

Short Sequences of Improvement Moves Lead to Approximate Equilibria in Constraint Satisfaction Games*

Ioannis Caragiannis¹, Angelo Fanelli², and Nick Gravin³

¹ Computer Technology Institute “Diophantus” &
Department of Computer Engineering and Informatics,
University of Patras, 26504 Rion, Greece

² CNRS, Laboratoire CREM (UMR 6211),
Université de Caen Basse-Normandie, France

³ Microsoft Research,
One Memorial Drive, Cambridge, MA 02142, USA

Abstract. We present an algorithm that computes approximate pure Nash equilibria in a broad class of constraint satisfaction games that generalize the well-known cut and party affiliation games. Our results improve previous ones by Bhalgat et al. (EC 10) in terms of the obtained approximation guarantee. More importantly, our algorithm identifies a polynomially-long sequence of improvement moves from any initial state to an approximate equilibrium in these games. The existence of such short sequences is an interesting structural property which, to the best of our knowledge, was not known before. Our techniques adapt and extend our previous work for congestion games (FOCS 11) but the current analysis is considerably simpler.

1 Introduction

Constraint satisfaction games are generalizations of the well-known cut games and party affiliation games. In a constraint satisfaction game, there is a set of boolean variables and a set of weighted constraints; each constraint depends on some of these variables. Each player controls the value of a distinct variable and has two possible strategies: setting the value of the variable to either 0 (false) or 1 (true). The payoff (or utility) of a player is the total weight in satisfied constraints where her variable appears. Constraint satisfaction games are potential games. The total weight of satisfied constraints serves as an exact potential function in the sense that the difference in the potential between two states that differ in the strategy of a single player equals the change in the utility of that player. Hence, pure Nash equilibria (i.e., states in which no player has an incentive to unilaterally move in order to improve her utility) can be computed by solving

* The work was partially supported by the European Social Fund and Greek national funds through the research funding program Thales on “Algorithmic Game Theory” and by a Caratheodory research grant from the University of Patras.

the local search problem (see [14] for a theoretical treatment of local search) of computing a local maximum of the potential function. Unfortunately, this is a computationally-hard problem [19]. In this paper, we resort to the question of whether relaxed solution concepts — namely, approximate (pure Nash) equilibria — can be computed efficiently.

In particular, we consider constraint satisfaction games where each constraint depends on the value of at most k variables and has the property that its value can change from false to true by a unilateral change in any of its variables. In general, we refer to such games as P_k -FLIP games following the terminology of Bhalgat et al. [3]. Particular examples of this type of constraints include “parity” and “not-all-equal” constraints. An odd (respectively, even) parity constraint requires that the number of its true variables is odd (respectively, even). A not-all-equal constraint consists of literals (i.e., variables or their negations) and requires that at least two of its literals have different values. We refer to P_k -FLIP games consisting of parity constraints as PARITY- k -FLIP games; P_k -FLIP games with not-all-equal constraints with at least \bar{k} literals are called NAE- (\bar{k}, k) -FLIP games. Party affiliation games are PARITY-2-FLIP games and, in particular, cut games are PARITY-2-FLIP games with odd constraints or NAE- $(2, 2)$ -FLIP games whose constraints have no negative literals.

By adapting and extending our techniques in [4] for congestion games, we present a polynomial-time algorithm that computes approximate equilibria in P_k -FLIP games. The approximation guarantee is related to the stretch θ of the potential function of games in a given class, defined as the maximum over all games in the class of the maximum ratio between the potential values in two equilibria. As we show, P_k -FLIP games have a stretch of $k + 1$; hence, for general P_k -FLIP games, the approximation guarantee $\theta + \varepsilon$ of our algorithm improves a previous one of $2k - 1 + \varepsilon$ by Bhalgat et al. [3] for $k \geq 3$. By bounding the stretch of NAE- (\bar{k}, k) -FLIP and PARITY- k -FLIP games, we are able to show further improvements. For NAE- (\bar{k}, k) -FLIP games, the approximation guarantee becomes $3 + \varepsilon$ for $\bar{k} = 2$ and $2 + \varepsilon$ for $\bar{k} \geq 3$; these results improve a bound of $\frac{2\bar{k}}{\bar{k}-1} + \varepsilon$ from [3]. For PARITY- k -FLIP games with odd k , the approximation guarantee is $k + \varepsilon$. The running time of the algorithm is bounded by a polynomial of the number of players, k , and $1/\varepsilon$. Our analysis follows the same general structure of [4] but uses different technical arguments and is considerably simpler due to the simplicity in the definition of P_k -FLIP games.

