Approximate Pure Nash Equilibria via Lovász Local Lemma

Thành Nguyen* and Éva Tardos**

Cornell University

Abstract. In many types of games, mixed Nash equilibria is not a satisfying solution concept, as mixed actions are hard to interpret. However, pure Nash equilibria, which are more natural, may not exist in many games. In this paper we explore a class of graphical games, where each player has a set of possible decisions to make, and the decisions have bounded interaction with one another. In our class of games, we show that while pure Nash equilibria may not exist, there is always a pure approximate Nash equilibrium. We also show that such an approximate Nash equilibrium can be found in polynomial time. Our proof is based on the Lovász local lemma and Talagrand inequality, a proof technique that can be useful in showing similar existence results for pure (approximate) Nash equilibria also in other classes of games.

1 Introduction

In his Nobel prize winning work, John Nash proved the existence of mixed equilibria, which are now called Nash equilibria. A nice property of Nash equilibria is that they exist in (almost) any game. On the other hand, in many applications, for example in network design problem [3] and facility locations [21], agents need to make a decision that then will be visible to all other players, such as locating a facility (network router, a store, a server, etc.). In such games, there is no natural interpretation of a randomized action of the players. Before the decision is made, it makes sense to think of two locations as equally likely for being the selected choice. But once a decision is made, and assuming this decision is observable for all other players, these other players will react to the actual decision made, and the fact that *a priori* another decision was equally likely becomes irrelevant. This is especially true if, once the decision is made, it becomes hard to undo, such as locating a facility that requires significant investment cost. In such contexts pure Nash equilibria are much more natural, but unfortunately, they may not exist. Even in games where pure Nash equilibria is known to exist (such as congestion games, for example), it was shown that finding a pure Nash equilibria is PLS-complete [11].

* Supported in part by ONR grant N00014-08-1-0031.

** Supported in part by ITR grant CCR-0325453, CCR-0729006, and ONR grant N00014-08-1-0031.

Our interest in this paper is to explore the existence of pure Nash equilibria (or approximate equilibria). Recent results on the hardness of computing Nash equilibria [8, 5] inspired work on finding approximate equilibria, see for example [7]. In this paper, we attempt to start similar investigations for approximate pure Nash equilibria. In a general model, we proved that there exists pure constant approximate Nash, and we can find it in polynomial time. In this paper, we also introduce a new technique for proving the existence of approximate Nash equilibria.

Our model The class of games we consider is a variant of graphical games, where different players interact in limited ways. Graphical game is a general class of games that successfully captures and exploits the locality and sparsity of direct influences in games with large number of players. In a graphical game we are given an undirected graph, in which players are identified with vertices, and a player's payoff function is entirely determined by the action of the player and his/her neighbors. In this paper we consider a slightly different class with limited direct interactions. We will assume that each player makes many decisions, and assume that each decision is directly influenced by only a limited number of other decisions (while a player can influence a large set of other players through his many decisions).

More precisely, in the game we consider, each player i will be represented as a set of vertices S_i, one vertex for each of the decisions the player has to make. In each vertex, the player has two options, to play 0 or 1. We can think of the decision of playing 1 as representing the strategy to locate a service at this location, or start a business, develop a new product, etc. In this game, each player has a complex strategy set: the number of strategies is $2^{|S_i|}$ corresponding to all possible subsets of S_i. We will describe the direct interactions of these decisions by an undirected graph G on all the vertices $\cup_i S_i$, with the set of edges connecting a vertex of a player to another player's vertex. We call this graph *interaction graph* .

We model the utility function of a player in two steps. First, on each possible location j (vertex in $j \in S_i$) there is an *outcome function*, the outcome of the decision player i made on this vertex, which is a function mapping decisions made at all vertices connected to the vertex j to a multiple dimensional vector in $\{0, 1\}^h$. We'll think of the outcome as the level of success of the decision made in several criteria. For example, a product can be a success in one part of the market but not in the other, or building a factory at a location can create profit for the company, offer jobs for the locals but can cause environmental problems as well. A player in this model sometimes needs to make a decision that balances the trade-off between many factors. We model the payoff function of each player i as a MAXSAT formula of the whole outcome vector on S_i.

