Convergence to Approximate Equilibria in Congestion Games among Coalitions

Dimitris Fotakis[1], Spyros Sarigiannidis[1], and Paul G. Spirakis[2]

[1] Dept. of Information and Communication Systems Engineering,
University of the Aegean, 83200 Samos, Greece.
fotakis@aegean.gr, ssar@aegean.gr

[2] Research Academic Computer Technology Institute,
P.O. Box 1382, N. Kazantzaki Str., 26500 Rio-Patra, Greece. spirakis@cti.gr

Abstract

We present upper bounds on the convergence time to (approximate) Nash equilibria (NE) for load balancing and congestion games with coalitions. We first consider single-commodity linear congestion games with static coalitions, where the selfish cost of each coalition is the total delay of its players. If every coalition is given the opportunity to improve its strategy within a bounded time interval, we show that an approximate NE is reached in polynomial time. This holds even for coalitions of different size and is the first non-trivial upper bound on the rate of convergence to approximate NE for a natural class of *asymmetric* congestion games. We also consider load balancing games on identical links with dynamic coalitions of size 2, where the selfish cost of each coalition is the maximum delay of its players. We prove that a natural family of improvement moves converges to a NE in pseudo-polynomial time and to an approximate NE in polynomial time. The latter is the first polynomial bound on the convergence time to approximate NE for load balancing games with coalitions.

1. Introduction

Congestion games provide a natural model for non-cooperative resource allocation and have been the subject of intensive research in algorithmic game theory. A *congestion game* is a non-cooperative game where selfish players compete over a set of resources. The players' strategies are subsets of resources. The delay of each player from selecting a particular resource is given by a non-negative and non-decreasing function of the load of the resource. The total delay of a player is equal to the sum of the delays for his selected resources. A natural solution concept is that of a (pure) Nash equilibrium (NE), a state where no player can decrease his delay by unilaterally changing his strategy. [Rosenthal (1973)] proved that congestion games NE correspond to the local optimal of a natural potential function. Many recent contributions have provided considerable insight into the structure and efficiency, tractability, and the rate of convergence to NE in congestion games (e.g. [Even-Dar et al (2003), Fabrikant et al (2004), Fotakis et al (2005), Ackerman et al (2006), Chien and Sinclair (2007)]).

In practice however, the competition for resources usually takes place among coalitions of players instead of individuals. For a typical example, one may consider a telecommunication network where antagonistic service providers seek to minimize their operational costs while meeting their customers' demands. In such settings, it is important to know how coalition formation affects the structure, the efficiency, and the rate of convergence to a NE. Motivated by similar considerations, [Hayrapetyan et al (2006)] and [Fotakis et al (2006)] proposed two essentially orthogonal models for investigating the effects of coalition formation in congestion games. On the one hand, the model of *congestion games with coalitions* [Hayrapetyan et al (2006)] uses total delay as coalitions' selfish cost, allows for arbitrary strategy spaces, but is restricted to static coalitions. On the other hand, the model of *load balancing games with coalitions* [Fotakis et al (2006)] uses maximum delay as coalitions' selfish cost, considers dynamic coalitions and weighted players, but is restricted to a set of singleton strategies shared by all players. In this work, we investigate the effect of coalition formation on the rate of convergence to (approximate) NE in both models.

Related Work. [Fabrikant et al (2004)] initiated the study of the complexity of computing a NE in congestion games. [Fabrikant et al (2004)] prove that it is PLS-complete to compute a NE in symmetric congestion games. On the positive side, [Fabrikant et al (2004)] show that in a symmetric *network* congestion game, a NE can be found in polynomial time by a min-cost flow computation.

Given the non-cooperative nature of congestion games, a natural question is whether the players trying to improve their delay converge to a NE in a reasonable amount of time. [Ackerman et al (2006)] introduce the class of matroid congestion games and proves that it essentially coincides with the class of asymmetric congestion games that guarantee fast convergence to a NE. [Chien and Sinclair (2007)] consider *symmetric* congestion games with a weak restriction on resource delays and proves that several natural families of improvement moves converge to an *approximate* NE in polynomial time. For load balancing games, [Even-Dar et al (2003)] show upper and lower bounds on the convergence time to a NE for several natural families of improvement moves.

