

A Note on Approximate Nash Equilibria

Constantinos Daskalakis¹, Aranyak Mehta², and Christos Papadimitriou¹

¹ University of California, Berkeley, USA*

² IBM Almaden Research Center, San Jose, USA

Abstract. In view of the intractability of finding a Nash equilibrium, it is important to understand the limits of approximation in this context. A subexponential approximation scheme is known [LMM03], and no approximation better than $\frac{1}{4}$ is possible by any algorithm that examines equilibria involving fewer than $\log n$ strategies [Alt94]. We give a simple, linear-time algorithm examining just two strategies per player and resulting in a $\frac{1}{2}$ -approximate Nash equilibrium in any 2-player game. For the more demanding notion of *well-supported approximate equilibrium* due to [DGP06] no nontrivial bound is known; we show that the problem can be reduced to the case of win-lose games (games with all utilities $0 - 1$), and that an approximation of $\frac{5}{6}$ is possible contingent upon a graph-theoretic conjecture.

1 Introduction

Since it was shown that finding a Nash equilibrium is PPAD-complete [DGP06], even for 2-player games [CD05], the question of approximate Nash equilibrium emerged as the central remaining open problem in the area of equilibrium computation. Assume that all utilities have been normalized to be between 0 and 1 (this is a common assumption, since scaling the utilities of a player by any positive factor, and applying any additive constant, results in an equivalent game). A set of mixed strategies is called an ϵ -approximate Nash equilibrium, where $\epsilon > 0$, if for each player all strategies have expected payoff that is at most ϵ more than the expected payoff of the given strategy. Clearly, any mixed strategy combination is a 1-approximate Nash equilibrium, and it is quite straightforward to find a $\frac{3}{4}$ -approximate Nash equilibrium by examining all supports of size two. In fact, [KPS06] provides a scheme that yields, for every $\epsilon > 0$, in time polynomial in the size of the game and $\frac{1}{\epsilon}$, a $\frac{2+\epsilon+\lambda}{4}$ -approximate Nash equilibrium, where λ is the minimum, among all Nash equilibria, expected payoff of either player. In [LMM03] it was shown that, for every $\epsilon > 0$, an ϵ -approximate Nash equilibrium can be found in time $O(n^{\frac{\log n}{\epsilon^2}})$ by examining all supports of size $\frac{\log n}{\epsilon^2}$. It was pointed out in [Alt94] that, even for zero-sum games, no algorithm that examines supports smaller than about $\log n$ can achieve an approximation better than $\frac{1}{4}$. Can this gap between $\frac{1}{4}$ and $\frac{3}{4}$ be bridged by looking at small supports? And how can the barrier of $\frac{1}{4}$ be broken in polynomial time?

* Supported by NSF grant CCF-0515259.

In this note we concentrate on 2-player games. We point out that a straightforward algorithm, looking at just three strategies in total, achieves a $\frac{1}{2}$ -approximate Nash equilibrium. The algorithm is very intuitive: For any strategy i of the row player let j be the best response of the column player, and let k be the best response of the row player to j . Then the row player plays an equal mixture of i and k , while the column player plays j . The proof of $\frac{1}{2}$ -approximation is rather immediate.

We also examine a more sophisticated approximation concept due to [GP06, DGP06], which we call here the *well-supported ϵ -approximate Nash equilibrium*, which does not allow in the support strategies that are suboptimal by at least ϵ . For this concept no approximation constant better than 1 is known. We show that the problem is reduced — albeit with a loss in the approximation ratio — to the case in which all utilities are either zero or one (this is often called the “win-lose case”). We also prove that, assuming a well-studied and plausible graph-theoretic conjecture, in win-lose games there is a well-supported $\frac{2}{3}$ -approximate Nash equilibrium with supports of size at most three (and of course it can be found in polynomial time). This yields a well-supported $\frac{5}{6}$ -approximate Nash equilibrium for any game.

2 Definitions

We consider normal form games between two players, the *row player* and the *column player*, each with n strategies at his disposal. The game is defined by two $n \times n$ payoff matrices, R for the row player, and C for the column player. The pure strategies of the row player correspond to the n rows and the pure strategies of the column player correspond to the n columns. If the row player plays row i and the column player plays column j , then the row player receives a payoff of R_{ij} and the column player gets C_{ij} . Payoffs are extended linearly to pairs of mixed strategies — if the row player plays a probability distribution x over the rows and column player plays a distribution y over the columns, then the row player gets a payoff of $x^T R y$ and the column player gets a payoff of $x^T C y$.