More importantly, for every initial state of the game, our algorithm identifies a polynomially-long sequence of improvement moves of the players that lead to an approximate equilibrium. The existence of such short sequence suggests an interesting structural property of P_k -FLIP games which, to the best of our knowledge, was not known before. Actually, Bhalgat et al. [3] argue about the limitations of (uncoordinated) improvement move sequences by presenting a particular cut game in which any sequence of ρ -moves (i.e., moves that improve the utility of the moving player by a factor of at least ρ) from some states to any ρ -approximate equilibrium has exponential length for any $\rho \in [1, 21/20)$. This negative result complements nicely with the structural property we prove.

Our algorithm is simple. Players are classified into blocks so that the players within the same block have polynomially-related maximum utility (i.e., total weight of the constraints a player can affect). Then, a set of phases is executed. In each phase the players in two consecutive blocks are allowed to move. The players in the block of higher maximum utility are allowed to make p -moves and the players of the other block are allowed to make q -moves. Then, the strategies of the players that were allowed to perform p -moves within a phase are irrevocably decided at its end. Clearly, this defines a sequence of improvement moves by the players. We show that by setting the parameters q and p appropriately, the algorithm terminates in polynomial time and, furthermore, the players whose strategies are irrevocably decided at the end of a phase will not be affected significantly by later moves. In order to do so, we select a value for parameter p that is slightly higher than the stretch of the class of games to which the input game belongs and a value for parameter q that is very close to 1.

Related Work. Schäffer and Yannakakis [19] proved that the problem of computing a pure Nash equilibrium in constraint satisfaction games is complete for the class PLS — standing for polynomial local search — that has been introduced by Johnson et al. [11]. The negative result of [19] covers all games considered in the current work and have been strengthened in [12,13] to capture instances in which each player participates in a constant number of constraints. Among the few rare non-trivial positive results is an algorithm by Poljac [17] who shows that a local maximum of the potential function in cut games can be computed in polynomial time when each player participates in at most three constraints.

The algorithm of [3] for approximate equilibria in P_k -FLIP games has the following structure. Players are partitioned into layers in a similar way to the block partitioning that we use in the current paper. Then, a rearrangement phase moves players across blocks in order to guarantee that the total weight of constraints, in which a player i participates together only with players in the same block or ones having lower maximum utility, is at least $1/k$ of player i 's maximum utility. This can be done in such a way that, eventually, each layer contains players with polynomially-related maximum utility. Then, a top-down layer dynamics phase takes place, where players within each layer play $(1 + \varepsilon/k)$ -moves in a restricted game among them until they reach an $(1 + \varepsilon/k)$ -approximate equilibrium in this restricted game. The authors of [3] show that the state computed in this way is a $(2k - 1 + \varepsilon)$ -approximate equilibrium for the original game. They also present a variation of their algorithm for NAE- (\bar{k}, k) -FLIP games that computes $(\frac{2\bar{k}}{k-1} + \varepsilon)$ -approximate equilibria. As the authors of [3] emphasize, in general, the moves during the top-down layer dynamics phase are not improvement moves in the original game. In contrast, our algorithm consists only of improvement moves.

Another class of potential games where the problem of computing an (approximate) equilibrium has received a lot of attention is that of congestion games. A classical potential function for these games has been defined by Rosenthal [18]. Fabrikant et al. [8] prove that computing a local minimum of this function (corresponding to a pure Nash equilibrium) is PLS-hard as well. Even worse,

for sufficiently general congestion games, Skopalik and Voecking [20] show that computing a ρ -approximate equilibrium is PLS-hard for every reasonable (i.e., polynomially-computable) value of ρ . In our previous work [4], we have presented an algorithm to compute $O(1)$ -approximate equilibria for congestion games under mild assumptions for the structure of the game. The current paper adapts and extends the main algorithmic techniques in that paper, which have also been applied to (non-potential) weighted variants of congestion games in [5]. Exact or almost exact equilibria can be computed in several special cases, e.g., see [6,8].

We remark that, even though it is hard to compute exactly, a local optimum of a potential function can be approximated with extremely low precision under very mild assumptions [15]. This does not imply that equilibria can be approximated with a similar precision, as the negative results of [20] show. Also, uncoordinated move sequences have been shown to reach states of high social value quickly [1,2,7], i.e., to states with low potential in the case of P_k -FLIP games. Unfortunately, these states are not approximate equilibria either, since some player typically has a high incentive to move.

Roadmap. The rest of the paper is structured as follows. We begin with preliminary definitions in Section 2. Section 3 is devoted to our upper bounds on the stretch of P_k -FLIP games. The algorithm and the statement of our main result are presented in Section 4 and the analysis follows in Section 5. We conclude with open problems in Section 6.

