

Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure

IOANNIS CARAGIANNIS, University of Patras & CTI
ANGELO FANELLI, Nanyang Technological University
NICK GRAVIN, Nanyang Technological University
ALEXANDER SKOPALIK, TU Dortmund

We consider structural and algorithmic questions related to the Nash dynamics of weighted congestion games. In weighted congestion games with linear latency functions, the existence of pure Nash equilibria is guaranteed by potential function arguments. Unfortunately, this proof of existence is inefficient and computing pure Nash equilibria in such games is a PLS-hard problem even when all players have unit weights. The situation gets worse when superlinear (e.g., quadratic) latency functions come into play; in this case, the Nash dynamics of the game may contain cycles and pure Nash equilibria may not even exist. Given these obstacles, we consider approximate pure Nash equilibria as alternative solution concepts. Do such equilibria exist? And if so, can we compute them efficiently?

We provide positive answers to both questions for weighted congestion games with polynomial latency functions by exploiting an “approximation” of such games by a new class of potential games that we call Ψ -games. This allows us to show that these games have d -approximate pure Nash equilibria, where d is the maximum degree of the latency functions. Our main technical contribution is an efficient algorithm for computing $O(1)$ -approximate pure Nash equilibria when d is a constant. For games with linear latency functions, the approximation guarantee is $\frac{3+\sqrt{5}}{2} + O(\gamma)$ for arbitrarily small $\gamma > 0$; for latency functions with maximum degree $d \geq 2$, it is $d^{2d+o(d)}$. The running time is polynomial in the number of bits in the representation of the game and $1/\gamma$. As a byproduct of our techniques, we also show the following interesting structural statement for weighted congestion games with polynomial latency functions of maximum degree $d \geq 2$: polynomially-long sequences of best-response moves from any initial state to a $d^{O(d^2)}$ -approximate pure Nash equilibrium exist and can be efficiently identified in such games as long as d is constant.

To the best of our knowledge, these are the first positive algorithmic results for approximate pure Nash equilibria in weighted congestion games. Our techniques significantly extend our recent work on unweighted congestion games through the use of Ψ -games. The concept of approximating non-potential games by potential ones is interesting in itself and might have further applications.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences-Economics

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Pure Nash equilibria, potential games

This work was partially supported by the grant NRF-RF2009-08 “Algorithmic aspects of coalitional games”, the EC-funded STREP Project FP7-ICT-258307 EULER and by Deutsche Forschungsgemeinschaft (DFG) within the Collaborative Research Center SFB 876 “Providing Information by Resource-Constrained Analysis”, project A2.

Author’s addresses: I. Caragiannis, Computer Technology Institute & Department of Computer Engineering and Informatics, University of Patras, Greece; A. Fanelli and N. Gravin, Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore; A. Skopalik, Department of Computer Science, TU Dortmund, Germany.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC’12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

1. INTRODUCTION

Among other solution concepts, the notion of the pure Nash equilibrium plays a central role in Game Theory. Pure Nash equilibria in a game characterize situations with non-cooperative deterministic players in which no player has any incentive to unilaterally deviate from the current situation in order to achieve a higher payoff. Unfortunately, it is well known that there are games that do not have pure Nash equilibria. Furthermore, even in games where the existence of equilibria is guaranteed, their computation can be a computationally hard task. Such negative results significantly question the importance of pure Nash equilibria as solution concepts that characterize the behavior of rational players.

Approximate pure Nash equilibria, which characterize situations where no player can *significantly improve* her payoff by unilaterally deviating from her current strategy, could serve as alternative solution concepts¹ provided that they exist and can be computed efficiently. In this paper, we present the first positive algorithmic results for approximate pure Nash equilibria in weighted congestion games. Our main contribution is a polynomial-time algorithm that computes $O(1)$ -approximate pure Nash equilibria under mild restrictions on the game parameters; these restrictions apply to important subclasses of games in which not even the existence of such approximate equilibria was known prior to our work.

Problem statement and related work. In a weighted congestion game, players compete over a set of resources. Each player has a positive weight. Each resource incurs a latency to all players that use it; this latency depends on the total weight of the players that use the resource according to a resource-specific, non-negative, and non-decreasing latency function. Among a given set of strategies (over sets of resources), each player aims to select one selfishly, trying to minimize her individual total cost, i.e., the sum of the latencies on the resources in her strategy. Typical examples include weighted congestion games in networks, where the network links correspond to the resources and each player has alternative paths that connect two nodes as strategies.

The case of unweighted congestion games (i.e., when all players have unit weight) has been widely studied in the literature. Rosenthal [1973] proved that these games admit a potential function with the following remarkable property: the difference in the potential value between two states (i.e., two snapshots of strategies) that differ in the strategy of a single player equals to the difference of the cost experienced by this player in these two states. This immediately implies the existence of a pure Nash equilibrium. Any sequence of improvement moves by the players strictly decreases the value of the potential and a state corresponding to a local minimum of the potential will eventually be reached; this corresponds to a pure Nash equilibrium. For weighted congestion games, potential functions are known only in the case where the latency functions are linear or exponential (see [Fotakis et al. 2005; Harks and Klimm 2010; Panagopoulou and Spirakis 2006]). Actually, in games with polynomial latency functions (of constant maximum degree higher than 1), pure Nash equilibria may not exist [Goemans et al. 2005; Harks and Klimm 2010]. In general, the problem of deciding whether a given weighted congestion game has a pure Nash equilibrium is NP-hard [Dunkel and Schulz 2008].

Potential functions provide only inefficient proofs of existence of pure Nash equilibria. Fabrikant et al. [2004] proved that the problem of computing a pure Nash equilibrium in a (unweighted) congestion game is PLS-complete (informally, as hard as

¹Actually, approximate pure Nash equilibria may be more desirable as solution concepts in practical decision making settings since they can accommodate small modeling inaccuracies due to uncertainty (e.g., see the arguments in [Christin et al. 2004]).

it could be, given that there is an associated potential function; see [Johnson et al. 1988]). This negative result holds even in the case of linear latency functions [Ackermann et al. 2008]. One consequence of PLS-completeness results is that almost all states in some congestion games are such that any sequence of players' improvement moves that originates from these states and reaches pure Nash equilibria is exponentially long. Such phenomena have been observed even in very simple weighted congestion games (see [Ackermann et al. 2009; Even-Dar et al. 2007]). Efficient algorithms are known only for special cases. For example, Fabrikant et al. [2004] show that the Rosenthal's potential function can be (globally) minimized efficiently by a flow computation in unweighted congestion games in networks when the strategy sets of the players are symmetric.

The above negative results have led to the study of the complexity of approximate pure Nash equilibria (or, simply, approximate equilibria). A ρ -approximate (pure Nash) equilibrium is a state, from which no player has an incentive to deviate so that she decreases her cost by a factor larger than ρ . In our recent work [Caragiannis et al. 2011], we present an algorithm for computing $O(1)$ -approximate equilibria in unweighted congestion games with polynomial latency functions of constant maximum degree. The restriction on the latency functions is necessary since, for more general latency functions, Skopalik and Vöcking [2008] show that the problem is still PLS-complete for any polynomially computable ρ (see also the discussion in [Caragiannis et al. 2011]). Improved bounds are known for special cases. For symmetric unweighted congestion games, Chien and Sinclair [2011] prove that the $(1 + \epsilon)$ -improvement dynamics converges to a $(1 + \epsilon)$ -approximate equilibrium after a polynomial number of steps; this result holds under mild assumptions on the latency functions and the participation of the players in the dynamics. Efficient algorithms for approximate equilibria have been recently obtained for other classes of games such as constraint satisfaction [Bhalgat et al. 2010; Nguyen and Tardos 2009], anonymous games [Daskalakis and Papadimitriou 2007], network formation [Anshelevich and Caskurlu 2009], and facility location games [Cardinal and Hoefer 2010].

In light of the negative results mentioned above, several authors have considered other properties of the dynamics of congestion games. The papers [Awerbuch et al. 2008; Fanelli and Moscardelli 2009; Goemans et al. 2005] consider the question of whether efficient states (in the sense that the total cost of the players, or social cost, is small compared to the optimum one) can be reached by best-response moves in linear weighted congestion games. In particular, Awerbuch et al. [2008] show that using almost unrestricted sequences of $(1 + \epsilon)$ -improvement best-response moves, the players rapidly converge to efficient states. Unfortunately, these states are not approximate equilibria, in general. Similar approaches have been followed in the context of other games as well, such as multicast [Charikar et al. 2008; Chekuri et al. 2007], cut [Christodoulou et al. 2006], and valid-utility games [Mirrokni and Vetta 2004].