On the other hand, the effects of coalition formation in congestion games are yet to be well understood. [Hayrapetyan et al (2006)] mostly consider congestion games on parallel links with identical users and convex delays. For this class of games, [Hayrapetyan et al (2006)] establish the existence and tractability of pure NE, presents examples where coalition formation deteriorates the efficiency of NE, and bounds the efficiency loss due to coalition formation.

[Fotakis et al (2006)] present a potential function for linear congestion games with coalitions. For load balancing games with weighted players, [Fotakis et al (2006)] present a generalized potential function and shows exponential-time-convergence to a NE for unrelated links. For identical links and dynamic coalitions of size 2, [Fotakis et al (2006)] show that a natural family of improvement moves converges to a NE in pseudo-polynomial time. In addition, [Fotakis et al (2006)] prove that coalition formation can only improve the efficiency of NE for load balancing games on identical links.

Contribution. In this work, we prove upper bounds on the rate of convergence to (approximate) NE in congestion and load balancing games with coalitions. We start with single-commodity linear congestion games with static coalitions. We restrict our attention to ε-moves, i.e. deviations that improve the coalition's delay by a factor of more than ε. Combining the approach of [Chien and Sinclair (2007)] with the potential function of [Fotakis et al (2006), Theorem 6], we show that if every coalition is given the opportunity to improve its delay within a bounded time interval, an approximate NE is reached in a polynomial number of steps. This bound holds even for coalitions of different size, in which case the game is *not symmetric*. Since the recent results of [Chien and Sinclair (2007)] hold for symmetric games only, this is the first non-trivial upper bound on the convergence time to approximate NE for a natural class of asymmetric congestion games.

We also consider load balancing games on identical links with dynamic coalitions. We introduce a natural family of improvement moves called BEST OF THE BEST RESPONSE, or BoBR in short. For dynamic coalitions of size 2, we prove that any sequence of BoBR moves converges to a refinement of NE in pseudo-polynomial time. In addition, we consider a certain sequence of BoBR ε-moves and show that it converges to an approximate equilibrium in polynomial time. To the best of our knowledge, this is the first polynomial bound on the convergence time to approximate NE for load balancing games with (either static or dynamic) coalitions.

2. Definitions and Notation

Congestion Games with Coalitions. A *congestion game with coalitions* consists of a set of identical players $N=[n]$ partitioned into k coalitions $\{C_1,\ldots,C_k\}$, a set of resources $E=\{e_1,\ldots,e_m\}$, a strategy space $\Sigma_i \subseteq 2^E$ for each player $i \in N$, and a non-negative and non-decreasing delay function $d_e : \mathbf{N} \mapsto \mathbf{N}$ associated with every resource e. In this work, we restrict our attention to games with linear delays of the form $d_e(x) = a_e x + b_e, a_e, b_e \geq 0$, and symmetric strategies (or *single-commodity* congestion games), where all players share the same strategy space, denoted Σ.

The congestion game is played among the coalitions instead of the individual players. We let n_j denote the number of players in coalition C_j. The strategy space of coalition C_j is Σ^{n_j} and the strategy space of the game is $\Sigma^{n_1} \times \ldots \times \Sigma^{n_k}$. A pure strategy $\sigma_j \in \Sigma^{n_j}$ determines a (pure) strategy $\sigma_j^i \in \Sigma$ for every player $i \in C_j$. For every resource $e \in E$, the load (or congestion) of e due to C_j in σ_j is $l_e(\sigma_j) = \left|\{i \in C_j : e \in \sigma_j^i\}\right|$. A tuple $\sigma = (\sigma_1,\ldots,\sigma_k)$ consisting of a pure strategy $\sigma_j \in \Sigma^{n_j}$ for every coalition C_j is a *state* of the game. For every resource $e \in E$, the load of e in σ is $l_e(\sigma) = \sum_{j=1}^{k} l_e(\sigma_j)$. The delay of a strategy $\alpha \in \Sigma$ in state σ is $d_a(\sigma) = \sum_{e \in a} d_e(l_e(\sigma))$.