2 Preliminaries

A constraint satisfaction game consists of a set N of n players, a set of at least n boolean variables $V = \{s_1, s_2, \dots, s_{|V|}\}$, and a set C of constraints (henceforth called clauses) over the variables in V . Each clause $c \in C$ has a non-negative weight w_c . Player $j \in N$ controls the value of a distinct variable s_j from V and has two possible strategies: setting the value of s_j to either 0 (false), or 1 (true). The variables of V that are not controlled by any player (if any) are frozen to certain boolean values. A state S of the game is simply a snapshot of variable values (or a snapshot of players' strategies complemented with the fixed values of the frozen variables), i.e., $S = (s_1, s_2, \dots, s_{|V|})$. Given a state S of the game, we denote by $SAT(S)$ the set of satisfied clauses. For a subset of players $R \subseteq N$, we denote by $SAT_R(S)$ the subset of $SAT(S)$ that consists of clauses in which the variable of some player from R appears. With some abuse of notation, we simplify $SAT_{\{j\}}(S)$ to $SAT_j(S)$. The utility of a player j is the total weight of the true clauses in which her variable appears, i.e., $u_j(S) = \sum_{c \in SAT_j(S)} w_c$. We also denote by C_R the set of clauses in which at least one player of R participates and simplify $C_{\{j\}}$ to C_j . We use U_j to denote the maximum possible utility that player j might have, i.e., $U_j = \sum_{c \in C_j} w_c$.

Given a state $S = (s_1, s_2, \dots, s_{|V|})$ and a player j , we denote by (S_{-j}, s'_j) the state obtained from S when player j unilaterally changes her strategy from s_j to its complement s'_j . This is an improvement move (or simply, a move) for player j if her utility increases, i.e., $u_j(S_{-j}, s'_j) > u_j(S)$. We call it a ρ -move when the

utility increases by more than a factor of ρ , i.e., $u_j(S_{-j}, s'_j) > \rho \cdot u_j(S)$. A state S is a pure Nash equilibrium (or simply, an equilibrium) if no player has a move to make. Similarly, S is a ρ -approximate (pure Nash) equilibrium if no player has a ρ -move.

We specifically consider clauses with the following property: any false clause can become true by changing the value of any of its variables. We will refer to games with clauses satisfying this property and with at most k variables per clause as P_k -FLIP games. This class is broad enough and contains (generalizations of) several well-studied games such as cut games and party affiliation games. We are particularly interested in two subclasses of P_k -FLIP games. A NAE-clause contains literals (i.e., variables or their negations) and equals 1 if and only if there are two literals with different values. We will refer to games consisting of NAE-clauses with at least $\bar{k} \geq 2$ and most k literals as NAE- (\bar{k}, k) -FLIP games. Observe that these games are P_k -FLIP games since changing the value of any variable that appears in a clause can change the value of the clause from 0 to 1. In PARITY- k -FLIP games, each clause is characterized as odd or even; an odd (respectively, even) clause is true if the number of its variables which are 1 is odd (respectively, even). An important property of P_k -FLIP games is that for any state S and any player j , it holds that $U_j \leq u_j(S) + u_j(S_{-j}, s'_j)$.

Given a state S of a P_k -FLIP game, we denote by $\Phi(S)$ the total weight of all true clauses, i.e., $\Phi(S) = \sum_{c \in SAT_N(S)} w_c$. The function Φ is a potential function for this game. In particular, it has the remarkable property that for every two states S and (S_{-j}, s'_j) differing only in the strategy of player j , the difference of the potential is equal to the difference of the utility of player j , i.e., $\Phi(S) - \Phi(S_{-j}, s'_j) = u_j(S) - u_j(S_{-j}, s'_j)$.

In the following, we will be often considering sequences of moves in which only players in a certain subset $R \subseteq N$ are allowed to move. We can view such moves as moves in a subgame among the players in R , with the set of clauses C_R (each clause in C_R has the same weight as in the original game), and with fixed values for the variables that are not controlled by players in R . Observe that any subgame of a P_k -FLIP game is a P_k -FLIP game as well. Similarly, any subgame of a NAE- (\bar{k}, k) -FLIP (respectively, PARITY- k -FLIP) game is a NAE- (\bar{k}, k) -FLIP (respectively, PARITY- k -FLIP) game as well. The function $\Phi_R(S) = \sum_{c \in SAT_R(S)} w_c$ is an exact potential function for the subgame among the players in R . The next lemma follows easily by the definitions.