Our contribution. To the best of our knowledge, no efficient algorithm for computing approximate equilibria is known for (any broad enough subclass of) weighted congestion games. We fill this gap by presenting an algorithm for computing $O(1)$ -approximate equilibria in weighted congestion games with polynomial latency functions of constant maximum degree. For games with linear latency functions, the approximation guarantee is $\frac{3+\sqrt{5}}{2} + O(\gamma)$ for arbitrarily small $\gamma > 0$; for latency functions of maximum degree $d \geq 2$, it is $d^{2d+o(d)}$. The algorithm runs in time that is polynomial in the number of bits in the representation of the game and $1/\gamma$.

This result is much more surprising than it looks at first glance. In particular, weighted congestion games with superlinear latency functions do not admit potential functions, the main tool that is exploited by all known positive algorithmic results

for (approximate) equilibria in congestion games. Given this, it is not even clear that $O(1)$ -approximate equilibria exist. In order to bypass this obstacle, we introduce a new class of potential games (that we call Ψ -games), which “approximate” weighted congestion games with polynomial latency functions in the following sense. Ψ -games of degree 1 are linear weighted congestion games. Each weighted congestion game of degree $d \geq 2$ has a corresponding Ψ -game of degree d defined in such a way that any ρ -approximate equilibrium in the latter is a $d!\rho$ -approximate equilibrium for the former. As an intermediate new result, we obtain that weighted congestion games with polynomial latency functions of degree d have $d!$ -approximate equilibria.

So, our algorithm is actually applied to Ψ -games. It has a simple general structure, similar to our recent algorithm for unweighted congestion games [Caragiannis et al. 2011], but has also important differences that are due to the dependency of the cost of each player on the weights of other players. Given a Ψ -game of degree d and an arbitrary initial state, the algorithm computes a sequence of best-response player moves of length that is bounded by a polynomial in the number of bits in the representation of the game and $1/\gamma$. The sequence consists of phases so that the players that participate in each phase experience costs that are polynomially related. This is crucial in order to obtain convergence in polynomial time. Within each phase, the algorithm coordinates the best-response moves according to two different but simple criteria; this is the main tool that guarantees that the effect of a phase to previous ones is negligible and, eventually, an approximate equilibrium is reached. The approximation guarantee is slightly higher than a quantity that characterizes the potential functions of Ψ -games; this quantity (which we call the *stretch*) is defined as the worst-case ratio of the potential value at an almost exact pure Nash equilibrium over the globally optimum potential value and is almost $\frac{3+\sqrt{5}}{2}$ for linear weighted congestion games and $d^{d+o(d)}$ for Ψ -games of degree $d \geq 2$. Our analysis follows the same main steps as in our recent paper [Caragiannis et al. 2011] but uses significantly more involved arguments due to the definition of Ψ -games.

We also present a similar but slightly inferior algorithm that is applied directly to weighted congestion games of maximum degree $d \geq 2$ and reveals a rather surprising structural property of their Nash dynamics: starting from any initial state, the algorithm identifies a polynomially-long sequence of best-response moves that lead to a $d^{O(d^2)}$ -approximate equilibrium. Even though the definition of this algorithm does not make any use of properties of Ψ -games, the analysis is heavily based on them, similarly to the analysis of our main algorithm.

We remark that, following the classical definition of polynomial latency functions in the literature, we assume that they have non-negative coefficients. This is a necessary limitation since the problem of computing a ρ -approximate equilibrium in (unweighted) congestion games with linear latency functions with negative offsets is PLS-complete for any polynomial-time computable $\rho \geq 1$ [Caragiannis et al. 2011].

Roadmap. We begin with preliminary general definitions in Section 2. Section 3 is devoted to Ψ -games and their properties. We present our algorithm and its analysis in Section 4 and conclude with open problems in Section 5. Due to lack of space, many proofs as well as our structural result have been omitted from this version.

2. DEFINITIONS AND PRELIMINARIES

In general, a *game* can be defined as follows. It has a set of n players \mathcal{N} ; each player $u \in \mathcal{N}$ has a set of available strategies Σ_u . A snapshot of strategies, with one strategy per player, is called a *state*. Each state $S \in \prod_{u \in \mathcal{N}} \Sigma_u$ incurs a positive cost $c_u(S)$ to player u . Players act selfishly; each of them aims to select a strategy that minimizes her cost, given the strategies of the other players. Given a state S and a strat-

egy $s'_u \in \Sigma_u$ for player u , we denote by (S_{-u}, s'_u) the state obtained from S when player u deviates to strategy s'_u . For a state S , an *improvement move* (or, simply, a *move*) for player u is the deviation to any strategy s'_u that (strictly) decreases her cost, i.e., $c_u(S_{-u}, s'_u) < c_u(S)$. For $\rho \geq 1$, such a move is called a ρ -*move* if it satisfies $c_u(S_{-u}, s'_u) < \frac{c_u(S)}{\rho}$. A *best-response move* is a move that minimizes the cost of the player (of course, given the strategies of the other players). So, from state S , a move of player u to strategy s_u is a best-response move (and is denoted by $BR_u(S)$) when $c_u(S_{-u}, s_u) = \min_{s \in \Sigma_u} c_u(S_{-u}, s)$. A state S is called a *pure Nash equilibrium* (or, simply, an *equilibrium*) when $c_u(S) \leq c_u(S_{-u}, s'_u)$ for every player $u \in \mathcal{N}$ and every strategy $s'_u \in \Sigma_u$. In this case, we say that no player has (any incentive to make) a move. Similarly, a state is called a ρ -*approximate pure Nash equilibrium* (henceforth called, simply, a ρ -*approximate equilibrium*) when no player has a ρ -move. Also, a state is called a ρ -approximate equilibrium for a subset of players $A \subseteq \mathcal{N}$ if no player in A has a ρ -move. We use the term *Nash dynamics* of a game in order to refer to the directed graph with nodes that correspond to the possible states of the game and directed edges that indicate improvement player moves; pure Nash equilibria correspond to sinks of the Nash dynamics.

A *weighted congestion game* \mathcal{G} can be represented by the tuple $(N, E, (w_u)_{u \in \mathcal{N}}, (\Sigma_u)_{u \in \mathcal{N}}, (f_e)_{e \in E})$. There is a set of n *players* \mathcal{N} and a set of *resources* E . Each player u has a positive weight w_u and a set of available *strategies* Σ_u ; each strategy s_u in Σ_u consists of a non-empty set of resources, i.e., $s_u \subseteq 2^E$. Each resource $e \in E$ has a non-negative and non-decreasing *latency function* f_e defined over non-negative reals, which denotes the latency incurred to the players using resource e ; this latency depends on the total weight of players whose strategies include the particular resource. For a state S , let us define $N_e(S)$ to be the multi-set of the weights of the players that use resource e in S , i.e., $N_e(S) = \{w_u : u \in \mathcal{N} \text{ such that } e \in s_u\}$. Also, we use the notation $L(A)$ to denote the sum of the elements of a finite multi-set of reals A . Then, the latency incurred by resource e to a player u that uses it is $f_e(L(N_e(S)))$. The *cost* of a player u at a state S is the total latency she experiences at the resources in her strategy s_u multiplied by her weight, i.e., $c_u(S) = w_u \sum_{e \in s_u} f_e(L(N_e(S)))$. We consider weighted congestion games in which the resources have polynomial latency functions with (integer) maximum degree $d \geq 1$ with non-negative coefficients. More precisely, the latency function of resource e is $f_e(x) = \sum_{k=0}^d a_{e,k} x^k$ with $a_{e,k} \geq 0$. The special case of linear weighted congestion games (i.e., with latency functions of degree 1) is of particular interest. In general, the size of the representation of a weighted congestion game is the number of bits required to represent the parameters $a_{e,k}$ of the latency functions, the weights of the players, and their strategy sets. In weighted congestion games in networks, the network links are the resources. Each player u aims to connect a pair of nodes (s_u, t_u) and her strategies are all paths connecting s_u with t_u in the network. Note that the representation of such games does not need to keep the whole set of strategies explicitly; it just has to represent the parameters $a_{e,k}$, the weight and the source-destination node pair of each player, and the network.

Unweighted congestion games (i.e., when $w_u = 1$ for each player $u \in \mathcal{N}$) as well as linear weighted congestion games are potential games. They admit a *potential function* $\Phi : \prod_u \Sigma_u \mapsto \mathbb{R}^+$, defined over all states of the game, with the following property: for any two states S and (S_{-u}, s'_u) that differ only in the strategy of player u , it holds that $\Phi(S_{-u}, s'_u) - \Phi(S) = c_u(S_{-u}, s'_u) - c_u(S)$. Clearly, the local minima of the potential function correspond to states that are pure Nash equilibria. The existence of a potential function also implies that the Nash dynamics of the corresponding game is acyclic. Potential functions for the two classes of games mentioned above have been presented by Rosenthal [1973] and Fotakis et al. [2005], respectively. Unfortunately, weighted

congestion games with polynomial latency functions of degree at least 2 are not potential games and may not even have pure Nash equilibria [Goemans et al. 2005; Harks and Klimm 2010].