The proposed model of outcomes, and utility functions via a MAXSAT expression is very general, allowing us to model interests by players in many aspects. A term in the MAXSAT formula can model the success of each of the player's ventures, but other terms can express combinatorial goals. For example, in the case of companies developing products, one of the goals can be: at least one of the given products needs to be successful in a given market. The outcome function is an arbitrary function and therefore can capture some complex situations. For example, when developing a strategy whether to invest and build facility in various locations or to develop some new products, a company needs to take into account many factors: whether there will be too many competitors in the location nearby, or too many similar products in the market. In some cases a product can only be successful if there are some other products available on the market. Our model captures some of these types of problems. A more precise description of the model is presented in Section 2.

Our result and technique In this paper we will prove that if we assume the utility functions satisfy a Lipschitz condition, that is $U_i(X_i) - U_i(X_i') \leq \Delta$ whenever X_i, X_i' differ on one coordinate, then the game has a pure approximate Nash equilibrium, at which each player i has an payoff at least

$$OPT_i/2 - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right),$$

where OPT_i is the optimal payoff he can get by a deviation, d is the maximum degree in the interaction graph G and n is the maximum size of a player's action set S_i, that is $n = \max_i |S_i|$. Furthermore, such a configuration can be found in polynomial time.
Note that when $OPT_i >> \Delta, \log n, \log d$ then this is a constant approximate Nash. Furthermore, in the case of using MAXSAT to describe user utility functions, finding a better than constant approximate Nash is unlikely, as finding a better than constant approximation for even a single player's optimization problem is NP hard [13].
In many cases, the assumption that OPT_i is relatively larger than $\log n$ and $\log d$ is reasonable, as the optimal solution of a MAXSAT formula is at least a constant times the number of its clauses. In our model if the number of clauses is much larger than $\log n$ and the outcome functions are not all "constant", that is, for a player by changing his decision on a vertex from 0 to 1 or from 1 to 0, the outcome of this decision also changes, then OPT_i satisfies the condition above.
At the heart of our proof is the use of Lovász local lemma and Talagrand's inequality. The Lipschitz condition is used extensively for a concentration bound using Talagrand's inequality. Lovász local lemma is used to prove the guarantee that every player has a payoff near the optimal. These techniques are nontrivial and general, and we hope that they will be useful in proving similar results about approximate Nash equilibria in many other settings. One such example is the case when instead of MAXSAT utility functions, we have a class of MAXSUBMODULAR utilities. It is known [12] that a random set gives a constant approximation. As to be seen later, our technique can be used in this setting as well.

A bad example To illustrate the difficulties and some of the ideas in our techniques, let us start with the following example:
There are two players in the game. The first player plays on n vertices $u_1, u_2, .., u_n$ and the second one plays on $v_1, v_2, .., v_n$. Let $o(u_i)$ and $o(v_i) \in \{0, 1\}$ be the outcome on vertex u_i and v_i respectively. The utility function of player 1 is a SAT formula consisting of the

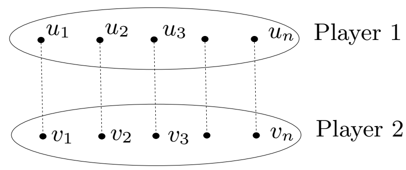

Fig. 1. An example

clauses: $(o(u_1) \vee o(u_2)), (o(u_1) \vee o(u_3)), .., (o(u_1) \vee o(u_n))$. Similarly, the utility function of player 2 is a SAT formula consisting of the clauses: $(o(v_1) \vee o(v_2)), (o(v_1) \vee o(v_3)), .., (o(v_1) \vee o(v_n))$. The underlying graph is a matching $\{u_i v_i\}$. Now the outcome functions are defined as follows: On all the vertices u_i, v_i, where $i \geq 2$, the outcome function is 0 no matter what the players' strategies are. On v_1 and u_1, the outcome function is the same as the payoff function of the "matching penny" game described by the following table:

	0	1
0	$(0,1)$	$(1,0)$
1	$(1,0)$	$(0,1)$

It is not hard to see that it is a mixed strategy Nash if player 1 plays 0 or 1 with 1/2 probability on u_1 and player 2 plays 0 or 1 with 1/2 probability on v_1 . Given any deterministic configuration of the game, there is always a player that can improve his payoff from 0 to $n-1$ by changing his strategy on his first vertex. Thus, the game not only has no pure Nash equilibria but also has no "reasonable" approximate pure Nash equilibria. It turns out that the main obstacle in this example is the property that in the utility functions there is a variable such that by changing its value the utility function changes rapidly.

Related works The complexity of mixed Nash equilibira is studied in [8, 5]. In many applications, such as most of the network formation games, pure Nash Equilibria are usually considered a more realistic model of rationality. See the [2, 3, 21] for detail.
Potential games [18] is essentially the only class of games that is known to have pure Nash equilibria. The complexity of finding a pure Nash in potential games is proved to be PLS-complete [11]. Inspired by these hardness results, many researchers have been investigating pure approximate Nash equilibria in various games [2, 19].
Graphical games were introduced in [15], and have been extensively studied. (See the survey [16].) The complexity of finding pure Nash equilibria in graphical games were studied in [6, 9]. Most of the results concerning pure NE in graphical games are either negative or for a small class of graphs such as trees or graphs with bounded tree-width. Our model is a variant of graphical games. Here, instead of representing each player as a vertex in the graph, we consider each player as a set of vertices and

thus, it models more complex strategy sets for players. In this paper, we provide a positive result for the existence of a pure approximate Nash in general graphs that can be found in polynomial time.
Lovász local lemma is proved in [10]. Lovász local lemma is a powerful tool in probability and has a wide range of applications, see the book of [1] for details and references. To the best of our knowledge, our paper is the first application of the local lemma in the area of algorithmic game theory.

Structure of the text In the next section, we will define our model more precisely. In Section 3 we will prove the existence of an approximate pure NE. Using an algorithmic version of the local lemma, we also obtain an algorithm to find such a solution.

2 The model and notations

Consider a game consisting N players, each player i has a set of possible *vertices* S_i, one vertex for each of the decisions that the player has to make. On each of his/her vertex, player i can choose one of two strategies 0 or 1. We call it the *decision* of player i on the vertex. In this game, each player i can have up to $2^{|S_i|}$ strategies corresponding to all possible subsets of S_i.
We assume that the size of S_i is at most n for every i, an all the set S_i are disjoint. The latter assumption does not affect the generality of our model, as each player can have his/her own copy of a vertex.
Similar to graphical games [15], we model the direct interactions of the players' decisions by an undirected graph G on the set of all the vertices $\cup_i S_i$. The edges of G connects a vertex of a player to another player's vertex. We call G the *interaction graph.* See Figure 2 for an example. In the rest of the paper, we use d as the maximum degree of the interaction graph.
The utility function of a player is modeled in two steps. First, at each vertex $j \in S_i$, there is an *outcome* function. The outcome function on a vertex j is a function mapping the decisions of the players on j and neighbors of j to a finite set of outcomes. In this paper, we consider the case where each outcome is a multiple dimensional vector in $\mathcal{O} = \{0,1\}^h$ for a constant h. For a vertex j, we denote $\Gamma(j)$ as the set of j' neighbors, and o_j the outcome function on j, we have:

$$o_j : \{0,1\}^{|\Gamma(j) \cup j|} \to \mathcal{O} = \{0,1\}^h$$

Now, given a configuration of the game, each player has an outcome on each of his vertices. The utility of player i denoted by U_i is a function mapping the outcome vector on S_i to a non negative number.