Lemma 1. *For every state S of a P_k -FLIP game and any set of players $R \subseteq N$, it holds that $\Phi_R(S) \leq \sum_{j \in R} u_j(S) \leq k\Phi_R(S)$. Furthermore, for every set of players $R' \subseteq R$, it holds that $\Phi_{R'}(S) \leq \Phi_R(S)$.*

3 The Stretch of P_k -FLIP Games

The approximation guarantee of our algorithm depends on a quantity related to the potential function of P_k -FLIP games that we call the stretch.

Definition 1. *Given $\eta \geq 0$, the $(1 + \eta)$ -stretch of a P_k -FLIP game is the ratio between the maximum and the minimum value of the potential function taken over all $(1 + \eta)$ -approximate pure Nash equilibria of the game.*

We use the term stretch as a synonym of 1-stretch; observe that it is simply the ratio between the maximum and minimum potentials of (exact) equilibria. In Theorem 1, we present upper bounds on the $(1 + \eta)$ -stretch of P_k -FLIP games. Note that these bounds may be of independent interest; bounds on the stretch of congestion games from our previous work [4] have been used by Piliouras et al. [16] in order to quantify the price of anarchy of congestion games in settings with uncertainty where players have particular risk attitudes.

Theorem 1. *For any $\eta > 0$, the $(1 + \eta)$ -stretch of P_k -FLIP games, NAE- $(3, k)$ -FLIP games, NAE- $(2, k)$ -FLIP games, and PARITY- k -FLIP games with odd k is at most $k + 1 + k\eta$, $2 + k\eta$, $3 + k\eta$, and $k + k\eta$, respectively.*

Due to lack of space, the proof (as well as counterexamples showing that the bounds are tight) is omitted. In the following, we use the notation $\theta(1 + \eta)$ to denote our upper bound on the $(1 + \eta)$ -stretch of P_k -FLIP games (and clarify when we refer to the stretch of particular subclasses of P_k -FLIP games). We use simply θ to denote the upper bound on the 1-stretch.

4 The Algorithm

The pseudocode of our algorithm appears below as Algorithm 1. We supplement this formal description with a detailed line-by-line explanation. The algorithm takes as input a P_k -FLIP game \mathcal{G} with n players, an initial state S_{in} , and an accuracy parameter $\varepsilon \in (0, 1]$. Starting from state S_{in} , it identifies a sequence of moves that lead to a state S_{out} ; this is the output of the algorithm. As we will prove later, S_{out} is an approximate equilibrium. The algorithm starts (lines 1 and 2) by setting the values of parameters q and p . Parameter q has a value very close to 1 (namely, $q = 1 + \frac{\varepsilon}{3k}$) and parameter p has a value slightly higher than the q -stretch of the class to which the input game belongs (namely, $p = \theta(q) + \varepsilon/3$). In particular, using our upper bounds on $\theta(q)$ from Theorem 1, p is set to be $k + 1 + 2\varepsilon/3$ in general, $2 + 2\varepsilon/3$ if \mathcal{G} is a NAE- $(3, k)$ -FLIP game, $3 + 2\varepsilon/3$ if it is a NAE- $(2, k)$ -FLIP games, and $k + 2\varepsilon/3$ if it is a PARITY- k game and k is odd. The algorithm also sets the value of parameter Δ to be a polynomial depending on n , k , p , and $1/\varepsilon$ (line 3). Then (lines 4-5), it implicitly partitions the players into blocks B_1, B_2, \dots, B_m according to their maximum utility. Denoting by U_{max} the maximum values among all players' maximum utilities, block B_i consists of the players j with maximum utility $U_j \in (U_{\text{max}}\Delta^{-i}, U_{\text{max}}\Delta^{1-i}]$. By the definition of Δ , the players in the same block have polynomially related maximum utilities.

The sequence of moves from state S_{in} to state S_{out} is computed by the code in the lines 6-15. The subsequence of moves described in lines 7-9 constitutes phase 0. During phase 0, the players in block B_1 make q -moves. After that, each phase i for $i \geq 1$ consists of p -moves of players in block B_i and q -moves of players

in block B_{i+1} . Strategies of players in block B_i are irrevocably decided at the end of phase i .

We are ready to state our main result which we will prove in the next section.