3. Ψ -GAMES

Our aim in this section is to define a new class of games which we call Ψ -games and study their properties. We will need the following interesting family of functions which have also been used in [Caragiannis 2009] in a slightly different context.

Definition 3.1. For integer $k \geq 0$, the function Ψ_k mapping finite multi-sets of reals to the reals is defined as follows: $\Psi_k(\emptyset) = 0$ for any integer $k \geq 1$, $\Psi_0(A) = 1$ for any (possibly empty) multi-set A , and for any non-empty multi-set $A = \{\alpha_1, \alpha_2, \dots, \alpha_\ell\}$ and integer $k \geq 1$,

$$\Psi_k(A) = k! \sum_{1 \leq d_1 \leq \dots \leq d_k \leq \ell} \prod_{t=1}^k \alpha_{d_t}.$$

So, $\Psi_k(A)$ is essentially the sum of all monomials of total degree k on the elements of A . Each term in the sum has coefficient $k!$. Clearly, $\Psi_1(A) = L(A)$. For $k \geq 2$, compare $\Psi_k(A)$ with $L(A)^k$ which can also be expressed as the sum of the same terms, albeit with different coefficients in $\{1, \dots, k!\}$, given by the multinomial theorem.

We are ready to define Ψ -games. A Ψ -game \mathcal{G} of (integer) degree $d \geq 1$ can be represented by the tuple $(\mathcal{N}, E, (w_u)_{u \in \mathcal{N}}, (\Sigma_u)_{u \in \mathcal{N}}, (a_{e,k})_{e \in E, k=0,1,\dots,d})$. Similarly to weighted congestion games, there is a set of n players \mathcal{N} and a set of resources E . Each player u has a weight w_u and a set of available strategies Σ_u ; each strategy $s_u \in \Sigma_u$ consists of a non-empty set of resources, i.e., $s_u \subseteq E$. Each resource e is associated with $d+1$ non-negative numbers $a_{e,k}$ for $k = 0, 1, \dots, d$. Again, for a state S , we define $N_e(S)$ to be the multi-set of weights of the players that use resource e at state S . Then, the cost of a player u at a state S is defined as

$$\hat{c}_u(S) = w_u \sum_{e \in s_u} \sum_{k=0}^d a_{e,k} \Psi_k(N_e(S)).$$

Of course, the general definitions in the beginning of Section 2 apply also to Ψ -games. With some abuse in notation, we also use $\mathbf{0}$ to refer to the pseudo-state in which no player selects any strategy and $\mathcal{BR}_u(\mathbf{0})$ to denote the best-response of player u assuming that no other player participates in the game.

Clearly, given a weighted congestion game with polynomial latency functions of maximum degree d , there is a corresponding Ψ -game with degree d , i.e., the one with the same sets of players, resources, and strategy sets, and parameters $a_{e,k}$ for each resource e and integer $k = 0, 1, \dots, d$ equal to the corresponding coefficient of the latency function f_e . Observe that Ψ -games of degree 1 are linear weighted congestion games. As we will see below, in a sense, a Ψ -game of degree $d \geq 2$ is an approximation of its corresponding weighted congestion game.

We remark here that a different approximation of weighted congestion games has been recently considered by Kollias and Roughgarden [2011]. Given a weighted congestion game, they define a new game by answering the following question: how should the product of the total weight of the players that use the resource times its latency be shared as cost among these players so that the resulting game is a potential game? Their games use a different sharing than the weight-proportional one used by weighted congestion games. In contrast, our approach is to define an artificial latency on each resource e (by replacing the term $a_{e,k} L(N_e(S))^k$ with $a_{e,k} \Psi_k(N_e(S))$ in the latency functions) so that weight-proportional sharing yields a potential game. This guarantees the

relation between approximate equilibria in weighted congestion games and Ψ -games presented in Lemma 3.5 below, which is crucial for our purposes. Also, the idea of approximating non-potential games by potential ones has been recently used in a slightly different context by Candogan et al. [2011a; 2011b].

3.1. Properties of Ψ -games

The following lemma is proved in (the full version of) [Caragiannis 2009] and is extensively used in our proofs.

LEMMA 3.2. *For any integer $k \geq 1$, any finite multi-set of non-negative reals A , and any non-negative real b the following hold:*

$$\begin{aligned} a. L(A)^k &\leq \Psi_k(A) \leq k!L(A)^k & d. \Psi_k(A \cup \{b\}) - \Psi_k(A) &= kb\Psi_{k-1}(A \cup \{b\}) \\ b. \Psi_{k-1}(A)^k &\leq \Psi_k(A)^{k-1} & e. \Psi_k(A) &\leq k\Psi_1(A)\Psi_{k-1}(A) \\ c. \Psi_k(A \cup \{b\}) &= \sum_{t=0}^k \frac{k!}{(k-t)!} b^t \Psi_{k-t}(A) & f. \Psi_k(A \cup \{b\}) &\leq (\Psi_k(\{b\})^{1/k} + \Psi_k(A)^{1/k})^k \end{aligned}$$

We now present a very important property of Ψ -games.

THEOREM 3.3. *The function $\Phi(S) = \sum_e \sum_{k=0}^d \frac{a_{e,k}}{k+1} \Psi_{k+1}(N_e(S))$ is a potential function for Ψ -games of degree d .*

PROOF. Consider a player u , a state S in which u plays strategy s_u and state (S_{-u}, s'_u) where u has deviated to strategy s'_u . Using the definition of the potential function, we have

$$\begin{aligned} \Phi(S) - \Phi(S_{-u}, s'_u) &= \sum_e \sum_{k=0}^d \frac{a_{e,k}}{k+1} \Psi_{k+1}(N_e(S)) - \sum_e \sum_{k=0}^d \frac{a_{e,k}}{k+1} \Psi_{k+1}(N_e(S_{-u}, s'_u)) \\ &= \sum_{e \in s_u \setminus s'_u} \sum_{k=0}^d \frac{a_{e,k}}{k+1} (\Psi_{k+1}(N_e(S)) - \Psi_{k+1}(N_e(S_{-u}, s'_u))) \\ &\quad + \sum_{e \in s'_u \setminus s_u} \sum_{k=0}^d \frac{a_{e,k}}{k+1} (\Psi_{k+1}(N_e(S)) - \Psi_{k+1}(N_e(S_{-u}, s'_u))) \\ &= \sum_{e \in s_u \setminus s'_u} \sum_{k=0}^d a_{e,k} w_u \Psi_k(N_e(S)) - \sum_{e \in s'_u \setminus s_u} \sum_{k=0}^d a_{e,k} w_u \Psi_k(N_e(S_{-u}, s'_u)) \\ &= w_u \sum_{e \in s_u} \sum_{k=0}^d a_{e,k} \Psi_k(N_e(S)) - w_u \sum_{e \in s'_u} \sum_{k=0}^d a_{e,k} \Psi_k(N_e(S_{-u}, s'_u)) \\ &= \hat{c}_u(S) - \hat{c}_u(S_{-u}, s'_u). \end{aligned}$$

The third equality follows by Lemma 3.2d and the facts that $N_e(S) = N_e(S_{-u}, s'_u) \cup \{w_u\}$ for every resource $e \in s_u \setminus s'_u$ and $N_e(S_{-u}, s'_u) = N_e(S) \cup \{w_u\}$ for every resource $e \in s'_u \setminus s_u$. The last equality follows by the definition of \hat{c}_u . \square

As a corollary, we conclude that the Nash dynamics of Ψ -games are acyclic; hence, these games admit pure Nash equilibria. Recall that Ψ -games of degree 1 are linear weighted congestion games; for this specific case, Theorem 3.3 has been proved in [Fotakis et al. 2005].

In the following, we study the relation between the approximation guarantee of a state for a Ψ -game and its corresponding weighted congestion game with polynomial latency functions.

CLAIM 3.4. Consider a weighted congestion game with polynomial latency functions of degree d and its corresponding Ψ -game. Then, for each player u and state S , $c_u(S) \leq \hat{c}_u(S) \leq d!c_u(S)$.

Using Claim 3.4, we can obtain a relation between approximate equilibria as well.

LEMMA 3.5. Any ρ -approximate pure Nash equilibrium for a Ψ -game of degree d is a $d!\rho$ -approximate pure Nash equilibrium for the corresponding weighted congestion game with polynomial latencies.