$$U_i : \{\mathcal{O}\}^{|S_i|} \to \mathbb{R}^+$$

One way to think about U_i is that it is function on $m_i = h \cdot |S_i|$ boolean variables measuring the success of the decisions made in many aspects. In many cases, there are trade-offs between these decisions and some of the

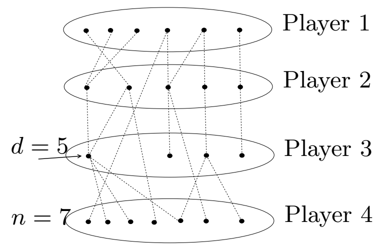

Fig. 2. An example of our model

goals of a player is a combination of several outcomes. We model U_i as a *MAXSAT function*, that is for each U_i there is a Normal Conjunctive Formula (NCF) Φ on boolean variables $x_1, .., x_{m_i}$ such that the value of $U_i(x_1, .., x_{m_i})$ is the number of satisfying clauses in Φ. Note that our results can be naturally extended to a weighted version of MAXSAT.
As we have seen in the counter example in the introduction, we will need a *Lipschitz condition* for the utility functions. This condition assumes that by changing the decision at one vertex, the player's utility does not change very much. More precisely, the utility function U_i, described above as a composition of a MAXSAT and several outcome functions, can be considered as a function of the player i's strategy $X_i \in \{0,1\}^{|S_i|}$ and all other players' $X_{-i} \in \{0,1\}^{|\cup_{j \neq i} S_j|}$. We assume that Δ is the Lipschitz constant for the class of utilities we consider. That is, for every i, $|U_i(X_i, X_{-i}) - U_i(X_i', X_{-i})| \leq \Delta$ whenever X_i, X_i' differ in only one coordinate.
In the rest of the paper, we call the class of games defined above *MAXSAT games* with the *Lipschitz constant* Δ.

3 Existence of approximate pure Nash equilibria

In this section we will prove the following result about the existence of a pure approximate Nash equilibrium for MAXSAT games.

Theorem 1. *In the MAX SAT game, there exists an approximate Nash, where each player i obtains a payoff at least*

$$OPT_i/2 - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right),$$

where OPT_i is the optimum that player i can achieve assuming that other players do not change their strategies, and Δ is the Lipschitz constant of the players' utilities.

We now give some intuition before proving the theorem. Consider an arbitrary configuration of the game, and a player i. If player i assumes that other players do not change their strategies, then on a vertex $j \in S_i$, the outcome function will be a function mapping his decision on j to a

vector in $\{0,1\}^h$. Let x_j be the player's decision on vertex j. The outcome o_j will be a vector whose coordinates can be a constant (independent of x_j) or x_j or $\neg x_j$. Because the player i's utility can be expressed as a MAXSAT function on the outcome vector, it is also a MAXSAT function on $\{x_j, j \in S_i\}$. Therefore, if player i tries to find a strategy to maximize his own utility, he needs to solve a MAXSAT problem. It is, however, known that by assigning each variable to 0 and 1 with probability 1/2, we get at least a 2-approximate solution *in expectation.* We note that in the worst case, this simple algorithm gives the best possible approximation on some instances of MAXSAT, as shown by [13].
With this intuition, let us consider a simple strategy where each player's decision on a location is to pick either 0 or 1 randomly with probability 1/2. This strategy is an 2-approximate *mixed strategy Nash.* However, by selecting the strategies randomly, players might end up in a situation when they can deviate and significantly increase their payoff. Because the number of players N can be arbitrarily large, independent of n and d, the number of such players can also be arbitrarily large. Our idea is to bound the probability that such event happens and then use Lovász local lemma to prove the existence of approximate Nash equilibria. Lovász local lemma, proved in [10], can be stated as follows:

Lemma 1 (Lovász). *Let $A_1, A_2, ..., A_N$ be a series of events such that each event occurs with probability at most p and such that each event is independent of all the other events except for at most D of them. If $e \cdot p \cdot (D+1) < 1$ (where e =2.718...), then there is a nonzero probability that none of the events occur.* □

In order to use Lovász local lemma, we will define A_i as the event that player i can "significantly" improve his payoff by deviating. Thus, to prove that the probability that none of A_i occur is positive, we will need to bound the probability of A_i and to show that each A_i is independent of all but few other A_js. We call the number of such events the *dependence number* of A_i. We first give a bound on the dependence numbers.