Theorem 2. *On input a P_k -FLIP game \mathcal{G} with n players, an initial state S_{in} , and $\varepsilon \in (0, 1]$, Algorithm 1 computes a sequence of at most $\text{poly}(n, k, 1/\varepsilon)$ moves that starts from S_{in} and converges to a $(k + 1 + \varepsilon)$ -approximate pure Nash equilibrium S_{out} . The approximation guarantee is at most $2 + \varepsilon$ when \mathcal{G} is a NAE- $(3, k)$ -FLIP game, at most $3 + \varepsilon$ when it is a NAE- $(2, k)$ -FLIP games, and at most $k + \varepsilon$ when it is a PARITY- k -FLIP game and k is odd.*

Input : A P_k -FLIP game \mathcal{G} with a set N of n players, an arbitrary initial state S_{in} , and $\varepsilon \in (0, 1]$

Output: A state S_{out} of \mathcal{G}

- 1 $q \leftarrow 1 + \frac{\varepsilon}{3k}$;
- 2 $p \leftarrow \theta(q) + \varepsilon/3$;
- 3 $\Delta = 200p^3nk/\varepsilon^2$;
- 4 Set $U_{\min} \leftarrow \min_{j \in N} U_j$, $U_{\max} \leftarrow \max_{j \in N} U_j$, and $m \leftarrow 1 + \lceil \log_{\Delta} (U_{\max}/U_{\min}) \rceil$;
- 5 (Implicitly) partition players into blocks B_1, B_2, \dots, B_m , such that $j \in B_i$ implies that $U_j \in (U_{\max}\Delta^{-i}, U_{\max}\Delta^{1-i}]$;
- 6 $S \leftarrow S_{in}$;
- 7 **while** there exists a player $j \in B_1$ such that $u_j(S_{-j}, s'_j) > q \cdot u_j(S)$ **do**
- 8 $S \leftarrow (S_{-j}, s'_j)$;
- 9 **end**
- 10 **for** phase $i \leftarrow 1$ **to** $m - 1$ such that $B_i \neq \emptyset$ **do**
- 11 **while** there exists a player j that either belongs to B_i and satisfies $u_j(S_{-j}, s'_j) > p \cdot u_j(S)$ or belongs to B_{i+1} and satisfies $u_j(S_{-j}, s'_j) > q \cdot u_j(S)$ **do**
- 12 $S \leftarrow (S_{-j}, s'_j)$;
- 13 **end**
- 14 **end**
- 15 $S_{out} \leftarrow S$;

Algorithm 1. Computing approximate equilibria in P_k -FLIP games.

5 Proof of Theorem 2

Before presenting the proof of Theorem 2, we give some intuition behind our analysis. The analysis uses two properties that are formally stated in Lemma 2. What this lemma essentially says is that, during each phase, the total utility of the moving players as well as an increase in the potential of the subgame among these players are small. The first property is used in Lemma 4 to prove that, once the strategy of a player is irrevocably decided, later phases may have only a negligible effect on her. And since no player has a p -move to make at the end of the phase when her strategy is decided, she cannot improve her utility

by a factor of (almost) p until the end of the algorithm. Together with the fact that each player's move increases her utility by some non-negligible amount, the second property is used in Lemma 5 to bound the total number of moves.

In our analysis, we denote by S^i the state reached at the end of phase $i \geq 0$, i.e., $S_{\text{Out}} = S^{m-1}$. We also denote by R_i the set of players that move during phase i . We also denote the upper boundary of block B_i by W_i and by W_{m+1} the lower boundary of block B_m , i.e., $W_i = U_{\max} \Delta^{1-i}$ for $i = 1, 2, \dots, m+1$. So, the players of block B_i are those with maximum utility $U_j \in (W_{i+1}, W_i]$.

Lemma 2. *For every phase $i \geq 1$, it holds that*

1. $\sum_{j \in R_i} U_j \leq 10pknW_{i+1}/\varepsilon$
2. $\Phi_{R_i}(S^i) - \Phi_{R_i}(S^{i-1}) \leq 3p^2nW_{i+1}/\varepsilon$.

Proof. First observe that players not in R_i have the same set of strategies in states S^{i-1} and S^i . Furthermore, the total weight of clauses depending on variables that are controlled by players from $R_i \cap B_{i+1}$ is at most nW_{i+1} . Hence, by the definition of the subgame potential, we have that the potential of the state $(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i)$ in which the players in $R_i \cap B_i$ play their strategies in state S^i and the remaining players play their strategies in S^{i-1} satisfies

$$\Phi_{R_i \cap B_i}(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i) \geq \Phi_{R_i}(S^i) - nW_{i+1}. \quad (1)$$

We will use inequality (1) in the proof of the next lemma that provides a bound on the potential $\Phi_{R_i}(S^{i-1})$ as well as later in the current proof.

Lemma 3. $\Phi_{R_i}(S^{i-1}) \leq 3pnW_{i+1}/\varepsilon$.