A *Nash equilibrium* in this setting is a pair of mixed strategies, x^* for the row player and y^* for the column player, such that neither player has an incentive to unilaterally defect. Note that, by linearity, the best defection is to a pure strategy. Let e_i denote the vector with a 1 at the i th coordinate and 0 elsewhere. A pair of mixed strategies (x^*, y^*) is a Nash equilibrium if

$$\forall i = 1..n, \quad e_i^T R y^* \leq x^{*T} R y^*$$

$$\forall i = 1..n, \quad x^{*T} C e_i \leq x^{*T} C y^*$$

It can be easily shown that every pair of equilibrium strategies of a game does not change upon multiplying all the entries of a payoff matrix by a constant, and upon adding the same constant to each entry. We shall therefore assume that the entries of both payoff matrices R and C are between 0 and 1.

For $\epsilon > 0$, we define an ϵ -approximate Nash equilibrium to be a pair of mixed strategies x^* for the row player and y^* for the column player, so that the incentive to unilaterally deviate is at most ϵ :

$$\forall i = 1..n, \quad e_i^T R y^* \leq x^{*T} R y^* + \epsilon$$

$$\forall i = 1..n, \quad x^{*T} C e_i \leq x^{*T} C y^* + \epsilon$$

A stronger notion of approximately equilibrium strategies was introduced in [GP06, DGP06]: For $\epsilon > 0$, a *well-supported ϵ -approximate Nash equilibrium*, or an ϵ -*well-supported Nash equilibrium*, is a pair of mixed strategies, x^* for the row player and y^* for the column player, so that a player plays only approximately best-response pure strategies with non-zero probability:

$$\forall i : x_i^* > 0 \Rightarrow e_i^T R y^* \geq e_j^T R y^* - \epsilon, \quad \forall j$$

$$\forall i : y_i^* > 0 \Rightarrow x^{*T} C e_i \geq x^{*T} C e_j - \epsilon, \quad \forall j$$

If only the first set of inequalities holds, we say that every pure strategy in the support of x^* is ϵ -well supported against y^* , and similarly for the second. This is indeed a stronger definition, in the sense that every ϵ -well supported Nash equilibrium is also an ϵ -approximate Nash equilibrium, but the converse need not be true. However, the following lemma from [CDT06] shows that there does exist a polynomial relationship between the two:

Lemma 1. [CDT06] *For every 2 player normal form game, for every $\epsilon > 0$, given an $\frac{\epsilon}{8n}$ -approximate equilibrium we can compute in polynomial time an ϵ -well-supported equilibrium.*

Since in this paper we are interested in constant ϵ , this lemma is of no help to us; indeed, our results for approximate equilibria are stronger and simpler than those for well-supported equilibria.

3 A Simple Algorithm

We provide here a simple way of computing a $\frac{1}{2}$ -approximate Nash equilibrium: Pick an arbitrary row for the row player, say row i . Let $j = \arg \max_{j'} C_{ij'}$. Let $k = \arg \max_{k'} R_{k'j}$. Thus, j is a best-response column for the column player to the row i , and k is a best-response row for the row player to the column j .

The equilibrium is $x^* = \frac{1}{2}e_i + \frac{1}{2}e_k$ and $y^* = e_j$, i.e., the row player plays row i or row k with probability $\frac{1}{2}$ each, while the column player plays column j with probability 1.

Theorem 1. *The strategy pair (x^*, y^*) is a $\frac{1}{2}$ -approximate Nash equilibrium.*

Proof. The row player’s payoff under (x^*, y^*) is $x^{*T} R y^* = \frac{1}{2}R_{ij} + \frac{1}{2}R_{kj}$. By construction, one of his best responses to y^* is to play the pure strategy on row k , which gives a payoff of R_{kj} . Hence his incentive to defect is equal to the difference:

$$R_{kj} - (\frac{1}{2}R_{ij} + \frac{1}{2}R_{kj}) = \frac{1}{2}R_{kj} - \frac{1}{2}R_{ij} \leq \frac{1}{2}R_{kj} \leq \frac{1}{2}$$

The column player’s payoff under (x^*, y^*) is $x^{*T}Cy^* = \frac{1}{2}C_{ij} + \frac{1}{2}C_{kj}$. Let j' be a best pure strategy response of the column player to x^* : this strategy gives the column player a value of $\frac{1}{2}C_{ij'} + \frac{1}{2}C_{kj'}$, hence his incentive to defect is equal to the difference:

$$\begin{aligned} (\frac{1}{2}C_{ij'} + \frac{1}{2}C_{kj'}) - (\frac{1}{2}C_{ij} + \frac{1}{2}C_{kj}) &= \frac{1}{2}(C_{ij'} - C_{ij}) + \frac{1}{2}(C_{kj'} - C_{kj}) \\ &\leq 0 + \frac{1}{2}(C_{kj'} - C_{kj}) \\ &\leq \frac{1}{2} \end{aligned} \tag{1}$$

Here the first inequality follows from the fact that column j was a best response to row i , by the first step of the construction.