Since pure Nash equilibria always exist in Ψ -games, the last statement (applied with $\rho = 1$) implies the following.

THEOREM 3.6. Every weighted congestion game with polynomial latency functions of maximum degree d has a $d!$ -approximate pure Nash equilibrium.

3.2. Subgames and partial potentials

We now define restrictions of the potential function of Ψ -games. Given a state S and a set of players $A \subseteq \mathcal{N}$, we denote by $N_e^A(S)$ the multiset of the weights of players in A that use resource e in S . Then, we define

$$\Phi^A(S) = \sum_e \sum_{k=0}^d \frac{a_{e,k}}{k+1} \Psi_{k+1}(N_e^A(S)).$$

We can think of Φ^A as the potential of a *subgame* in which only the players of A participate.

We also use the notion of the *partial* potential to account for the contribution of subsets of players to the potential function. Consider sets of players A and B with $B \subseteq A \subseteq \mathcal{N}$. Then, the B -partial potential of the subgame among the players in A is defined as

$$\Phi_B^A(S) = \Phi^A(S) - \Phi^{A \setminus B}(S).$$

When $A = \mathcal{N}$, we remove the superscript from partial potentials, i.e., $\Phi_B(S) = \Phi_B^{\mathcal{N}}(S)$. Also, when B is a singleton containing player u , we simplify the notation of the partial potential to $\Phi_u^A(S)$. Furthermore, observe that $\Phi_A^A(S) = \Phi^A(S)$.

The next four claims present basic properties of partial potentials.

CLAIM 3.7. Let S be a state of a Ψ -game and let $B \subseteq A \subseteq \mathcal{N}$. Then, $\Phi_B^A(S) \leq \Phi_B(S)$.

CLAIM 3.8. Let $A \subseteq \mathcal{N}$ be a set of players and let S and S' be states such that each player in A uses the same strategy in S and S' . Then, for every set of players $B \subseteq A$, $\Phi_B^A(S) = \Phi_B^A(S')$.

CLAIM 3.9. Let S be a state of a Ψ -game and let u be a player. Then, $\Phi_u(S) = \hat{c}_u(S)$.

CLAIM 3.10. Let u be a player and $A \subseteq \mathcal{N}$ a set of players that contains u . Then, for any two states S and S' that differ only in the strategy of player u , it holds that $\Phi_A(S) - \Phi_A(S') = \hat{c}_u(S) - \hat{c}_u(S')$.

In particular, Claim 3.10 implies that the A -partial potential can be thought of as a potential function defined over all states in which each player in $\mathcal{N} \setminus A$ uses the same strategy.

We proceed with the following interesting property that shows that the potential function of Ψ -games is cost-revealing. It also implies (for $A = \mathcal{N}$) that the potential of a state lower-bounds the total cost of all players.

LEMMA 3.11. *For every state S of a Ψ -game and any set of players $A \subseteq N$, it holds that $\Phi_A(S) \leq \sum_{u \in A} \hat{c}_u(S)$.*

3.3. The stretch of the potential function

An important quantity for our purposes is the *stretch* of the potential function of Ψ -games; a general definition that applies to every potential game follows.

Definition 3.12. Consider a potential game with a positive potential function Φ and let S^* be the state of minimum potential. The ρ -stretch of the potential function of the game is the maximum over all ρ -approximate pure Nash equilibria S of the ratio $\Phi(S)/\Phi(S^*)$.

The next two statements provide bounds on the ρ -stretch of the potential function of Ψ -games of degree 1 (i.e., linear weighted congestion games) and $d \geq 2$, respectively.

LEMMA 3.13. *For every $\rho \in [1, 11/10]$, the ρ -stretch of the potential function of a linear weighted congestion game is at most $\frac{3+\sqrt{5}}{2} + 6(\rho - 1)$.*

PROOF. Let S^* be the state of minimum potential and S be a ρ -approximate equilibrium. For each player u , we denote by s_u and s_u^* the strategies she plays at states S and S^* , respectively. Using the ρ -approximate equilibrium condition $c_u(S) \leq \rho \cdot c_u(S_{-u}, s_u^*)$, the definition of the cost of player u , and the definition of function Ψ_1 , we obtain

$$\begin{aligned} \sum_u c_u(S) &\leq \rho w_u \sum_{e \in s_u^*} (a_{e,1} \Psi_1(N_e(S_{-u}, s_u^*)) + a_{e,0}) \\ &\leq \rho w_u \sum_{e \in s_u^*} (a_{e,1} \Psi_1(N_e(S) \cup \{w_u\}) + a_{e,0}) \\ &= \rho w_u \sum_{e \in s_u^*} (a_{e,1} \Psi_1(N_e(S)) + a_{e,1} w_u + a_{e,0}). \end{aligned}$$

By summing over all players, by exchanging sums, and using the definition of $N_e(S^*)$, we obtain

$$\begin{aligned} \sum_u c_u(S) &\leq \rho \sum_u w_u \sum_{e \in s_u^*} (a_{e,1} \Psi_1(N_e(S)) + a_{e,1} w_u + a_{e,0}) \\ &= \rho \sum_e \left(a_{e,1} \Psi_1(N_e(S)) \sum_{u: e \in s_u^*} w_u + a_{e,1} \sum_{u: e \in s_u^*} w_u^2 + a_{e,0} \sum_{u: e \in s_u^*} w_u \right) \\ &= \rho \sum_e \left(a_{e,1} \Psi_1(N_e(S)) \Psi_1(N_e(S^*)) + a_{e,1} \sum_{u: e \in s_u^*} w_u^2 + a_{e,0} \Psi_1(N_e(S^*)) \right). \end{aligned}$$

We now apply the inequality $xy \leq \frac{\sqrt{5}-1}{2(3-\sqrt{5})}y^2 + \frac{\sqrt{5}-2}{3-\sqrt{5}}x^2$ that holds for any pair of non-negative x and y on the rightmost part of the above derivation to obtain

$$\begin{aligned} &\sum_u c_u(S) \\ &\leq \rho \sum_e \left(\frac{\sqrt{5}-1}{2(3-\sqrt{5})} a_{e,1} \Psi_1(N_e(S^*))^2 + \frac{\sqrt{5}-2}{3-\sqrt{5}} a_{e,1} \Psi_1(N_e(S))^2 \right) \end{aligned}$$

$$\begin{aligned}
& + a_{e,1} \sum_{u:e \in s_u^*} w_u^2 + a_{e,0} \Psi_1(N_e(S^*)) \Big) \\
= & \rho \sum_e \left(\frac{5 - \sqrt{5}}{4(3 - \sqrt{5})} a_{e,1} \left(\Psi_1(N_e(S^*))^2 + \sum_{u:e \in s_u^*} w_u^2 \right) + a_{e,0} \Psi_1(N_e(S^*)) \right) \\
& - \rho \sum_e \frac{7 - 3\sqrt{5}}{4(3 - \sqrt{5})} a_{e,1} \left(\Psi_1(N_e(S^*))^2 - \sum_{u:e \in s_u^*} w_u^2 \right) + \frac{\sqrt{5} - 2}{3 - \sqrt{5}} \rho \sum_e a_{e,1} \Psi_1(N_e(S))^2.
\end{aligned}$$

Now, observe that $\Psi_1(N_e(S^*))^2 \geq \sum_{u:e \in s_u^*} w_u^2$ for every resource e . Furthermore, $\Psi_1(N_e(S^*))^2 + \sum_{u:e \in s_u^*} w_u^2 = \Psi_2(N_e(S^*))$. Hence, we have

$$\begin{aligned}
& \sum_u c_u(S) \\
\leq & \rho \sum_e \left(\frac{5 - \sqrt{5}}{4(3 - \sqrt{5})} a_{e,1} \Psi_2(N_e(S^*)) + a_{e,0} \Psi_1(N_e(S^*)) \right) + \frac{\sqrt{5} - 2}{3 - \sqrt{5}} \rho \sum_e a_{e,1} \Psi_1(N_e(S))^2 \\
\leq & \frac{5 - \sqrt{5}}{2(3 - \sqrt{5})} \rho \sum_e \left(\frac{a_{e,1}}{2} \Psi_2(N_e(S^*)) + a_{e,0} \Psi_1(N_e(S^*)) \right) + \frac{\sqrt{5} - 2}{3 - \sqrt{5}} \rho \sum_e a_{e,1} \Psi_1(N_e(S))^2 \\
= & \frac{5 - \sqrt{5}}{2(3 - \sqrt{5})} \rho \Phi(S^*) + \frac{\sqrt{5} - 2}{3 - \sqrt{5}} \rho \sum_e a_{e,1} \Psi_1(N_e(S))^2. \tag{1}
\end{aligned}$$

