vol. 55, 2004.
6. A. Fabrikant, A. Luthra, E. N. Maneva, C. H. Papadimitriou, S. Shenker. On a network creation game. *In Proc. of 22nd Symposium on Principles of Distributed Computing (PODC 2003)*, pp 347-351, 2003.
7. S. Fischer, and B. Vöcking. On the Evolution of Selfish Routing. *In Proc. 12th European Symposium on Algorithms (ESA)*, pp 323-334, 2004.
8. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. *In Proceedings of the 18th Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp 1-20, 2002.
9. J. C. Harsanyi. Games with Incomplete Information Played by 'Bayesian Players', Parts I, II, and III. *In Management Science*, (1967) 14: pp 159-182, pp 320-334, and pp 486-502.
10. M. S. Kearns and S. Suri. Networks preserving evolutionary equilibria and the power of randomization. *ACM Conference on Electronic Commerce 2006*, ACM, pp. 200-207.
11. S. Kontogiannis and P. Spirakis The Contribution of Game Theory to Complex Systems. *PCI 2005*, P. Bozaris and E.N. Houstis (Eds.): LNCS 3746, pp. 101111, 2005, Springer-Verlag Berlin Heidelberg.
12. E. Koutsoupias and C.H. Papadimitriou. Worst-Case Equilibria. *In Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pp.404-413, 1999.
13. M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. *In Proc. of the 33rd Annual Symposium on Theory of Computing (STOC)*, pp. 510-519, 2001.
14. J. Maynard Smith. Evolution and the Theory of Games. *In Cambridge University Press*, 1982.
15. J.F. Nash. Equilibrium Points in N -Person Games. *In Proc. of the National Academy of Sciences of the United States of America*, 36:48-49.
16. T. Roughgarden and É. Tardos. How Bad is Selfish Routing? *In Journal of the ACM*, 49(2), pp. 236-259, 2002.
17. W. H. Sandholm. Evolutionary Implementation and Congestion Pricing. *In Review of Economic Studies*, 69:667-689, 2002.
18. E. van Damme. Stability and Perfection of Nash Equilibria. *In 2nd Edition*, Berlin: Springer-Verlag 1991.
19. J. Weibull. Evolutionary Game Theory. *Cambridge, MA: The M.I.T. Press*, 1995.