



Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2005-010
MIT-LCS-TM-648

February 8, 2005

Complexity of finding Nash equilibria in
0-1 bimatrix games

Tim Abbott, Daniel Kane, and Paul Valiant

Finding Nash equilibria in 0-1 bimatrix games is as hard as finding equilibria in rational bimatrix games

Tim Abbott, Daniel Kane, Paul Valiant*

February 8, 2005

Abstract

We exhibit a polynomial reduction from the problem of finding a Nash equilibrium of a bimatrix game with rational coefficients to the problem of finding a Nash equilibrium of a bimatrix game with 0-1 coefficients.

1 Introduction

We study the problem of finding Nash equilibria in two-player matrix games, a problem that has found increasing applications in both economics and the internet. It has been shown that it is #P-hard to count all Nash equilibria of a two-player game, even if all the matrix entries are 0 or 1 [CS1]. However, the complexity of finding a Nash equilibrium is wide open, and has been proposed as one of the most important open problems in complexity theory today [Pap1].

We give a polynomial reduction from finding Nash equilibria in general bimatrix games to finding Nash equilibria in games where all payoffs are either 0 or 1, resolving an open problem posed in [CS2].

2 Definitions and General Lemmas

Definition 1 (Game) *A bimatrix game is a two-player game defined by a pair (R, C) where R and C are $m \times n$ matrices. R and C are the payoff matrices for the row and column players, respectively. When the game is played, the row player picks a row i to play and the column player picks a column j to play, and each player gets a payoff equal to the element (i, j) of his payoff matrix.*

The goal of the game is to maximize one's expected payoff.

Definition 2 *A pure strategy for the row or column player is a row or column index of the payoff matrix, respectively. A (mixed) strategy is a probability distribution over the pure strategies, denoted by a vector x . To be a probability distribution, each entry must be in $[0, 1]$ and their sum must be 1. The support of a strategy $\text{Supp}(x)$ is the subset of the pure strategies which the player sometimes plays.*

Note that the (expected) payoff for a player with payoff matrix M if the row player is playing strategy x and the column player is playing strategy y is $x^T M y$.

Given such a game, a natural question that arises is what a “rational” player should do. The notion of rationality that has become widely accepted and almost ubiquitous is that of mutual *best responses*. The concept is that a player should play a strategy that maximizes his payoff, given what the other player is playing.

*MIT Computer Science and Artificial Intelligence Laboratory, The Stata Center, Building 32, 32 Vassar Street, Cambridge, MA 02139. Email: pvaliant@mit.edu.

Definition 3 A strategy x is a best response (for the row player) to a strategy y if for every strategy x' , $x'^T R y \leq x^T R y$.

Given such a game, one might ask what the outcome might reasonably be when both the row and column players are “rational”. In this situation, both players will be playing the best response to the other’s strategy. This is exactly the notion that the pair of row and column strategies are in a *Nash equilibrium*.

Definition 4 A Nash equilibrium is a pair of (mixed) strategies x^* and y^* such that each player’s strategy is a best response to the other’s, i.e. for any strategies x, y ,

$$x^T R y^* \leq x^{*T} R y^*, \text{ and } x^{*T} C y \leq x^{*T} C y^*$$

Definition 5 A zero-sum game is a game in which $R + C = 0$, where by 0 we mean the $m \times n$ zero matrix. More generally, a constant-sum game is one where $R + C$ is an $m \times n$ matrix all of whose entries are the same. We also will consider the notions of a $\{0, 1\}$ -game, where all payoffs are either zero or one.

Note that constant-sum games are equivalent to zero-sum games, since subtracting the constant from all of the entries of one player’s payoff matrix M gives a zero-sum game. We apply the following lemma:

Lemma 6 The best responses for a player with payoff matrix M are the same as those if we replace the matrix by $z(M - \mathbb{1}r)$, where $\mathbb{1}$ is the $m \times n$ matrix with all entries equal to one, for any $r, z \in \mathbb{R}, z > 0$.

Proof: We have that

$$x^T z(M - r\mathbb{1})y = z(x^T M y - r x^T \mathbb{1}y) = z(x^T M y - r)$$

so that the payoffs are simply shifted by a constant and then scaled by another constant, which by linearity does not affect the notion of a best response. \blacksquare

Lemma 7 x is a best response to y if and only if for every pure strategy e_i ,

$$e_i^T R y \leq x^T R y$$

Proof: Suppose every pure strategy e_i satisfied

$$e_i^T R y \leq x^T R y.$$

Then by linearity of matrix multiplication, for an arbitrary $x' = \sum_{i=1}^m x'_i e_i$, we have

$$x'^T R y = \left(\sum_{i=1}^m x'_i e_i \right)^T R y = \sum_{i=1}^m x'_i (e_i^T R y) \leq \sum_{i=1}^m x'_i (x^T R y) = \left(\sum_{i=1}^m x'_i \right) (x^T R y) = x'^T R y$$

which implies that x is a best response. The converse is obvious, since any pure strategy is a strategy. \blacksquare

The following corollary simplifies the process of testing whether a pair of strategies is a Nash equilibrium.

Corollary 8 (x^*, y^*) is a Nash equilibrium if and only if neither player has a pure strategy that gives that player a higher payoff.