Lemma 2. *For every i, the dependence number of A_i is at most $2nd^2$.*

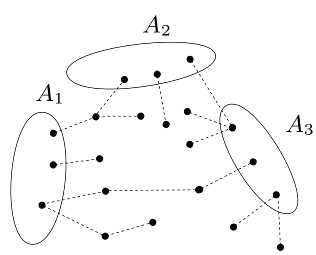

Fig. 3. A_1 and A_3 are independent, A_1 and A_2 are not independent.

Proof. A_i is the event that a player i can find a "significantly" better move. This event depends on the decision of player i on his vertices and the decision of all other players on the neighboring vertices. Consider two events A_i and A_j. If each path connecting two vertices of player i and player j has a distance at least 3, then the set of neighboring vertices of player i and player j are disjoint. Hence A_i and A_j are independent of each other. See Figure 3 for an example. The dependence number of A_i, therefore is bounded by the number of A_j that has a path of length at most 2 connecting from a vertex of j to a vertex of i. Since each vertex can belong to at most one player, the number of such A_j's that might not be independent of A_i is at most the number of vertices having a path of length 1 or 2 to any vertex of player i. Now the maximum degree of the interaction graph is bounded by d. Thus the number of vertices within a distance 2 from player i's vertices is at most $n.d + n.d.d \leq 2nd^2$. □

We now give a simple way to define and prove a bound on $Pr(A_i)$. This approach only gives a weaker result than Theorem 1. We then use a more powerful technique to prove our main theorem. The correct definition of A_i and the formal proof of Theorem 1 will be given at the end of this section.
For a player i, assume that other players do not change their strategies, the utility U_i is a function on the strategy set $\{0,1\}^{|S_i|}$. As discussed at the beginning of this section, a random strategy gives at least a 2- approximate solution in expectation. Because U_i satisfies a Lipschitz condition, one can use a concentration inequality to bound the probability that the value of U_i is much smaller than the expected value. In particular we can use the following Hoeffding-Azuma inequality [14, 4]:

Lemma 3 (Hoeffding-Azuma). *Let $U : \{0,1\}^n \rightarrow R$ be a function such that $|U(x) - U(x')| \leq \Delta$ whenever x, x' differ in only one coordinate, then if $x_1, .., x_n$ are independent boolean variables, we have : $Pr(|U(x_1, .., x_n) - E(U(x_1, .., x_n))| \geq \lambda\Delta\sqrt{n}) \leq 2e^{\frac{-\lambda^2}{2}}$.* □

If we define A_i to be the event that:

$$U_i(x_1, .., x_{|S_i|}) - E(U_i(x_1, .., x_{|S_i|})) \geq \lambda\Delta\sqrt{|S_i|}$$

with a λ chosen such that

$$(2e^{\frac{-\lambda^2}{2}})e(2d^2n + 1) \leq 1$$

then we can apply the local lemma to prove that there the game has a configuration where each player i get at least $\frac{1}{2}OPT_i - \Delta\lambda\sqrt{|S_i|}$. Now:

$$2e^{\frac{-\lambda^2}{2}}e(2d^2n + 1) \leq 1 \Leftrightarrow 2e^{\frac{\lambda^2}{2}} \geq e(2d^2n + 1)$$

One can choose

$$\lambda = \sqrt{2\log(2e(2d^2n + 1))} = O(\sqrt{\log n + \log d}),$$

and obtain the following result:

Claim. There exists a configuration such that each player obtains a payoff at least $\frac{1}{2}OPT_i - O(\Delta\sqrt{|S_i|(\log n + \log d)})$, where OPT_i is the optimum that that player can obtain assuming other players do not change their strategies. □

The additive error in the result above is of order $O(\sqrt{|S_i|(\log n + \log d)})$, thus when OPT_i is relatively small compred with $\sqrt{|S_i|}$ the result is rather weak. To prove the stronger result in Theorem 1, we will apply a stronger concentration technique developed by Talagrand [20]. We start by describing a class of functions that can be used in this approach.