4 Well Supported Nash Equilibria

The algorithm of the previous section yields equilibria that are, in the worst case, as bad as 1-well supported. In this section we address the harder problem of finding ϵ -well supported equilibria for some $\epsilon < 1$.

Our construction has two components. In the first we transform the given 2-player game into a new game by rearranging and potentially discarding or duplicating some of the columns of the original game. The transformation will be such that well supported equilibria in the new game can be mapped back to well supported equilibria of the original game; moreover, the mapping will result in some sort of *decorrelation of the players*, in the sense that computation of well supported equilibria in the decorrelated game can be carried out by looking at the row player. The second part of the construction relies in mapping the original game into a win-lose game (game with 0 – 1 payoffs) and computing equilibria on the latter. The mapping will guarantee that well supportedness of equilibria is preserved, albeit with some larger ϵ .

4.1 Player Decorrelation

Let (R, C) be a 2-player game, where the set of strategies of both players is $[n]$.

Definition 1. A mapping $f : [n] \rightarrow [n]$ is a best response mapping for the column player iff, for every $i \in [n]$,

$$C_{if(i)} = \max_j C_{ij}.$$

Definition 2 (Decorrelation Transformation). The decorrelated game (R^f, C^f) corresponding to the best response mapping f is defined as follows

$$\begin{aligned} \forall i, j \in [n] : \quad R_{ij}^f &= R_{if(j)} \\ C_{ij}^f &= C_{if(j)} \end{aligned}$$

Note that the decorrelation transformation need not be a permutation of the columns of the original game. Some columns of the original game may very well

be dropped and others duplicated. So, it is not true in general that (exact) Nash equilibria of the decorrelated game can be mapped into Nash equilibria of the original game. However, some specially structured well supported equilibria of the decorrelated game can be mapped into well supported equilibria of the original game as we explore in the following lemmas.

In the following discussion, we assume that we have fixed a best response mapping f for the column player and that the corresponding decorrelated game is (R^f, C^f) . Also, if $S \subseteq [n]$, we denote by $\Delta(S)$ the set of probability distributions over the set S . Moreover, if $x \in \Delta(S)$, we denote by $\text{supp}(x) \triangleq \{i \in [n] | x(i) > 0\}$, the support of x .

Lemma 2. *In the game (R^f, C^f) , for all sets $S \subseteq [n]$, the strategies of the column player in S are $\frac{|S|-1}{|S|}$ -well supported against the strategy x^* of the row player, where x^* is defined in terms of the set $S' = \{i \in S | C_{ii}^f = 0\}$ as follows*

- if $S' \neq \emptyset$, then x^* is uniform over the set S'
- if $S' = \emptyset$, then

$$x^*(i) = \begin{cases} \frac{1}{Z} \frac{1}{C_{ii}^f}, & \text{if } i \in S \\ 0, & \text{otherwise} \end{cases}$$

where $Z = \sum_{i \in S} \frac{1}{C_{ii}^f}$ is a normalizing constant.

Proof. Suppose that $S' \neq \emptyset$. By the definition of S' , it follows that $\forall i \in S', j \in [n], C_{ij}^f = 0$. Therefore, for all $j \in [n]$,

$$x^{*T} C^f e_j = 0,$$

which proves the claim.

The proof of the $S' = \emptyset$ case is based on the following observations.

- $\forall j \in S$:

$$x^{*T} C^f e_j \geq x^*(j) C_{jj}^f = \frac{1}{Z} \frac{1}{C_{jj}^f} C_{jj}^f = \frac{1}{Z}$$

- $\forall j \in [n]$:

$$x^{*T} C^f e_j = \frac{1}{Z} \sum_{i \in S} \frac{1}{C_{ii}^f} C_{ij}^f \leq \frac{1}{Z} \sum_{i \in S} \frac{1}{C_{ii}^f} C_{ii}^f = \frac{|S|}{Z}$$

- $Z = \sum_{i \in S} \frac{1}{C_{ii}^f} \geq |S|$ because every entry of C^f is at most 1.