We now use the definition of $\Phi(S)$, the fact that for every player u and resource $e \in s_u$, it holds that $w_u \leq \Psi_1(N_e(S))$, and the definition of the cost of player u . We have

$$\begin{aligned}
\Phi(S) &= \sum_e \left(\frac{a_{e,1}}{2} \Psi_2(N_e(S)) + a_{e,0} \Psi_1(N_e(S)) \right) \\
&= \sum_e \left(\frac{a_{e,1}}{2} \sum_{u:e \in s_u} (w_u \Psi_1(N_e(S)) + w_u^2) + a_{e,0} \sum_{u:e \in s_u} w_u \right) \\
&\leq \sum_e \left(\frac{a_{e,1}}{2} \sum_{u:e \in s_u} \left((6 - 2\sqrt{5}) w_u \Psi_1(N_e(S)) + (2\sqrt{5} - 4) w_u^2 \right) + a_{e,0} \sum_{u:e \in s_u} w_u \right) \\
&= (3 - \sqrt{5}) \sum_u w_u \sum_{e \in s_u} (a_{e,1} \Psi_1(N_e(S)) + a_{e,0}) + (\sqrt{5} - 2) \sum_e a_{e,1} \sum_{u:e \in s_u} w_u^2 \\
&\quad + (\sqrt{5} - 2) \sum_e a_{e,0} \sum_{u:e \in s_u} w_u \\
&= (3 - \sqrt{5}) \sum_u c_u(S) + (\sqrt{5} - 2) \sum_e a_{e,1} \sum_{u:e \in s_u} w_u^2 + (\sqrt{5} - 2) \sum_e a_{e,0} \sum_{u:e \in s_u} w_u.
\end{aligned}$$

By applying inequality (1) to the rightmost part of this derivation, we obtain

$$\begin{aligned}
\Phi(S) &\leq \frac{5 - \sqrt{5}}{2} \rho \Phi(S^*) + (\sqrt{5} - 2) \rho \sum_e a_{e,1} \Psi_1(N_e(S))^2 + (\sqrt{5} - 2) \sum_e a_{e,1} \sum_{u:e \in s_u} w_u^2 \\
&\quad + (\sqrt{5} - 2) \sum_e a_{e,0} \Psi_1(N_e(S))
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{5 - \sqrt{5}}{2} \rho \Phi(S^*) \\
&\quad + (2\sqrt{5} - 4) \rho \sum_e \left(\frac{a_{e,1}}{2} \left(\Psi_1(N_e(S))^2 + \sum_{u: e \in s_u} w_u^2 \right) + a_{e,0} \Psi_1(N_e(S)) \right) \\
&= \frac{5 - \sqrt{5}}{2} \rho \Phi(S^*) + (2\sqrt{5} - 4) \rho \sum_e \left(\frac{a_{e,1}}{2} \Psi_2(N_e(S)) + a_{e,0} \Psi_1(N_e(S)) \right) \\
&= \frac{5 - \sqrt{5}}{2} \rho \Phi(S^*) + (2\sqrt{5} - 4) \rho \Phi(S).
\end{aligned}$$

The last inequality implies that $\Phi(S)$ is not larger than $\frac{(5-\sqrt{5})\rho}{2(1-(2\sqrt{5}-4)\rho)}\Phi(S^*)$ which can be easily proved to be at most $\left(\frac{3+\sqrt{5}}{2} + 6(\rho-1)\right)\Phi(S^*)$ when $\rho \in [1, 11/10]$. \square

LEMMA 3.14. *The ρ -stretch of the potential function of a Ψ -game of degree $d \geq 2$ is at most $\rho(\rho+1)^d(d+1)^{d+1}$.*

In the rest of the paper, we denote by $\theta_d(\rho)$ the upper bounds on the ρ -stretch given by Lemmas 3.13 and 3.14, namely $\theta_1(\rho) = \frac{3+\sqrt{5}}{2} + 6(\rho-1)$ and $\theta_d(\rho) = \rho(\rho+1)^d(d+1)^{d+1}$. The next lemma extends these bounds to partial potentials.

LEMMA 3.15. *Consider a Ψ -game of degree d and a state S which is a ρ -approximate pure Nash equilibrium for a set of players $R \subseteq \mathcal{N}$. Then, $\Phi_R(S) \leq \theta_d(\rho)\Phi_R(S^*)$ for any state S^* such that each player in $\mathcal{N} \setminus R$ uses the same strategy in S and S^* .*

4. THE ALGORITHM

In this section we describe our algorithm (Algorithm 1; see the table below). The algorithm takes as input a Ψ -game \mathcal{G} of degree d with n players, an arbitrary initial state S of the game, and a small positive parameter γ . It produces as output a state of \mathcal{G} . The algorithm starts by initializing its parameters, namely \hat{c}_{\max} , \hat{c}_{\min} , m , g , q , and p (lines 1-6). It first computes the minimum possible cost \hat{c}_{\min} among all players and the maximum cost \hat{c}_{\max} experienced by players in the initial state S . Then, it sets the parameter m equal to $\log(\hat{c}_{\max}/\hat{c}_{\min})$; in this way, m is polynomial in the number of bits in the representation of the game (i.e., polynomial in the number of bits necessary to store the parameters $a_{e,k}$ and the weights of the players). Then, the parameter q is set close to 1 (namely, $q = 1 + \gamma$) and parameter p is set close to $\theta_d(q)$ (namely, $p = \left(\frac{1}{\theta_d(q)} - 2\gamma\right)^{-1}$).

Recall that $\theta_d(q)$ is the bound on the q -stretch of the potential function of Ψ -games of degree d in the statements of Lemmas 3.13 (for $d = 1$) and 3.14 (for $d \geq 2$).

Then, the algorithm runs a sequence of phases; within each phase, it coordinates best-response moves of the players. This process starts (line 7) by computing a decreasing sequence of boundaries $b_0, b_1, b_2, \dots, b_m$ that will be used to define the sets of players that are considered to move within each phase. Then, it executes phase 0 (lines 8-10). During this phase, as long as there are players of cost at least b_1 that have a q -move, they play a best-response strategy. Hence, after the end of the phase, all players with cost higher than b_1 are in a q -approximate equilibrium. Then, the algorithm uses set F to keep the players whose strategies have been irrevocably decided; F is initialized to \emptyset in line 11. Phases 1 to $m-1$ (lines 12-17) constitute the heart of our algorithm. During each such phase i , the algorithm repeatedly checks whether, in the current state, there is a player that either has cost higher than b_i that has a p -move or her cost is in $[b_{i+1}, b_i)$ and has a q -move. While such a player is found, she deviates

input : A Ψ -game \mathcal{G} of degree d with a set \mathcal{N} of n players, an arbitrary initial state S , and $\gamma > 0$ with $\gamma \in (0, 1/10]$ if $d = 1$ and $\gamma \in (0, \frac{1}{8\theta_d(2)}]$, otherwise

output: A state of \mathcal{G}

```

1  $\hat{c}_{\min} \leftarrow \min_{u \in \mathcal{N}} \hat{c}_u(\mathbf{0}_{-u}, \mathcal{BR}_u(\mathbf{0}));$ 
2  $\hat{c}_{\max} \leftarrow \max_{u \in \mathcal{N}} \hat{c}_u(S);$ 
3  $m \leftarrow \log(\hat{c}_{\max}/\hat{c}_{\min});$ 
4  $g \leftarrow 2(1 + m(1 + \gamma^{-1}))^d d^d n \gamma^{-3};$ 
5  $q \leftarrow 1 + \gamma;$ 
6  $p \leftarrow \left(\frac{1}{\theta_d(q)} - 2\gamma\right)^{-1};$ 
7 for  $i \leftarrow 0$  to  $m$  do  $b_i \leftarrow \hat{c}_{\max} g^{-i};$ 
8 while there exists a player  $u \in \mathcal{N}$  such that  $\hat{c}_u(S) \in [b_i, +\infty)$  and
    $\hat{c}_u(S_{-u}, \mathcal{BR}_u(S)) < \hat{c}_u(S)/q$  do
9    $S \leftarrow (S_{-u}, \mathcal{BR}_u(S));$ 
10 end
11  $F \leftarrow \emptyset;$ 
12 for phase  $i \leftarrow 1$  to  $m - 1$  do
13   while there exists a player  $u \in \mathcal{N} \setminus F$  such that either  $\hat{c}_u(S) \in [b_i, +\infty)$  and
      $\hat{c}_u(S_{-u}, \mathcal{BR}_u(S)) < \hat{c}_u(S)/p$  or  $\hat{c}_u(S) \in [b_{i+1}, b_i)$  and  $\hat{c}_u(S_{-u}, \mathcal{BR}_u(S)) < \hat{c}_u(S)/q$ 
     do
14      $S \leftarrow (S_{-u}, \mathcal{BR}_u(S));$ 
15     end
16    $F \leftarrow F \cup \{u \in \mathcal{N} \setminus F : \hat{c}_u(S) \in [b_i, +\infty)\};$ 
17 end

```

ALGORITHM 1: Computing approximate equilibria in Ψ -games.

to her best-response strategy. The phase terminates when no such player exists and the algorithm irrevocably decides the strategy of the players that have cost at least b_i . These players are included in set F ; at this point, they are guaranteed to be at a p -approximate equilibrium. Subsequent moves by other players may either increase their cost or decrease the cost they could experience by deviating to another strategy. As we will show, these changes are not significant and each player will still be at an almost p -approximate equilibrium at the end of all phases. The fact that plays a crucial role towards proving such a claim is that, at the end of each phase i , any player with cost in $[b_{i+1}, b_i)$ is guaranteed to be in a q -approximate equilibrium. Note that $b_m \leq \hat{c}_{\min}$ and, eventually, all players will be included in set F .