Definition 1. *We call a non negative function f defined on a set Ω of n dimensional vectors a* c-configuration function *if it has the following property: for each $x \in \Omega$ there is a non-negative unit n-dimensional vector α such that for each $y \in \Omega$ we have:*

$$f(y) > f(x) - \sqrt{cf(x)}d_\alpha(x, y),$$

where

$$d_\alpha(x, y) = \sum_{i \in \{1,..,n\}; x_i \neq y_i} \alpha_i.$$

Lemma 4 (Talagrand). *Let f be a c-configuration function, and let m be a median for $f(X)$, where X is a random variable taking each coordinate independently from any distribution. Then for any $t > 0$*

$$Pr(f(X) \leq m - t) \leq 2e^{-t^2/4cm}.$$

$$Pr(f(X) \geq m + t) \leq 2e^{-t^2/4c(m+t)}.$$

$$Pr(|f(X) - m| \geq t) \leq 2e^{-t^2/4cm} + 2e^{-t^2/4c(m+t)}.$$

□

The Talagrand's inequality exploits the structure of the function f. The concentration error can be given as a function of the value of the median instead of the number of the variables. Therefore, Talagrand's inequality gives a stronger result when the median is relatively small.
We will show later that the class of utility functions considered in this paper are Δ^2-configuration functions, where Δ is the Lipschitz constant of the utilities. We will then apply the Talagrand's inequality. A technical problem here is that, the concentration is however bounded around the *median* of the variable. Some technical work will be needed to express the concentration inequality around the expected value. We will use the following result instead of Lemma 4:

Lemma 5. *Let f be a c-configuration function, and let μ be the expected value for $f(X)$, where X is a random variable taking each coordinate independently from any distribution, then for all $\lambda > 10$*

$$Pr\left(f(X) < \mu - 60\lambda(\sqrt{c\mu} + c)\right) \leq 2e^{-\lambda^2}$$

Proof. Before proving the lemma, let us remark that all the constants are chosen in the calculation for convenience. We did not attempt to optimize them. Now, recall that μ and m is the expected value and median of $f(X)$. We have:

$$|\mu - m| = |E(f(X) - m)| \leq E(|f(X) - m|) = \int_0^\infty Pr(|f(X) - m| > t)dt$$

Because of the Lemma 4:

$$\int_0^\infty Pr(|f(X)-m|>t)dt \le 2\int_0^\infty e^{-t^2/4cm}dt + 2\int_0^\infty e^{-t^2/4c(m+t)}dt$$

$$= 2\sqrt{\pi cm} + 2\int_0^\infty e^{-t^2/4c(m+t)}dt$$

$$\le 2\sqrt{\pi cm} + 2\int_0^m e^{-t^2/4c(m+m)}dt + 2\int_m^\infty e^{-t^2/4c(t+t)}dt$$

$$= 2\sqrt{\pi cm} + 2\int_0^m e^{-t^2/8cm}dt + 2\int_m^\infty e^{-t^2/8ct}dt$$

$$< 2\sqrt{\pi cm} + 2\sqrt{2\pi cm} + 16c \le 10\sqrt{cm} + 16c$$

Thus we have

$$|\mu - m| < 10\sqrt{cm} + 16c \tag{1}$$

From here it is not hard to see that

$$15(\sqrt{\mu} + \sqrt{c}) \ge \sqrt{m} \tag{2}$$

Now, applying the first inequality in Lemma 4 for $t = 2\lambda\sqrt{cm}$, we have:

$$2e^{-\lambda^2} \ge Pr(f(X) < m - 2\lambda\sqrt{cm}) \tag{3}$$

Using the fact that $m \ge \mu - |\mu - m|$, we have:

$$Pr(f(X) < m - 2\lambda\sqrt{cm}) \ge Pr\left(f(X) < \mu - |\mu - m| - 2\lambda\sqrt{cm}\right) \tag{4}$$