Therefore, $\forall j_1 \in S, j_2 \in [n]$,

$$x^{*T} C^f e_{j_2} - x^{*T} C^f e_{j_1} \leq \frac{|S|-1}{Z} \leq \frac{|S|-1}{|S|},$$

which completes the claim.

The following lemma is an immediate corollary of Lemma 2.

Lemma 3 (Player Decorrelation). *In the game (R^f, C^f) , if there exists a set $S \subseteq [n]$ and a mixed strategy $y \in \Delta(S)$ for the column player such that the strategies in S are $\frac{|S|-1}{|S|}$ -well supported for the row player against the distribution y , then there exists a strategy $x \in \Delta(S)$ for the row player so that the pair (x, y) is an $\frac{|S|-1}{|S|}$ -well supported Nash equilibrium.*

The next lemma describes how well supported equilibria in the games (R, C) and (R^f, C^f) are related.

Lemma 4. *$\forall S \subseteq [n]$, if the pair (x^*, y^*) , where x^* is defined as in the statement of Lemma 2 and y^* is the uniform distribution over S , constitutes an $\frac{|S|-1}{|S|}$ -well supported Nash equilibrium for the game (R^f, C^f) , then the pair of distributions (x^*, y') is an $\frac{|S|-1}{|S|}$ well supported Nash equilibrium for the game (R, C) , where y' is the distribution defined as follows*

$$y'(i) = \sum_{j \in S} y^*(j) \mathcal{X}_{f(j)=i}, \forall i \in [n],$$

where $\mathcal{X}_{f(j)=i}$ is the indicator function of the condition “ $f(j) = i$ ”.

Proof. We have to verify that the pair of distributions (x^*, y') satisfies the conditions of well supportedness for the row and column player in the game (R, C) .

Row Player: We show that the strategies y^* and y' of the column player give to every pure strategy of the row player the same payoff in the two games. And, since the support of the row player stays the same set S in the two games, the fact that the strategy of the row player is well supported in the game (R^f, C^f) guarantees that the strategy of the row player will be well supported in the game (R, C) as well.

$$\begin{aligned} \forall i \in [n] : \quad e_i^T R y' &= \sum_{k=1}^n R_{ik} \cdot y'(k) \\ &= \sum_{k=1}^n R_{ik} \cdot \sum_{j \in S} y^*(j) \mathcal{X}_{f(j)=k} \\ &= \sum_{j \in S} y^*(j) \sum_{k=1}^n R_{ik} \cdot \mathcal{X}_{f(j)=k} \\ &= \sum_{j \in S} y^*(j) R_{i f(j)} \\ &= \sum_{j \in S} y^*(j) R_{ij}^f = e_i^T R^f y^* \end{aligned}$$

Column Player: As in the proof of Lemma 2, the analysis proceeds by distinguishing the cases $S' \neq \emptyset$ and $S' = \emptyset$. The case $S' \neq \emptyset$ is easy, because, in

this regime, it must hold that $\forall i \in S', j \in [n], C_{ij}^f = 0$ which implies that, also, $C_{ij} = 0, \forall i \in S', j \in [n]$. And, since, the row player plays the same distribution as in the proof of Lemma 2, we can use the arguments applied there.

So it is enough to deal with the $S' = \emptyset$ case. The support of y' is clearly the set $S'' = \{j | \exists i \in S \text{ such that } f(i) = j\}$. Moreover, observe the following:

$$\begin{aligned} \forall j \in S'' : \quad x^{*T} C e_j &= \sum_{i \in S} x^*(i) C_{ij} \\ &\geq \sum_{\substack{i \in S \text{ s.t.} \\ f(i)=j}} x^*(i) C_{if(i)} \\ &= \sum_{\substack{i \in S \text{ s.t.} \\ f(i)=j}} x^*(i) C_{ii}^f \geq \frac{1}{Z} \end{aligned}$$

The final inequality holds because there is at least one summand, since $j \in S''$. On the other hand,

$$\begin{aligned} \forall j \notin S'' : \quad x^{*T} C e_j &= \sum_{i \in S} x^*(i) C_{ij} \\ &\leq \sum_{i \in S} x^*(i) C_{if(i)} \\ &= \sum_{i \in S} x^*(i) C_{ii}^f \\ &= \sum_{i \in S} \frac{1}{Z} \frac{1}{C_{ii}^f} C_{ii}^f \\ &= \frac{|S|}{Z} \end{aligned}$$

Moreover, as we argued in the proof of Lemma 2, $\frac{|S|}{Z} - \frac{1}{Z} \leq \frac{|S|-1}{|S|}$. This completes the proof, since the strategy of the column player is, thus, also well supported.