We remark that the sequence of the phases is similar to the one in our algorithm for unweighted congestion games with polynomial latency functions of constant degree d in [Caragiannis et al. 2011]. However, there is an important difference. In that context, each player is considered to move during only two consecutive phases; these phases are defined statically based only on the characteristics of the particular player. The main reason that allows this is that the cost that a player may experience by following a specific strategy may change by at most a polynomial factor (namely, at most n^d) during the execution of the algorithm. This is not the case in the context of Ψ -games since the fact that the cost of a player depends on the weights of the other players does not satisfy this polynomial relation. So, in the current algorithm, the players that are considered to move within each phase are decided *dynamically* based on the cost they experience during a phase. In this way, a player may (be considered to) move in many different phases.

Below, we will prove the following statement.

THEOREM 4.1. *Algorithm 1 computes a $\hat{\rho}_d$ -approximate equilibrium for every Ψ -game of constant degree d , where $\hat{\rho}_1 = \frac{3+\sqrt{5}}{2} + O(\gamma)$ and $\hat{\rho}_d \in d^{d+o(d)}$. The running time is polynomial in γ^{-1} and in the number of bits in the representation of the game.*

Combined with Lemma 3.5, Theorem 4.1 immediately yields the following result for weighted congestion games.

THEOREM 4.2. *When Algorithm 1 is applied to the Ψ -game corresponding to a weighted congestion game with polynomial latency functions of constant degree d , it computes a state which is a ρ_d -approximate equilibrium for the latter, where $\rho_1 = \frac{3+\sqrt{5}}{2} + O(\gamma)$ and $\rho_d \in d^{2d+o(d)}$ for $d \geq 2$.*

The rest of this section is devoted to proving Theorem 4.1. Throughout the section we consider the application of the algorithm on a Ψ -game of degree d and denote by S^i the state computed by the algorithm after the execution of phase i for $i = 0, 1, \dots, m-1$. Also, we use R_i to denote the set of players that make at least one move during phase i . Our arguments are split in three parts. First, we present a key property maintained by our algorithm stating that the R_i -partial potential is small when the phase $i \geq 1$ starts. Then, we use this fact together with the parameters of the algorithm to prove that the running time is polynomial. The proof of the approximation guarantee follows. Recall that the players whose strategies are irrevocably decided during phase $j \geq 1$ are at a p -approximate equilibrium at the end of the phase. The purpose of the third part of the proof is to show that for each such player, neither her cost increases significantly nor the cost she would experience by deviating to another strategy decreases significantly after phase j . Hence, the approximation guarantee in the final state computed by the algorithm is slightly higher than p .

We remark that the analysis follows the same general steps as in our recent paper on unweighted congestion games [Caragiannis et al. 2011]. However, due to the definition of Ψ -games and the dependency of players' cost on the weights, different and significantly more involved arguments are required, especially in the first and third step.

4.1. Proving the key property

The key property maintained by our algorithm is the following.

LEMMA 4.3. *For every phase $i \geq 1$, it holds that $\Phi_{R_i}(S^{i-1}) \leq \gamma^{-1}nb_i$.*

PROOF. In order to prove the key property maintained by our algorithm, we will need the following lemma which relates the R_i -partial potential to the cost the players in R_i experience when they make their last move within phase i .

LEMMA 4.4. *Let $\hat{c}(u)$ denote the cost of player $u \in R_i$ just after making her last move within phase $i \geq 1$. Then,*

$$\Phi_{R_i}(S^i) \leq \sum_{u \in R_i} \hat{c}(u).$$

PROOF. Rename the players in R_i as $u_1, u_2, \dots, u_{|R_i|}$ so that u_j is the j -th player that performed her last move within phase $i \geq 1$. Also, denote by $S^{i,j}$ the state in which player u_j performed her last move. Let $R_i^{|R_i|} = \emptyset$ and $R_i^j = \{u_{j+1}, u_{j+2}, \dots, u_{|R_i|}\}$ for $j = 0, 1, 2, \dots, |R_i| - 1$. Then,

$$\Phi_{R_i}(S^i) = \Phi(S^i) - \Phi^{\mathcal{N} \setminus R_i}(S^i) = \sum_{j=1}^{|R_i|} \left(\Phi^{\mathcal{N} \setminus R_i^j}(S^i) - \Phi^{\mathcal{N} \setminus R_i^{j-1}}(S^i) \right) = \sum_{j=1}^{|R_i|} \Phi_{u_j}^{\mathcal{N} \setminus R_i^j}(S^i)$$

$$= \sum_{j=1}^{|R_i|} \Phi_{u_j}^{\mathcal{N} \setminus R_i^j}(S^{i,j}) \leq \sum_{j=1}^{|R_i|} \Phi_{u_j}(S^{i,j}) = \sum_{u \in R_i} \hat{c}(u).$$

The first three equalities follow by the definition of the partial potential functions and the definition of sets R_i^j . The fourth equality follows by Claim 3.8 since players in $\mathcal{N} \setminus R_i^j$ do not move after state $S^{i,j}$ and until the end of the phase. The inequality follows by Claim 3.7 and the last equality follows by Claim 3.9 and the definition of $\hat{c}(u)$. \square

We now proceed to the proof of Lemma 4.3. For the sake of contradiction, we assume that $\Phi_{R_i}(S^{i-1}) > \gamma^{-1}nb_i$ and we denote by P_i and Q_i the set of players in R_i whose last move was a p -move and q -move, respectively. Since each player in P_i decreases her cost by at least $(p-1)\hat{c}(u)$ during her last move within phase i (see Claim 3.10), we have

$$\Phi_{R_i}(S^i) \leq \Phi_{R_i}(S^{i-1}) - (p-1) \sum_{u \in P_i} \hat{c}(u).$$

By Lemma 4.4 and the fact that each player in Q_i experiences a cost of at most b_i when she makes her last move within phase i , we have

$$\sum_{u \in P_i} \hat{c}(u) \geq \Phi_{R_i}(S^i) - \sum_{u \in Q_i} \hat{c}(u) \geq \Phi_{R_i}(S^i) - nb_i.$$

Using the last two inequalities and our assumption, we obtain that

$$\begin{aligned} \Phi_{R_i}(S^i) &\leq \Phi_{R_i}(S^{i-1}) - (p-1)\Phi_{R_i}(S^i) + (p-1)nb_i \\ &< (1 + (p-1)\gamma)\Phi_{R_i}(S^{i-1}) - (p-1)\Phi_{R_i}(S^i) \end{aligned}$$

which implies that

$$\Phi_{R_i}(S^i) < \left(\frac{1}{p} + \gamma\right) \Phi_{R_i}(S^{i-1}).$$

Now, consider state S^{i-1} and let X_i and Y_i be the sets of players in R_i with cost at least b_i and smaller than b_i , respectively. Notice that, by the definition of phase $i-1$, S^{i-1} is a q -approximate equilibrium for the players in X_i . We construct a new Ψ -game of degree d among the players in \mathcal{N} as follows. The new game has all resources of the original game; the parameters $a_{e,k}$ for these resources are the same as in the original game. In addition, the new game has a new resource e_u for each player $u \in Y_i$; the parameters for this resource are $a_{e_u,0} = b_i/w_u$ and $a_{e_u,k} = 0$ for $k = 1, \dots, d$. Each player in $\mathcal{N} \setminus Y_i$ has the same set of strategies in the two games. The strategy set of player $u \in Y_i$ consists of the strategy s_u she uses in S^{i-1} as well as strategy $s'_u \cup \{e_u\}$ for each strategy $s'_u \neq s_u$ she has in the original game.