Now, due to (1) $|\mu - m| \le 10\sqrt{c\pi m} + 16c$, we have:

$$Pr\left(f(X) < \mu - |\mu - m| - 2\lambda\sqrt{cm}\right) \ge Pr\left(f(X) < \mu - 10\sqrt{cm} - 16c - 2\lambda\sqrt{cm}\right) \tag{5}$$

Combining (3),(4) and (5) we obtain

$$2e^{-\lambda^2} \ge Pr\left(f(X) < \mu - 10\sqrt{cm} - 16c - 2\lambda\sqrt{cm}\right)$$

$$\Leftrightarrow 2e^{-\lambda^2} \ge Pr\left(f(X) < \mu - \sqrt{cm}(10 + 2\lambda) - 16c\right)$$

Using the fact from (2) that $\sqrt{m} < 15(\sqrt{\mu} + \sqrt{c})$ we obtain:

$$2e^{-\lambda^2} \ge Pr\left(f(X) < \mu - 15\sqrt{c}(\sqrt{\mu} + \sqrt{c})(10 + 2\lambda) - 16c\right)$$

Therefore if λ big enough (> 10), we can simplify the last expression to get:

$$2e^{-\lambda^2} \ge Pr\left(f(X) < \mu - 60\lambda(\sqrt{c\mu} + c)\right)$$

This is what we need to prove. $\square$

We are now ready to prove the main theorem:

Proof (Proof of Theorem 1).
Given a MAXSAT utility function U with the Lipschitz constant Δ, we first show that U is a Δ^2-configuration function. And hence we can apply Lemma 5. Given a configuration of the game, let X be the strategy of a player, and k be the value of his payoff. In other words, by playing X, the player can satisfy k clauses in the corresponding Normal conjunctive formula. From each satisfying clause, pick a variable. It is possible that we pick the same variable for different clauses, thus the number of variables picked is $k' \leq k$. Each of these variables corresponds to a decision made on a vertex. We can consider these decisions as a set of "witnesses". The reason is that if the player makes the same decisions as these witnesses, he is guaranteed to obtain a payoff of at least k. Furthermore, because of the Lipschitz condition, if at most l witness-decisions are made differently, then the player obtains a pay of of at least $k - \Delta l$.
Now let $\alpha \in \mathbb{R}^n$ be a non negative vector that takes $\frac{1}{\sqrt{k'}}$ on the coordinates of the witnesses and 0 elsewhere. Let Y be an alternative strategy such that $d_\alpha(X,Y) = \frac{l}{\sqrt{k'}}$. This means that X and Y differs on l witness. Thus the payoff of the player playing Y is at least $k - \Delta l$. This shows that:

$$U(Y) \geq U(X) - \Delta\sqrt{k'}d_\alpha(X,Y) \geq U(X) - \sqrt{\Delta^2 U(X)}d_\alpha(X,Y).$$

The last inequality is from the fact that $U(X) = k \geq k'$. This proves that the utility function is a Δ^2-configuration function.
Let $\mu = E(U(X))$ and $\lambda > 10$ be a parameter, we now apply Lemma 5,

$$Pr\left(U(X) < \mu - 60\lambda(\Delta\sqrt{\mu} + \Delta^2)\right) \leq 2e^{-\lambda^2}$$

We need to choose λ such that $e \cdot 2e^{-\lambda^2} \cdot (2nd^2 + 1) < 1$. Again without optimizing the constant, let us choose $\lambda = 2\sqrt{\log n + \log d} + 10$. And we have:

$$Pr\left(U(X) \leq \mu - 60(2\sqrt{\log n + \log d} + 10)(\Delta\sqrt{\mu} + \Delta^2)\right) \leq \frac{1}{e(2nd^2 + 1)} \tag{6}$$