4.2 Reduction to Win-Lose Games

We now describe a mapping from a general 2-player game to a win-lose game so that well supported equilibria of the win-lose game can be mapped to well supported equilibria of the original game. A bit of notation first. If A is an $n \times n$ matrix with entries in $[0, 1]$, we denote by $round(A)$ the 0-1 matrix defined as follows, for all $i, j \in [n]$,

$$round(A)_{ij} = \begin{cases} 1, & \text{if } A_{ij} \geq \frac{1}{2} \\ 0, & \text{if } A_{ij} < \frac{1}{2} \end{cases}$$

The following lemma establishes a useful connection between well supported equilibria of the 0-1 game and those of the original game.

Lemma 5. *If (x, y) is an ϵ -well supported Nash equilibrium of the game $(\text{round}(R), \text{round}(C))$, then (x, y) is a $\frac{1+\epsilon}{2}$ -well supported Nash equilibrium of the game (R, C) .*

Proof. We will show that the strategy of the row player in game (R, C) is well supported; similar arguments apply to the second player. Denote $R' = \text{round}(R)$ and $C' = \text{round}(C)$.

The following claim follows easily from the rounding procedure.

Claim. $\forall i, j \in [n] : \frac{R'_{ij}}{2} \leq R_{ij} \leq \frac{1}{2} + \frac{R'_{ij}}{2}$

Therefore, it follows that, $\forall i \in [n]$,

$$\frac{1}{2}e_i^T R' y \leq e_i^T R y \leq \frac{1}{2} + \frac{1}{2}e_i^T R' y \tag{2}$$

We will use (2) to argue that the row player is well supported. Indeed, $\forall j \in \text{supp}(x)$, and $\forall i \in [n]$

$$e_i^T R y - e_j^T R y \leq \frac{1}{2} + \frac{1}{2}e_i^T R' y - \frac{1}{2}e_j^T R' y \leq \frac{1}{2} + \frac{1}{2} \cdot (e_i^T R' y - e_j^T R' y) \leq \frac{1}{2} + \frac{1}{2} \cdot \epsilon$$

where the last implication follows from the fact that (x, y) is an ϵ -well supported Nash equilibrium.

4.3 Finding Well Supported Equilibria

Lemmas 2 through 5 suggest the following algorithm, *ALG-WS*, to compute well supported Nash equilibria for a given two player game (R, C) :

1. Map game (R, C) to the win-lose game $(\text{round}(R), \text{round}(C))$.
2. Map game $(\text{round}(R), \text{round}(C))$ to the game $(\text{round}(R)^f, \text{round}(C)^f)$, where f is any best response mapping for the column player.
3. Find a subset $S \subseteq [n]$ and a strategy $y \in \Delta(S)$ for the column player such that all the strategies in S are $\frac{|S|-1}{|S|}$ well supported for the row player in $(\text{round}(R)^f, \text{round}(C)^f)$ against the strategy y for the column player.
4. By a successive application of lemmas 3, 4 and 5, get an $\frac{1}{2} + \frac{1}{2} \frac{|S|-1}{|S|} = 1 - \frac{1}{2|S|}$ well supported Nash equilibrium of the original game.

The only non-trivial step of the algorithm is step 3. Let us paraphrase what this task entails:

“Given a 0-1 matrix $\text{round}(R)^f$, find a subset of the columns $S \subset [n]$ and a distribution $y \in \Delta(S)$, so that all rows in S are $\frac{|S|-1}{|S|}$ well supported against the distribution y over the columns.”

It is useful to consider the $0-1$ matrix $\text{round}(R)^f$ as the adjacency matrix of a directed graph G on n vertices. We shall argue next that the task above is easy in two cases: When G has a small cycle, and when G has a small *undominated* set of vertices, that is, a set of vertices such that no other vertex has edges to all of them.