The selfish cost of each coalition C_j in state σ is given by the *total delay* of its players, denoted $r_j(\sigma)$. Formally, $r_j(\sigma) \equiv \sum_{i \in C_j} d_{\sigma_j^i}(\sigma) = \sum_{e \in E} l_e(\sigma_j) d_e(l_e(\sigma))$.

A state σ is a *Nash equilibrium* if for every coalition C_j and every strategy $\sigma'_j \in \Sigma^{n_j}$, $r_j(\sigma) \le r_j(\sigma_{-j}, \sigma'_j)$, i.e. the total delay of coalition C_j cannot decrease by C_j unilaterally changing its strategy. For every $\varepsilon \in (0,1)$, a state σ is an ε*-Nash equilibrium* if for every coalition C_j and every strategy $\sigma'_j \in \Sigma^{n_j}$, $(1-\varepsilon) r_j(\sigma) \le r_j(\sigma_{-j}, \sigma'_j)$. An ε*-move* of coalition C_j is a deviation from σ_j to σ'_j that decreases the total delay of C_j by more than $\varepsilon \cdot r_j(\sigma)$.

Load Balancing Games with Coalitions. In *load balancing games with coalitions* each player i is associated with an positive integer weight w_i, which can be assigned to any of the parallel links $E = \{e_1,\ldots,e_m\}$. All players share the same strategy space E. We allow the coalitions to be formed even dynamically. The coalitions' strategy spaces, the coalitions' pure strategies, and the states of the games are defined as in congestion games with coalitions. A pure strategy σ_j of coalition C_j determines a link $\sigma_j^i \in E$ for each player $i \in C_j$ to which w_i is assigned. For every link $e \in E$, the load of e due to C_j in σ_j is $l_e(\sigma_j) = \sum_{i \in C_j : e = \sigma_j^i} w_i$.

For every link $e \in E$, the load of e in state σ is $l_e(\sigma) = \sum_{j=1}^{k} l_e(\sigma_j)$. In this paper, we restrict our attention to load balancing games on identical parallel links, where the delay of every link e in σ is equal to its load $l_e(\sigma)$.

The selfish cost of each coalition C_j in state σ is given by the *maximum delay* of its players, denoted $\lambda_j(\sigma)$. Formally, $\lambda_j(\sigma) \equiv \max_{i \in C_j}\{l_{\sigma_j^i}(\sigma)\}$. For a dynamic coalition $C \subseteq [n]$, let $\sigma_C = (\sigma_C^i)_{i \in C}$ denote the combined strategies of C's players in σ, and let $\lambda_C(\sigma) = \max_{i \in C}\{l_{\sigma_C^i}(\sigma)\}$ denote the maximum delay of C in σ.

For static and singleton coalitions, Nash equilibrium provides a natural notion of stability. A state σ is a *Nash equilibrium* if for every coalition C_j and every strategy $\sigma'_j \in E^{n_j}, \lambda_j(\sigma) \leq \lambda_j(\sigma_{-j}, \sigma'_j)$. ε-Nash equilibria and ε-moves are defined as before.

For dynamic coalitions, we use a stronger notion of stability called r-robust equilibria (see also [Fotakis et al (2006)]). A state σ is a r*-robust equilibrium* if for every (even dynamic) coalition $C \subseteq [n]$ of size r and every strategy $\sigma'_C \in E^{|C|}$, $\lambda_C(\sigma) \leq \lambda_C(\sigma_{-C}, \sigma'_C)$. A r-robust equilibrium is a r'-robust equilibrium for every integer $r' \leq r$. An (improvement) *move* of a (dynamic) coalition C is a deviation from strategy σ_C to strategy σ'_C that decreases the maximum delay of C.