Proof. We assume on the contrary that $\Phi_{R_i}(S^{i-1}) > 3pnW_{i+1}/\varepsilon$ and we are going to conclude that the potential of the state $(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i)$ satisfies $\Phi_{R_i \cap B_i}(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i) > \theta(q) \cdot \Phi_{R_i \cap B_i}(S^{i-1})$. By Theorem 1, this would contradict the fact that S^{i-1} is the output of phase $i-1$, i.e., a q -approximate equilibrium of the subgame among the players in $R_i \cap B_i$, since there is another q -approximate equilibrium (the one that can be reached from $(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i)$ with q -moves by the players in $R_i \cap B_i$) with a potential that is higher than $\theta(q)$ times the potential at state S^{i-1} .

We denote by $\ell(j)$ the utility of player $j \in R_i \cap B_i$ right after she makes her last move in phase i . Then we have

$$\Phi_{R_i}(S^i) - \Phi_{R_i}(S^{i-1}) \geq (1 - 1/p) \cdot \sum_{j \in R_i \cap B_i} \ell(j). \quad (2)$$

Indeed, the last move of a player $j \in R_i \cap B_i$ increases her utility by a factor of at least p and the difference $\Phi_{R_i}(S^i) - \Phi_{R_i}(S^{i-1})$ equals to the total increase in the utility of the deviating players within the phase.

Furthermore, we claim that

$$\sum_{j \in R_i \cap B_i} \ell(j) + nW_{i+1} \geq \Phi_{R_i}(S^i). \quad (3)$$

To see why (3) is true, observe that the right-hand side is the sum of the weights of the clauses in $SAT_{R_i}(S^i)$. The term nW_{i+1} is an upper bound on the total weight of the clauses in $SAT_{R_i \cap B_{i+1}}(S^i)$. The weight of each of the remaining ones (i.e., the clauses in $SAT_{R_i}(S^i) \setminus SAT_{R_i \cap B_{i+1}}(S^i)$) is accounted for at least once in the sum $\sum_{j \in R_i \cap B_i} \ell(j)$, as part of the utility of some player from $R_i \cap B_i$ after her last move.

By (2) and (3) (i.e., by multiplying (2) by p and (3) by $p - 1$ and summing them), we obtain that

$$\Phi_{R_i}(S^i) \geq p \cdot \Phi_{R_i}(S^{i-1}) - (p - 1)nW_{i+1}. \quad (4)$$

Hence, using (1), (4), the definition of p , and the second inequality of Lemma 1, we obtain

$$\begin{aligned} \Phi_{R_i \cap B_i}(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i) &\geq \Phi_{R_i}(S^i) - nW_{i+1} \geq p \cdot \Phi_{R_i}(S^{i-1}) - pnW_{i+1} \\ &> (p - \varepsilon/3) \cdot \Phi_{R_i}(S^{i-1}) \geq \theta(q) \cdot \Phi_{R_i \cap B_i}(S^{i-1}). \end{aligned}$$

We have obtained the desired contradiction. \square

Using the observation that no player in $R_i \cap B_i$ has a q -move to make at the end of phase $i - 1$ (i.e., at state S^{i-1}) as well as the first inequality of Lemma 1, we obtain that

$$\begin{aligned} \sum_{j \in R_i \cap B_i} U_j &\leq \sum_{j \in R_i \cap B_i} (u_j(S^{i-1}) + u_j(S_{-j}^{i-1}, s'_j)) \leq \sum_{j \in R_i \cap B_i} (1 + q)u_j(S^{i-1}) \\ &\leq (1 + q)k \cdot \Phi_{R_i \cap B_i}(S^{i-1}) \leq 9pknW_{i+1}/\varepsilon. \end{aligned}$$

The proof of the first inequality in the statement of the lemma follows by observing that the total utility of the players in $R_i \cap B_{i+1}$ is at most nW_{i+1} .

In order to prove the second inequality we use inequality (1), the q -stretch bound for the subgame among the players in $R_i \cap B_i$, the fact that $\theta(q) \leq p$, the second inequality of Lemma 1, and the bound on $\Phi_{R_i}(S^{i-1})$ from Lemma 3.

$$\begin{aligned} \Phi_{R_i}(S^i) - \Phi_{R_i}(S^{i-1}) &\leq \Phi_{R_i \cap B_i}(S_{-R_i \cap B_i}^{i-1}, S_{R_i \cap B_i}^i) - \Phi_{R_i}(S^{i-1}) + nW_{i+1} \\ &\leq \theta(q) \cdot \Phi_{R_i \cap B_i}(S^{i-1}) - \Phi_{R_i}(S^{i-1}) + nW_{i+1} \\ &\leq (p - 1) \cdot \Phi_{R_i}(S^{i-1}) + nW_{i+1} \\ &\leq 3p^2nW_{i+1}/\varepsilon. \end{aligned}$$