Let \bar{S}^{i-1} be the state of the new game in which all players play their strategies in S^{i-1} . Clearly, state \bar{S}^{i-1} is a q -approximate equilibrium for the players in X_i . Also, at state \bar{S}^{i-1} , each player $u \in Y_i$ experiences a cost equal to the cost she experiences at state S^{i-1} of the original game, i.e., smaller than b_i . In the new game, any deviation of u would include resource e_u and would increase the cost of player u to at least $w_u a_{e_u,0} = b_i$. Hence, \bar{S}^{i-1} is a q -approximate equilibrium for the players of Y_i as well. We use $\bar{\Phi}$ to denote the potential of the new game. Since the players use the same strategies in states S^{i-1} and \bar{S}^{i-1} and the parameters $a_{e,k}$ of the original resources are the same in both games, we have $\bar{\Phi}_{R_i}(\bar{S}^{i-1}) = \Phi_{R_i}(S^{i-1})$.

Now, let \bar{S}^i be the state in which each player in $\mathcal{N} \setminus Y_i$ uses her strategy in S^i and the strategies for the players in Y_i are defined as follows. Let u be a player of Y_i and s'_u be the strategy she uses at state S^i of the original game. Her strategy in state \bar{S}^i of

the new game is $s'_u \cup \{e_u\}$ if $s'_u \neq s_u$ and s_u otherwise. Observe that, by the definition of the partial potential, we have that the partial potential $\bar{\Phi}_{R_i}(\bar{S}^i)$ of the new game at state \bar{S}^i is by at most $\sum_{u \in Y_i} a_{e_u,0} \Psi_1(N_{e_u}(\bar{S}^i)) \leq nb_i$ higher than the partial potential of the original game at state S^i (due to the contribution of the additional resources to the potential value). Hence,

$$\begin{aligned} \bar{\Phi}_{R_i}(\bar{S}^i) &\leq \Phi_{R_i}(S^i) + nb_i < \left(\frac{1}{p} + 2\gamma\right) \Phi_{R_i}(S^{i-1}) \\ &= \left(\frac{1}{p} + 2\gamma\right) \bar{\Phi}_{R_i}(\bar{S}^{i-1}) = \frac{1}{\theta_d(q)} \bar{\Phi}_{R_i}(\bar{S}^{i-1}). \end{aligned}$$

So, we have identified a state \bar{S}^{i-1} of the new game which is a q -approximate equilibrium for the players in R_i and another state \bar{S}^i such that the players in $\mathcal{N} \setminus R_i$ use the same strategies in \bar{S}^{i-1} and \bar{S}^i and $\bar{\Phi}_{R_i}(\bar{S}^{i-1}) > \theta_d(q) \bar{\Phi}_{R_i}(\bar{S}^i)$. This contradicts Lemma 3.15 and, subsequently, it also contradicts our assumption $\Phi_{R_i}(S^{i-1}) > \gamma^{-1}nb_i$. The lemma follows. \square

4.2. Bounding the running time

We will now use Lemma 4.3 and the properties of Ψ -games to prove that the algorithm terminates quickly.

LEMMA 4.5. *The algorithm terminates after a number of steps that is polynomial in the number of bits in the representation of the game and γ^{-1} .*

PROOF. Clearly, if the number of strategies is polynomial in the number of resources, computing a best-response strategy for a player u can be trivially performed in polynomial time (by the definition of \hat{c}_u). This is also the case for weighted congestion games in networks (where the number of strategies of a player can be exponential) using a shortest path computation. So, it remains to bound the total number of player moves.

At the initial state, the total cost of the players and, consequently (by Lemma 3.11), its potential is at most $n\hat{c}_{\max}$. Each of the players that move during phase 0 decreases her cost and, consequently (by Theorem 3.3), the potential by at least $(q-1)b_1 = \gamma g^{-1}\hat{c}_{\max}$. Hence, the total number of moves in phase 0 is at most $n\gamma^{-1}g$. For $i \geq 1$, we have $\Phi_{R_i}(S^i) \leq nb_i\gamma^{-1}$ (by Lemma 4.3). Each of the players in R_i that move during phase i decreases her cost and, consequently (by Claim 3.10), the R_i -partial potential by at least $(q-1)b_{i+1} = b_i g^{-1}\gamma$. Hence, phase i completes after at most $ng\gamma^{-2}$ moves. In total, we have at most $mng\gamma^{-2}$ moves. The theorem follows by observing that g depends polynomially on m , n , and γ^{-1} . \square

4.3. Proving the approximation guarantee

It remains to prove that our algorithm computes approximate equilibria. Our proofs will exploit Lemma 4.3 as well as the following lemma which relates the cost of a player in a state to the partial potential of two different subgames.

LEMMA 4.6. *Consider a Ψ -game of degree d , a player u and a set of players $R \subseteq \mathcal{N} \setminus \{u\}$. Then, for every state S and every $\epsilon > 0$, it holds that*

$$\hat{c}_u(S) \leq (1 + \epsilon)\Phi_u^{\mathcal{N} \setminus R}(S) + \xi_\epsilon \Phi_R^{\mathcal{N} \setminus \{u\}}(S),$$

where $\xi_\epsilon = (1 + 1/\epsilon)^d d^d - 1$.

Using Lemmas 4.3 and 4.6, we will show that neither the cost of a player increases significantly after the phase at the end of which her strategy was irrevocably decided

(in Lemma 4.7), nor the cost she would experience by deviating to another strategy decreases significantly (in Lemma 4.8).

LEMMA 4.7. *Let u be a player whose strategy was irrevocably decided at phase j . Then, $\hat{c}_u(S^{m-1}) \leq (1 + 2\gamma)\hat{c}_u(S^j)$.*

PROOF. For every $i > j$ and $\epsilon > 0$, we apply Lemma 4.6 for strategy S^i , player u , and the set of players R_i that move during phase i to obtain

$$\begin{aligned}\hat{c}_u(S^i) &\leq (1 + \epsilon)\Phi_u^{\mathcal{N} \setminus R_i}(S^i) + \xi_\epsilon \Phi_{R_i}^{\mathcal{N} \setminus \{u\}}(S^i) \\ &= (1 + \epsilon)\Phi_u^{\mathcal{N} \setminus R_i}(S^{i-1}) + \xi_\epsilon \Phi_{R_i}^{\mathcal{N} \setminus \{u\}}(S^i) \\ &\leq (1 + \epsilon)\Phi_u(S^{i-1}) + \xi_\epsilon \Phi_{R_i}(S^i) \\ &\leq (1 + \epsilon)\hat{c}_u(S^{i-1}) + \xi_\epsilon \Phi_{R_i}(S^{i-1}).\end{aligned}$$

The equality holds by Claim 3.8 since the players in $\mathcal{N} \setminus R_i$ do not move during phase i . The second inequality follows by Claim 3.7. The last one follows by Claim 3.9 and since the R_i -partial potential decreases during phase i .

We now set $\epsilon = (1 + \gamma)^{1/m} - 1$. This implies that $(1 + \epsilon)^m = 1 + \gamma$ and $\epsilon \geq \frac{\gamma}{m}(1 + \gamma)^{1/m-1} \geq (m(1 + \gamma^{-1}))^{-1}$ and, by the definition of the parameters g and γ , $\xi_\epsilon = (1 + m(1 + \gamma^{-1}))^d d^d - 1 \leq \frac{g\gamma^3}{2n} \leq \frac{g\gamma}{2(1 + \gamma^{-1})n}$. Using the above inequality together with these observations, we obtain

$$\begin{aligned}\hat{c}_u(S^{m-1}) &\leq (1 + \epsilon)^{m-1-j}\hat{c}_u(S^j) + \xi_\epsilon \sum_{i=j+1}^{m-1} (1 + \epsilon)^{m-1-i}\Phi_{R_i}(S^{i-1}) \\ &\leq (1 + \epsilon)^m \hat{c}_u(S^j) + (1 + \epsilon)^m \xi_\epsilon \sum_{i=j+1}^{m-1} \Phi_{R_i}(S^{i-1}) \\ &\leq (1 + \gamma)\hat{c}_u(S^j) + (1 + \gamma)\xi_\epsilon \sum_{i=j+1}^{m-1} nb_i \gamma^{-1} \\ &= (1 + \gamma)\hat{c}_u(S^j) + (1 + \gamma^{-1})\xi_\epsilon nb_j \sum_{i=1}^{m-1-j} g^{-i} \\ &\leq (1 + \gamma)\hat{c}_u(S^j) + 2(1 + \gamma^{-1})\xi_\epsilon nb_j g^{-1} \\ &\leq (1 + \gamma)\hat{c}_u(S^j) + \gamma b_j \\ &\leq (1 + 2\gamma)\hat{c}_u(S^j).\end{aligned}$$

The second inequality is obvious, the third one follows by Lemma 4.3 and by the relation between ϵ and γ , the equality follows by the definition of b_i , the fourth inequality follows since $g \geq 2$ which implies that $\sum_{i \geq 1} g^{-i} \leq 2g^{-1}$, the fifth one follows by our observation about ξ_ϵ above, and the last one follows since, by the definition of the algorithm, the fact that the strategy of player u is irrevocably decided at phase j implies that $\hat{c}_u(S^j) \geq b_j$. \square

LEMMA 4.8. *Let u be a player whose strategy was irrevocably decided at phase j and let s'_u be any of her strategies. Then, $\hat{c}_u(S_{-u}^{m-1}, s'_u) \geq (1 - 2\gamma)\hat{c}_u(S_{-u}^j, s'_u)$.*

We are now ready to use the last two lemmas in order to prove the approximation guarantee of the algorithm. This will complete the proof of Theorem 4.1.