For every $\varepsilon \in (0,1)$, a state σ is an (ε, r)*-robust equilibrium* if for every (even dynamic) coalition $C \subseteq [n]$ of size r and every strategy $\sigma'_C \in E^{|C|}$, $(1-\varepsilon)\lambda_C(\sigma) \leq \lambda_C(\sigma_{-C}, \sigma'_C)$. An ε*-move* of a (dynamic) coalition C is a deviation from strategy σ_C to strategy σ'_C that decreases the maximum delay of C by more than $\varepsilon \cdot \lambda_C(\sigma)$.

3. Linear Congestion Games with Static Coalitions

In this section, we consider single-commodity linear congestion games with identical (unit size) players and static coalitions. The selfish cost of each coalition is given by the total delay of its players.

To bound the convergence time to ε-Nash equilibria, we use the potential function $\Phi(\sigma) = \frac{1}{2}[\sum_{e \in E} l_e(\sigma) \cdot d_e(l_e(\sigma)) + \sum_{j=1}^{k} l_e(\sigma_j) \cdot d_e(l_e(\sigma_j))]$. [Fotakis et al (2006)] prove that Φ is an exact potential function for (multi-commodity) linear congestion games with coalitions.

We consider sequences of ε-moves where every coalition with an ε-move available is given a chance to move in a (possibly large but) finite amount of time (see also [Chien and Sinclair (2007), Section 3]). To formalize this notion, we may think of a schedule $j_1, j_2, \ldots, j_t, \ldots$ of coalitions, where coalition j_t is given the opportunity to move at step t. At step t, coalition j_t performs an ε-move if it has one available. Otherwise, nothing happens. For some integer $T \geq k$, a schedule of coalitions $j_1, j_2, \ldots, j_t, \ldots$ is called T*-restricted* if in every subsequence of length T, every coalition appears at least once. The schedule can be determined by an adaptive adversary whose goal is to delay the convergence to an ε-equilibrium as much as possible. The following theorem shows that as long as the schedule is T-restricted, an ε-Nash equilibrium is reached in a polynomial number of steps.

Theorem 1. *In a single-commodity linear congestion game with coalitions, any sequence of* ε*-moves determined by a* T*-restricted schedule converges to an* ε*-Nash equilibrium after at most* $\left\lceil \frac{kn(n+1)}{\varepsilon(1-\varepsilon)} \log \Phi(\sigma^{init}) \right\rceil T$ *steps, where* σ^{init} *is the initial state.*

Proof. The outline of the proof is similar to that of [Chien and Sinclair (2007), Theorem 4.1], which holds for *symmetric* congestion games only. However, coalitions may be of different size, in which case the game is *asymmetric*. Hence, we have to extend the technique of [Chien and Sinclair (2007)] and bound the effect of coalitions of different size.

We start with a proposition showing that a decrease in the total delay of a coalition implies a decrease in the potential, even if the coalition does not move.

Proposition 1. *Let* σ' *be a state reachable from* σ *by a sequence of* ε*-moves in which coalition* j *does not move. Then,* $\Phi(\sigma)-\Phi(\sigma') \geq \varepsilon(r_j(\sigma)-r_j(\sigma'))/n$.

Let us consider a phase $j_0, j_2, \ldots, j_{T-1}$ of length T of a T-restricted schedule. Let σ^t denote the state at time step $t = 0,1,\ldots,T$, with σ^0 denoting the initial state. Successive states may not be distinct, since if coalition j_t does not have an ε-move available in σ^{t-1}, j_t does not move and $\sigma^{t-1} = \sigma^t$. The proof of the theorem follows from the following claim:

Claim. In any phase of length T, the potential decreases by at least $\frac{\varepsilon(1-\varepsilon)}{kn(n+1)}\Phi(\sigma^0)$.