\square

The first property of Lemma 2 indicates that the total weight of the moving players in phase i is significantly smaller than the upper boundary of block B_i . In Lemma 4 we combine this with the fact that the upper boundary of subsequent blocks decreases exponentially and formally prove that, after the strategy of a player is irrevocably decided, subsequent phases may have only a negligible effect on her. Recall that θ is the stretch of the class of games to which the input game belongs to and equals $k + 1$ for P_k -FLIP games, 3 for NAE-(2, k)-FLIP games, and 2 for NAE-(3, k)-FLIP games, and $k - 1$ for PARITY- k -FLIP games with odd k .

Lemma 4. *The state S_{out} is a $(\theta + \varepsilon)$ -approximate pure Nash equilibrium.*

Proof. By the definition of phase $m - 1$, the players in blocks B_{m-1} and B_m have no p -move to make at the end of phase $m - 1$. We will consider a player j belonging to block B_t whose strategy is irrevocably decided at the end of phase t with $t \leq m - 2$, and will show that she has no $(p + \varepsilon/3)$ -move to make at the end of phase $m - 1$ (i.e., at state $S^{m-1} = S_{\text{out}}$). The lemma will then follow since $p + \varepsilon/3 = \theta(1 + \frac{\varepsilon}{3k}) + 2\varepsilon/3 = \theta + \varepsilon$.

Let s_j be the strategy used by player j at the end of phase t . Using Lemma 2 and the definition of the block boundaries, we can bound the quantity $\sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r$. Thus, we get an upper bound on the total weight of clauses with players that move in phases $t + 1, \dots, m - 1$, as follows:

$$\begin{aligned} \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r &\leq \sum_{i=t+1}^{m-1} 10pnkW_{i+1}/\varepsilon \leq \frac{10pnkW_{t+1}}{\varepsilon} \sum_{i=1}^{\infty} \Delta^{-i} \\ &= \frac{10pnkW_{t+1}}{\varepsilon(\Delta - 1)} \leq \frac{W_{t+1}\varepsilon}{10p^2}. \end{aligned} \quad (5)$$

The last inequality follows by the definition of Δ and the fact that $\Delta - 1 \geq \Delta/2$.

Now observe that since player j has no p -move at the end of phase t (i.e., at state S^t), it holds that $u_j(S^t) \geq u_j(S_{-j}^t, s'_j)/p$ and $W_{t+1} \leq u_j(S^t) + u_j(S_{-j}^t, s'_j) \leq (1 + p)u_j(S^t)$, i.e., $u_j(S^t) \geq \frac{W_{t+1}}{1+p}$. Furthermore, during phases $t + 1, \dots, m - 1$, the total change in the utility of player j or in the utility player j would have by deviating is at most $\sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r$. Using these observations and inequality (5), we have

$$\begin{aligned} u_j(S^{m-1}) &\geq u_j(S^t) - \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r \\ &\geq \frac{p}{p + \varepsilon/3} u_j(S^t) + \frac{\varepsilon/3}{p + \varepsilon/3} \frac{W_{t+1}}{1 + p} - \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r \\ &\geq \frac{1}{p + \varepsilon/3} u_j(S_{-j}^t, s'_j) + \frac{W_{t+1}\varepsilon}{5p(p + \varepsilon/3)} - \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r \\ &\geq \frac{1}{p + \varepsilon/3} u_j(S_{-j}^{m-1}, s'_j) + \frac{W_{t+1}\varepsilon}{5p(p + \varepsilon/3)} - \left(1 + \frac{1}{p + \varepsilon/3}\right) \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r \\ &\geq \frac{1}{p + \varepsilon/3} u_j(S_{-j}^{m-1}, s'_j) + \frac{W_{t+1}\varepsilon}{5p(p + \varepsilon/3)} - \frac{2p}{p + \varepsilon/3} \sum_{i=t+1}^{m-1} \sum_{r \in R_i} U_r \\ &\geq \frac{1}{p + \varepsilon/3} u_j(S_{-j}^{m-1}, s'_j), \end{aligned}$$

as desired. In the third and fifth inequalities we have used the inequalities $3(1 + p) \leq 5p$ and $p + 1 + \varepsilon/3 \leq 2p$ which follow since $p \geq 2$ and $\varepsilon \in (0, 1]$. This completes the proof of the lemma. \square

To conclude the proof of Theorem 2, it remains to bound the running time of the algorithm; the proof of the next statement is omitted due to lack of space.