1. Suppose first that G has a cycle of length k , and let S be the vertices on the cycle. Then it is easy to see that all the k strategies in S are $\frac{k-1}{k}$ -well supported for the row player against y , where y is the uniform strategy for the column player over the set S . The reason is that each strategy in S has expected payoff $\frac{1}{k}$ against y , and thus no other strategy can dominate it by more than $\frac{k-1}{k}$. This, via the above algorithm, implies a $(1 - \frac{1}{2k})$ -well-supported Nash equilibrium.
2. Second, suppose that there is a set S of ℓ undominated vertices. Then every strategy in S is $(1 - \frac{1}{\ell})$ -well supported for the row player against the uniform strategy y of the column player on S , simply because there is no row that has payoff better than $1 - \frac{1}{\ell}$ against y . Again, via the algorithm, this implies that we can find a $(1 - \frac{1}{2\ell})$ -well-supported Nash equilibrium.

This leads us to the following graph theoretic conjecture:

Conjecture 1. There are integers k and ℓ such that every digraph either has a cycle of length at most k or an undominated set of ℓ vertices.

Now, the next result follows immediately from the preceding discussion:

Theorem 2. *If Conjecture 1 is true for some values of k and ℓ , then Algorithm ALG-WS returns in polynomial time (e.g. by exhaustive search) a $\max\{1 - \frac{1}{2k}, 1 - \frac{1}{2\ell}\}$ -well-supported Nash equilibrium which has support of size $\max\{k, \ell\}$.*

The statement of the conjecture is false for $k = \ell = 2$, as can be seen by a small construction (7 nodes). The statement for $k = 3, \ell = 2$ is already non-trivial. In fact, it was stated as a conjecture by Myers [Mye03] in relation to solving a special case of the Caccetta-Häggkvist Conjecture [CH78]. Moreover, it has recently been proved incorrect in [Cha05] via an involved construction. The case of a constant bound on k for $\ell = 2$ has been left open.

While stronger forms of Conjecture 1 seem to be related to well-known and difficult graph theoretic conjectures, we believe that the conjecture itself is true, and even that it holds for some small values of k and ℓ , such as $k = \ell = 3$.

What we can prove is the case of $\ell = \log n$ by showing that every digraph has a set of $\log n$ undominated vertices. This gives a $(1 - \frac{1}{2 \log n})$ -well-supported equilibrium, which does not seem to be easily obtained via other arguments. We can also prove that the statement is true for $k = 3, \ell = 1$ in the special case of tournament graphs; this easily follows from the fact that every tournament is either transitive or contains a directed triangle.

5 Open Problems

Several open problems remain: Can we achieve a better than $1/2$ -approximate Nash equilibrium using constant sized supports, and if so, what is the limit, in view of the lower bound of $1/4$ -approximate equilibria [Alt94]? If constant supports do not suffice, then can we extend our techniques for larger supports? One attempt would be a natural extension of our simple algorithm from Section 3: Continue the iterations for a larger number of steps - in every step, if there is a good defection for either player, then give that pure strategy a non-zero probability and include it in the support. A second idea is to run the Lemke-Howson algorithm for some polynomial number of steps and return the best pair of strategies (note that our algorithm may be interpreted as running three steps of the Lemke-Howson algorithm with an extension or truncation of the last step). Towards this, we have the following result: Recall that an imitation game is a two-player game in which the column player's matrix is the identity matrix. We can show that we can find a $\frac{1}{4}$ -approximate Nash equilibrium in an imitation game by running the Lemke-Howson method for 6 steps. As a final question, can we find in polynomial time a constant well-supported equilibrium, either by proving our conjecture, or independent of it?

References

- [Alt94] I. Althofer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199, 1994.
- [CH78] L. Caccetta and R. Haggkvist. On minimal digraphs with given girth. *Congressus Numerantium*, XXI, 1978.
- [Cha05] P. Charbit. Circuits in graphs and digraphs via embeddings. *Doctoral dissertation*, University of Lyon I, 2005.
- [CD05] Xi Chen and Xiaotie Deng. Settling the complexity of two-player nash equilibrium. *FOCS*, 2006.
- [CDT06] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing nash equilibria: approximation and smoothed complexity. *ECCC*, 2006.
- [DGP06] Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a nash equilibrium. *STOC*, 2006.
- [GP06] Paul Goldberg and Christos Papadimitriou. Reducibility among equilibrium problems. *STOC*, 2006.
- [LMM03] Richard Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. *ACM Electronic Commerce*, 2003.
- [Mye03] J.S. Myers. Extremal theory of graph minors and directed graphs. *Doctoral dissertation*, 2003.
- [KPS06] Spyros C. Kontogiannis, Panagiota N. Panagopoulou, and Paul G. Spirakis. *Workshop on Internet and Network Economics, WINE*, 2006.