Sketch of Proof. We pick the coalition C_j with the largest total delay in σ^0 and mark that $r_j(\sigma^0) \geq \Phi(\sigma^0)/k$. We examine two possibilities for the first time σ^t that coalition C_j is given the opportunity to move in the particular phase. If there is a coalition C_i such that $r_j(\sigma^t) > \frac{n^2}{1-\varepsilon} r_i(\sigma^t)$, then adopting the strategy of C_i in σ^t is an ε-move for C_j which leads to a decreases of the potential by at least $\varepsilon \cdot r_j(\sigma^t)$. By Proposition 1, $\Phi(\sigma^0)-\Phi(\sigma^t) \geq \varepsilon(r_j(\sigma^0)-r_j(\sigma^t))/n$, and therefore,

$$\Phi(\sigma^0)-\Phi(\sigma^{t+1}) \geq \varepsilon \cdot r_j(\sigma^0)/n \geq \frac{\varepsilon}{kn}\Phi(\sigma^0).$$

If there is no coalition C_i with $r_j(\sigma^t) > \frac{n^2}{1-\varepsilon} r_i(\sigma^t)$, we distinguish between the case where no coalition moves before t and the case where some coalition moves before

t. In the former case , the decrease in the potential due to the first coalition C_p making an ε-move in the current phase at step $t+k$ is at least $\frac{\varepsilon(1-\varepsilon)}{kn^2}\Phi(\sigma^0)$. In the latter case the decrease in the potential up to step t is bounded from below by both the decrease in the potential due to the last coalition C_p making an ε-move before t, at step $t-k$ and the decrease in the potential due to the change in C_j 's total delay (Proposition 1). This leads to a decrease in the potential function by at least $\frac{\varepsilon(1-\varepsilon)}{n(n+1-\varepsilon)} r_j(\sigma^0) \geq \frac{\varepsilon(1-\varepsilon)}{kn(n+1)}\Phi(\sigma^0)$ □

The claim above implies that starting from state σ^{init}, it takes at most $\left\lceil \frac{kn(n+1)}{\varepsilon(1-\varepsilon)} \log \Phi(\sigma^{init}) \right\rceil T$ steps to reach an ε-Nash equilibrium. □

4. Load Balancing Games with Dynamic Coalitions

In this section, we consider load balancing games on identical links with weighted players and dynamic coalitions of size 2. The selfish cost of each coalition is given by the maximum delay of its players. To bound the convergence time to (approximate) 2-robust equilibria, we use the potential function $F(\sigma) = \sum_{e\in E}(l_e(s))^2$. $F(\sigma)$ is a weighted potential function for load balancing games among individual players on identical links [Even-Dar et al (2003)]. Even though, $F(\sigma)$ is not a potential function for load balancing games with coalitions of any size, [Fotakis et al (2006)] show that for dynamic coalitions of size 2, $F(\sigma)$ is a generalized potential for a family of improvement moves called SMALLER COALITIONS FIRST.

4.1 Best of the Best Response

In this section, we introduce a natural family of improvement moves called BEST OF THE BEST RESPONSE (BoBR). We prove that $F(\sigma)$ is a generalized potential for BoBR(2) moves, i.e. $F(\sigma)$ decreases every time a coalition of size (at most) 2 makes a BoBR move and improves its maximum delay. Therefore, every sequence of BoBR(2) moves converges to a 2-robust equilibrium in at most $W_{tot}^2/2$ steps.

Let (i,j) be a (possibly dynamic) coalition of two players. Wlog. we assume that $w_i \geq w_j$. Let σ be the current state where players i and j are assigned to links e_i and e_j respectively, and let $l_e(\sigma_{-\{i,j\}})$ be the load of link e in σ due to players other than i and j.

The best of the best response of coalition (i, j) is to assign the largest weight w_i to the link with the minimum load caused by other players, and the smallest weight w_j to the link with the minimum load caused by other players and w_i's new assignment.

We call BEST OF THE BEST RESPONSE, or BoBR(2) in short, the family of improvement moves where every coalition (i, j) switches to the links (e'_i, e'_j) defined as above. Assuming integer weights, we prove that any sequence of BoBR(2) moves converges to a 2-robust equilibrium in at most $W_{tot}^2/2$ steps.