Lemma 5. *On input of P_k -FLIP (in particular, NAE- (\bar{k}, k) -FLIP) game, the algorithm identifies a sequence of at most $\mathcal{O}(n^3 k^7 / \varepsilon^4)$ (in particular, $\mathcal{O}(n^3 k^2 / \varepsilon^4)$) moves.*

6 Open Problems

A challenging open problem is to improve the approximation guarantee of our algorithm. Our analysis indicates that a state with lower stretch at the beginning of each phase would allow us to use an even smaller value for parameter p and, subsequently, to obtain a better approximation guarantee. One idea that comes immediately to mind is to replace the q -moves of the players of block B_{i+1} within phase i with the execution of an algorithm that computes states with approximately-optimal potential. For example, a random assignment to players of B_{i+1} would yield a 2-approximation to the potential of the subgame among them. Furthermore, for more structured P_k -FLIP games such as cut games, one might think to use the famous algorithm of [9] that is based on semi-definite programming. Unfortunately, we do not see how to include these ideas into our algorithm at this point. The main difficulty is that the low-stretch property should hold for the subgame among the players that will move during the next phase which we do not know in advance. An algorithm that approximates the potential of all subgames simultaneously would be ideal here but, besides the local search approach implied by the q -moves, neither the random assignment nor the SDP-based algorithms satisfy this property.

Even if we could bypass these obstacle, our technique has limitations since computing states with low-stretch in P_k -FLIP games includes famous hard-to-approximate problems (e.g., see [10]). So, in order to compute almost exact equilibria, we need new techniques. Of course, we have no idea whether this is at all possible. To put the question differently, is there some inapproximability threshold for approximate equilibria? We remark that such negative statements are not known in the literature: the only known negative results are either specific to exact equilibria (such as the PLS-hardness results of [8,19]) or rule out any reasonable approximation guarantee in games with very general structure (e.g., in [20]). We believe that such questions that are related to the computational complexity of approximate pure Nash equilibria deserve further attention.

References

1. Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V.S., Skopalik, A.: Fast convergence to nearly optimal solutions in potential games. In: Proceedings of the 9th ACM Conference on Electronic Commerce (EC), pp. 264–273 (2008)
2. Balcan, M.-F., Blum, A., Mansour, Y.: Circumventing the price of anarchy: leading dynamics to good behavior. *SIAM Journal on Computing* 42(1), 230–264 (2013)

3. Bhargat, A., Chakraborty, T., Khanna, S.: Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In: Proceedings of the 11th ACM Conference on Electronic Commerce (EC), pp. 73–82 (2010)
4. Caragiannis, I., Fanelli, A., Gravin, N., Skopalik, A.: Efficient computation of approximate pure Nash equilibria in congestion games. In: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 532–541 (2011)
5. Caragiannis, I., Fanelli, A., Gravin, N., Skopalik, A.: Approximate pure Nash equilibria in weighted congestion games: existence, efficient computation, and structure. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 284–301 (2012)
6. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior* 71(2), 315–327 (2011)
7. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. *Theoretical Computer Science* 438, 13–27 (2012)
8. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC), pp. 604–612 (2004)
9. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM* 42(6), 1115–1145 (1995)
10. Hastad, J.: Some optimal inapproximability results. *Journal of the ACM* 48(4), 798–859 (2001)
11. Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How easy is local search? *Journal of Computer and System Sciences* 37, 79–100 (1988)
12. Klauck, H.: On the hardness of global and local approximation. In: Karlsson, R., Lingas, A. (eds.) SWAT 1996. LNCS, vol. 1097, pp. 88–99. Springer, Heidelberg (1996)
13. Krentel, M.W.: Structure in locally optimal solutions. In: Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS), pp. 216–221 (1989)
14. Michiels, W., Aarts, E., Korst, J.: Theoretical aspects of local search. In: EATCS Monographs in Theoretical Computer Science. Springer (2007)
15. Orlin, J.B., Punnen, A.P., Schulz, A.S.: Approximate local search in combinatorial optimization. *SIAM Journal on Computing* 33, 1201–1214 (2004)
16. Piliouras, G., Nikolova, E., Shamma, J.S.: Risk sensitivity of price of anarchy under uncertainty. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 715–732 (2013)
17. Poljak, S.: Integer linear programs and local search for max-cut. *SIAM Journal on Computing* 24, 822–839 (1995)
18. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
19. Schäffer, A.A., Yannakakis, M.: Simple local search problems that are hard to solve. *SIAM Journal on Computing* 20, 56–87 (1991)
20. Skopalik, A., Vöcking, B.: Inapproximability of pure Nash equilibria. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), pp. 355–364 (2008)