Note that μ is the expected value of a MAXSAT function, therefore, μ is at least 1/2 the optimal value. Therefore, we define A_i to be the event that the player i gets a payoff less than

$$\begin{aligned}&\frac{OPT_i}{2} - 60(2\sqrt{\log n + \log d} + 10)(\Delta\sqrt{\frac{OPT_i}{2}} + \Delta^2)\\ &= \frac{OPT_i}{2} - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right)\end{aligned}$$

where OPT_i is the maximum payoff that player i can get assuming other players do not change their strategies. Because of (6), we have

$$Pr(A_i) \leq \frac{1}{e(2nd^2 + 1)}.$$

Thus, according to Lovász local lemma, there exists a configuration where none of A_i occur. This is what we need to prove. □

Note that using algorithmic versions of Lovász local lemma for example, [17], we can give a polynomial time algorithm for finding such an approximate equilibrium.

References

1. N. Alon and J.H. Spencer, *The probabilistic method.* Wiley-Interscience Ser. Discrete Math and Optimization, Wiley, New York, 1992.
2. E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler, *Near-Optimal Network Design with Selfish Agents.* Theory of Computing, Volume 4 (2008), pp. 77-109.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden, *The Price of Stability for Network Design with Fair Cost Allocation.* In FOCS 2004
4. K. Azuma, *Weighted sums of certain dependent random variables.* Tokuku Math. J. 19 (1967), 357-367.
5. X. Chen and X. Deng, *Settling the Complexity of 2-Player Nash-Equilibrium.* FOCS 06.
6. C. Daskalakis and C. Papadimitriou, *Computing Nash in graphical games via Markov random field.* EC 2006.
7. C. Daskalakis, A. Mehta and C. Papadimitriou, *Progress in Approximate Nash Equilibria.* EC 2007.
8. C. Daskalakis, P.W. Goldberg and C. Papadimitriou,*The Complexity of Computing a Nash Equilibrium.* STOC 2006.
9. E. Elkind, L.A. Goldenberg and P.W Goldenberg, *Nash Equilibria in Graphical Games on Tree Revisited.* EC 2006.
10. P. Erdős and L. Lovász, *Problems and results on 3-chromatic hypergraphs and some related questions.* Infinite and finite sets (to Paul Erdős on his 60th birthday), Vol II, (A. Hajnal, R. Rado, and V.T. Sós, eds.) North-Holland, Amsterdam, 1975, pp. 609-627.
11. A. Fabrikant, C. Papadimitriou and K. Talwar, *The Complexity of Pure Nash Equilibria.* STOC 2004.
12. Uriel Feige, Vahab S. Mirrokni, Jan Vondrak, *Maximizing Non-Monotone Submodular Functions.* FOCS 2007.
13. J. Hastad, *Some optimal inapproximability results.* Journal of ACM, Vol. 48, 2001, pp 798–859.
14. W. Hoeffding, *Probability inequalities for sums of bounded random variables.* J. Amer. Statist. Assoc. 58, 713-721 (1963)
15. M. Kearns, M. Littman ans S. Singh, *Graphical models for game theory,* In UAI, 2001.
16. M. Kearns, *Graphical Games* in *Algorithmic Game Theory,* (N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, eds), Cambridge University Press, 2007.
17. R. Moser, *A Constructive Proof of the Lovász Local Lemma.* STOC 2009.
18. R.W. Rosenthal, *A class of games possessing pure-strategy Nash equilibria.* International journal of Game Theory, 2:65-67, 1973.
19. A. Skopalik and B. Vöcking, *Inapproximability of Pure Nash Equilibria*, STOC 2008.
20. M. Talagrand, *Concentration of measure and isoperimetric inequalities in product spaces.* Publications Mathématiques de L'IHÉS, 81, 73-205, 1995.
21. A. Vetta, *Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions.* FOCS 2002.