Theorem 2. $F(\sigma) = \sum_{e \in E} (l_e(s))^2$ *is a generalized potential for BoBR(2). Additionally, BoBR(2) converges to a 2-robust equilibrium in at most* $W_{tot}^2/2$ *steps.*

Sketch of Proof. Let σ be the current state where players i and j are assigned to links e_i and e_j respectively. We show that $F(\sigma)$ decreases by at least 2 when coalition (i, j) changes its strategy from (e_i, e_j) to (e'_i, e'_j). The proof proceeds by case analysis. For succinctness of notation, we let

$$x_i = l_{e_i}(\sigma_{-\{i,j\}}) = l_{e_i}(\sigma) - w_i \qquad y_i = l_{e'_i}(\sigma_{-\{i,j\}})$$
$$x_j = l_{e_j}(\sigma_{-\{i,j\}}) = l_{e_j}(\sigma) - w_j \qquad y_j = l_{e'_j}(\sigma_{-\{i,j\}})$$

Most of the cases are easily reduced to the special case of an individual player's improvement move, for which $F(\sigma)$ decreases by at least 2. The most interesting case is when $e_i \neq e_j$ and $e'_i \neq e'_j$. The change in the potential due to (i, j)'s move from (e_i, e_j) to (e'_i, e'_j) is

$$F(\sigma) - F(\sigma') = 2w_i(x_i - y_i) + 2w_j(x_j - y_j) \qquad \textbf{(1)}$$

Since e'_i minimizes the load caused by players other than i and j, $y_i \leq x_i$. If $y_j \leq x_j$, at least one of the inequalities $y_i \leq x_i$ and $y_j \leq x_j$ is strict, because the maximum delay of i and j decreases when (i, j) moves from (e_i, e_j) to (e'_i, e'_j). Then $F(\sigma) - F(\sigma') \geq 2$ follows by weight integrality. On the other hand, if $y_j > x_j$, then $y_i = x_j < y_j \leq x_i$. Since the maximum delay of i and j decreases when (i, j) moves from (e_i, e_j) to (e'_i, e'_j), $x_i + w_i > y_j + w_j$. Then, by weight integrality, at least one of $w_i \geq w_j + 1$ and $x_i \geq y_j + 1$ holds. In both cases, starting from (1) and using simple algebra, we conclude that $F(\sigma) - F(\sigma') \geq 2$.

Therefore, in all cases, the potential decreases by at least 2. Initially, it is at most W_{tot}^2. Since it remains positive, any sequence of BoBR(2) takes less than $W_{tot}^2/2$ moves. □

4.2 Fast Convergence to Approximate Equilibria

Seeking for a family of moves that converges fast to approximate 2-robust equilibria, we consider a subfamily of BoBR(2), called BoBR$(\varepsilon,2)$. Given the current state σ, the next BoBR$(\varepsilon,2)$ move is determined by the following rule:

1. If σ is not a 1-equilibrium, let i be the maximum weight player who can improve his delay by switching to his best response link $e^{\min}(\sigma)$.
2. If σ is a 1-equilibrium, let (i,j) be any dynamic coalition of size 2 with an ε-move available. Coalition (i,j) moves from (e_i,e_j) to its BoBR (e'_i,e'_j).

Theorem 2 implies that every BoBR$(\varepsilon,2)$ move decreases the potential $F(\sigma)$ and BoBR$(\varepsilon,2)$ converges to a $(\varepsilon,2)$-robust equilibrium. The following theorem shows that BoBR$(\varepsilon,2)$ reaches a $(\varepsilon,2)$-robust equilibrium after a polynomial number steps.

Theorem 3. *Starting from any initial state, BoBR$(\varepsilon,2)$ converges to $(\varepsilon,2)$-robust equilibrium in at most $n\cdot\lceil 8m/\varepsilon^2\rceil$ steps.*

Proof. The first step of BoBR$(\varepsilon,2)$ reaches a 1-equilibrium in at most n steps [Even-Dar et al (2003)] and can only decrease the potential. Hence, it suffices to show that BoBR$(\varepsilon,2)$ makes at most $\lceil 8m/\varepsilon^2\rceil$ BoBR ε-moves for coalitions of size 2. To establish this claim, we show that every time a coalition of 2 players makes a BoBR ε-move, the potential decreases by at least $\frac{\varepsilon^2}{2m^2}W_{tot}^2$.