LEMMA 4.9. *Given a Ψ -game of degree d , the algorithm computes a $\hat{\rho}_d$ -approximate equilibrium with $\hat{\rho}_1 \leq \frac{3+\sqrt{5}}{2} + O(\gamma)$ and $\hat{\rho}_d \leq d^{d+o(d)}$.*

PROOF. Consider the application of the algorithm to a Ψ -game and let u be any player whose strategy is irrevocably decided at the end of phase j of the algorithm. Also, let s'_u be any other strategy of this player. By Lemmas 4.7 and 4.8 and since, by the definition of the algorithm, player u has no incentive to make a p -move at state S^j , we have

$$\frac{\hat{c}_u(S^{m-1})}{\hat{c}_u(S_{-u}^{m-1}, s'_u)} \leq \frac{(1+2\gamma)}{(1-2\gamma)} \cdot \frac{\hat{c}_u(S^j)}{\hat{c}_u(S_{-u}^j, s'_u)} \leq \frac{1+2\gamma}{1-2\gamma} \left(\frac{1}{\theta_d(1+\gamma)} - 2\gamma \right)^{-1}.$$

Hence, the right-hand side of the above inequality upper-bounds the approximation guarantee of the algorithm. For $d = 1$, the parameter γ takes values in $(0, 1/10]$. Since $\gamma \in (0, 1/10]$ and $\theta_1(1+\gamma) = \frac{3+\sqrt{5}}{2} + 6\gamma$ (see Lemma 3.13), by making simple calculations, we obtain that the algorithm computes a $\hat{\rho}_1$ -approximate equilibrium with

$$\hat{\rho}_1 \leq \frac{3+\sqrt{5}}{2} + 110\gamma.$$

For larger values of d , the algorithm uses $\gamma \in (0, \frac{1}{8\theta_d(2)}]$. Since $\theta_d(1+\gamma)$ is non-decreasing in γ , we have that $\left(\frac{1}{\theta_d(1+\gamma)} - 2\gamma \right)^{-1} \leq \frac{4}{3}\theta_d(2)$. Also, we have that $\gamma < 1/34$ and hence $\frac{1+2\gamma}{1-\gamma} \leq \frac{9}{8}$. By using the value for $\theta_d(2)$ from Lemma 3.14, we have that the algorithm computes a $\hat{\rho}_d$ -approximate equilibrium with $\hat{\rho}_d \leq 3^{d+1}(d+1)^{d+1} \in d^{d+o(d)}$. \square

5. CONCLUSIONS AND OPEN PROBLEMS

Our work reveals interesting open problems. The obvious one is whether approximate equilibria with a better approximation guarantee can be computed in polynomial time. We believe that our techniques have reached their limits for linear weighted congestion games. However, in the case of superlinear latency functions, approximations of weighted congestion games by potential games different than Ψ -games might yield improved (existential or algorithmic) approximation guarantees. On the conceptual level, it is interesting to further explore applications of approximations of non-potential games by potential ones like the one we have exploited in the current paper.

REFERENCES

- ACKERMANN, H., RÖGLIN, H., AND VÖCKING, B. 2008. On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55, 6.
- ACKERMANN, H., RÖGLIN, H., AND VÖCKING, B. 2009. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science* 410, 17, 1552–1563.
- ANSHELEVICH, E. AND CASKURLU, B. 2009. Exact and approximate equilibria for optimal group network formation. In *Proceedings of the 17th Annual Symposium on Algorithms (ESA)*, LNCS 5757, Springer, 239–250.
- AWERBUCH, B., AZAR, Y., EPSTEIN, A., MIRROKNI, V. S., AND SKOPALIK, A. 2008. Fast convergence to nearly optimal solutions in potential games. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, ACM, 264–273.
- BHALGAT, A., CHAKRABORTY, T., AND KHANNA, S. 2010. Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, ACM, 73–82.
- CANDOGAN, O., OZDAGLAR, A. E., AND PARRILO, P. A. 2011a. Learning in near-potential games. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 2428–2433.
- CANDOGAN, O., OZDAGLAR, A. E., AND PARRILO, P. A. 2011b. Dynamics in near-potential games. *arxiv:1107.4386*.

- CARAGIANNIS, I. 2009. Efficient coordination mechanisms for unrelated machine scheduling. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 815–824. Full version available as *arxiv:1107.1814*.
- CARAGIANNIS, I., FANELLI, A., GRAVIN, N., AND SKOPALIK, A. 2011. Efficient computation of approximate pure Nash equilibria in congestion games. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 532–541. Full version available as *arxiv:1104.4423*.
- CARDINAL, J. AND HOEFER, M. 2010. Non-cooperative facility location and covering games. *Theoretical Computer Science* 411, 16–18, 1855–1876.
- CHARIKAR, M., KARLOFF, H. J., MATHIEU, C., NAOR, J., AND SAKS, M. E. 2008. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, ACM, 70–76.
- CHEKURI, C., CHUZHUY, J., LEWIN-EYTAN, L., NAOR, J., AND ORDA, A. 2007. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications* 25, 6, 1193–1206.
- CHIEN, S. AND SINCLAIR, A. 2011. Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior* 71, 2, 315–327.
- CHRISTIN, N., GROSSKLAGS, J., AND CHUANG, J. 2004. Near rationality and competitive equilibria in networked systems. In *Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, ACM, 213–219.
- CHRISTODOULOU, G., MIRROKNI, V. S., AND SIDIROPOULOS, A. 2006. Convergence and approximation in potential games. In *Proceedings of the 23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 3884, Springer, 349–360.
- DASKALAKIS, C. AND PAPADIMITRIOU, C. H. 2007. Computing equilibria in anonymous games. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, 83–93.
- DUNKEL, J. AND SCHULZ, A. S. 2008. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research* 33, 4, 851–868.
- EVEN-DAR, E., KESSELMAN, A., AND MANSOUR, Y. 2007. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms* 3, 3.
- FOTAKIS, D., KONTOGIANNIS, S., AND SPIRAKIS, P. G. 2005. Selfish unsplittable flows. *Theoretical Computer Science* 340, 3, 514–538.
- FABRIKANT, A., PAPADIMITRIOU, C. H., AND TALWAR, K. 2004. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, 604–612.
- FANELLI, A. AND MOSCARDELLI, L. 2009. On best response dynamics in weighted congestion games with polynomial delays. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*, LNCS 5929, Springer, 55–66.
- GOEMANS, M. X., MIRROKNI, V. S., AND VETTA, A. 2005. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 142–154.
- HARKS, T. AND KLIMM, M. 2010. On the existence of pure Nash equilibria in weighted congestion games. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP)*, Part 1, LNCS 6198, Springer, 79–89.
- JOHNSON, D. S., PAPADIMITRIOU, C. H., AND YANNAKAKIS, M. 1988. How easy is local search? *Journal of Computer and System Sciences* 37, 79–100.
- KOLLIAS, K. AND ROUGHGARDEN, T. 2011. Restoring pure equilibria to weighted congestion games. In *Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP)*, Part II, LNCS 6756, Springer, 539–551.
- MIRROKNI, V. S. AND VETTA, A. 2004. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, LNCS 3122, Springer, 183–192.
- NGUYEN, T. AND TARDOS, E. 2009. Approximate pure Nash equilibria via Lovász local lemma. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*, LNCS 5929, Springer, 160–171.
- PANAGOPOULOU, P. N. AND SPIRAKIS, P. G. 2006. Algorithms for pure Nash equilibria in weighted congestion games. *ACM Journal of Experimental Algorithmics* 11.
- ROSENTHAL, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67.
- SKOPALIK, A. AND VÖCKING, B. 2008. Inapproximability of pure Nash equilibria. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, ACM, 355–364.