Let σ be a 1-equilibrium, and let (i,j) be a coalition of 2 players with an ε-move available. As before, we assume that $w_i \ge w_j$ and we use the notation e_i,e_j,e'_i,e'_j introduced in Section 4.1 with exactly the same meaning. Additionally, without loss of generality we assume that $w_{\max} < W_{tot}/m$. We start with two simple propositions claims concerning the links loads and the best of the best response of (i,j) in σ.

Proposition 2. *Let σ be any 1-equilibrium for a load balancing game on identical links. If $w_{\max} < W_{tot}/m$, then $l^{\max}(\sigma) < 2W_{tot}/m$ and $l^{\min}(\sigma) \ge \frac{1}{2}W_{tot}/m$.*

Proposition 3. *Let σ be any 1-equilibrium for a load balancing game on identical links, and let (i,j) be a (dynamic) coalition with an improvement move available in σ. Then the BoBR of (i,j) is i and j to swap their links, i.e. $e'_i = e_j$ and $e'_j = e_i$.*

Let σ' be the new state after coalition (i,j) best of the best response. For simplicity of notation, let $x_i = l_{e_i}(\sigma) - w_i$, and let $x_j = l_{e_j}(\sigma) - w_j$. Using (1), we obtain that the potential decreases by $F(\sigma) - F(\sigma') = 2(w_i - w_j)(x_i - x_j)$. Since coalition (i, j) has an ε-move available in σ, its best of the best response decreases the maximum delay of (i, j) by a factor of more than ε. By Proposition 3, σ being an 1-equilibrium implies that $l_{e_i}(\sigma) \geq l_{e_j}(\sigma)$ and therefore the maximum delay of (i, j) in σ is $x_i + w_i$. Moreover, the maximum delay of (i, j) in σ' is $\max\{x_i + w_j, x_j + w_i\}$. Hence, $(1-\varepsilon)(x_i + w_i) > \max\{x_i + w_j, x_j + w_i\}$. This implies that $x_i - x_j > \varepsilon(x_i + w_i)$ and $w_i - w_j > \varepsilon(x_i + w_i)$. Therefore, $F(\sigma) - F(\sigma') > 2\varepsilon^2 (l_{e_i}(\sigma))^2 \geq \frac{\varepsilon^2}{2m^2} W_{tot}^2$, since σ is a 1-equilibrium and $l^{\min}(\sigma) \geq \frac{1}{2} W_{tot}/m$ (Proposition 2).

By Proposition 2, the potential before the first coalitional ε-move is less than $4W_{tot}^2/m$. Additionally, the potential cannot drop below W_{tot}^2/m, a value achieved when a load of W_{tot}/m is assigned to every link. Since every coalitional ε-move decreases the potential by at least $\frac{\varepsilon^2}{2m^2} W_{tot}^2$, an $(\varepsilon,2)$-robust equilibrium is reached after at most $n \cdot \lceil 8m/\varepsilon^2 \rceil$ coalitional BoBR ε-moves. □

References

H. Ackerman, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. FOCS ’06, pp. 613–622, 2006.

S. Chien and A. Sinclair. Convergence to Approximate Nash Equilibria in Congestion Games. SODA ’07, 2007.

E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence Time to Nash Equilibria. ICALP ’03, pp. 502–513, 2003.

A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. STOC ’04, pp. 604–612, 2004.

D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. Theoretical Computer Science, 348:226–239, 2005.

D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic Congestion Games among Coalitions. ICALP ’06, pp. 573–584, 2006.

A. Hayrapetyan, É. Tardos, and T. Wexler. The Effect of Collusion in Congestion Games. STOC ’06, pp. 89–98, 2006.

R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. International Journal of Game Theory, 2:65–67, 1973.