In this paper, we consider the following problem:

Given a bimatrix game G , find a Nash equilibrium

which defines the complexity class NASH. We assume here that the elements of the matrices that define the game G are represented as rational numbers in the canonical way. A related complexity class is the class $\text{NASH}_{\{0,1\}}$, which is the subset of the class NASH where the matrices are restricted to have values in $\{0, 1\}$.

The main result of this note is that there is a polynomial reduction from NASH to $\text{NASH}_{\{0,1\}}$.

3 Reduction to Mimicking Games

The first step of our reduction is the well-known reduction from general games (R, C) to *mimicking* games (M, I) , where I is an identity matrix. These games are called “mimicking” because the payoff of the second player is non-zero iff she plays the same move as the first player. This implies the following simple lemma.

Lemma 9 *In any Nash equilibrium (x^*, y^*) of the mimicking game (M, I) , $\text{Supp}(y^*) \subset \text{Supp}(x^*)$.*

Proof: Suppose for the sake of contradiction that the second player sometimes plays a strategy not in the support of x^* . The second player will get 0 payoff in this case, and could profitably change her play to mimic some strategy in x^* . Thus (x^*, y^*) is not a Nash equilibrium, the desired contradiction. ■

The next lemma trivially implies the reduction to mimicking games.

Lemma 10 [CS2] *For any bimatrix game (R, C) with R and C $m \times n$ matrices, there is a mimicking game (M, I) with $M = \begin{pmatrix} 0 & C^T \\ R & 0 \end{pmatrix}$ and I the $(m+n) \times (m+n)$ identity matrix such that the Nash equilibria of (R, C) correspond exactly to the Nash equilibrium strategies of the column player in the game (M, I) .*

Proof: (sketch, see [CS2] for details) Given an equilibrium (x^*, y^*) of the game (M, I) , we decompose y^* as $y^* = (c^*, r^*)$ where r^* and c^* have m and n entries respectively. We then note that the condition that (x^*, y^*) is an equilibrium of (M, I) trivially implies the conditions that make (r^*, c^*) an equilibrium of (R, C) , after we scale r^* and c^* to have sum 1.

Similarly, if (r^*, c^*) is an equilibrium of (R, C) , we may scale the two vectors to vectors $(\alpha r^*, \beta c^*)$ with $\alpha, \beta > 0$ so that the scaled incentives are equal:

$$\alpha \max_i (C r^{*T})_i = \beta \max_i (R c^{*T})_i,$$

and so that $(\alpha r^*, \beta c^*)$ sums to 1. Let $y^* = (\alpha r^*, \beta c^*)$ and let x^* be uniform on the support of y^* . It is straightforward to show that (x^*, y^*) is a Nash equilibrium of the game (M, I) . ■

The above lemma implies that NASH is equivalent to the mimicking-game version of NASH.

4 Reduction to $\{0, 1\}$ Games

We now describe the reduction to $\{0, 1\}$ games. The first ingredient of our construction is a method of representing each rational entry of a payoff matrix using only zeros and ones.

We first note that given an $n \times n$ game (M, I) , by Lemma 6 we can scale M by any positive constant, or add any number to its entries without changing the Nash equilibrium strategies. Thus we may put all the rational entries in M under a common denominator to produce a new matrix M' whose entries are integers. We further note that while each entry may now take more bits to express, the number of new bits that are needed per entry is at most the number of bits in the common denominator, which is at most the number of bits needed to express M . Thus the total number of bits has increased at most quadratically.

We now have a game (M', I) where each entry is integral. Our strategy from here is as follows: we will express each entry in binary by replacing the $n \times n$ matrix M' with an $kn \times kn$ block matrix M'' , for some k greater than the binary length of any entry in M' . We then encode each entry of M' into a binary string, and place it in the corresponding $k \times k$ block of M'' , using the rest of the entries to ensure that the binary string is correctly “interpreted” as representing an integer payoff.

We now describe the part of the construction that enforces this interpretation: a $(k-1) \times (k-1)$ matrix G with the property that the game (G, I) , for a $(k-1) \times (k-1)$ identity matrix I , has a *unique* Nash equilibrium (r^*, c^*) , and furthermore, there exist $\frac{k-1}{6}$ columns whose probabilities of being played are in the ratio

$$1 : 2 : 4 : \dots : 2^{\frac{k-1}{6}}.$$

This property is proven in the following lemma.

Lemma 11 *Define matrices A, B as*

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}.$$

For $j = \frac{k-1}{6}$ define the $\frac{k-1}{2} \times \frac{k-1}{2}$ matrix G' to have the following $j \times j$ block form:

$$G' = \begin{pmatrix} A & A & \cdots & A & B \\ A & A & \cdots & B & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ A & B & \cdots & 0 & 0 \\ B & 0 & \cdots & 0 & 0 \end{pmatrix}.$$

Explicitly, G' has block B on the minor diagonal, block A above, and 0 below. Further, let $G = \begin{pmatrix} 0 & 1 - G'^T \\ G' & 0 \end{pmatrix}$.

Then the game (G, I) (equivalent to the constant-sum game $(G', 1 - G')$) has a unique equilibrium (r^*, c^*) with

$$r^* = \frac{1}{6j}(1, 1, \dots, 1)$$

and

$$c^* = \left(\frac{2 \cdot 2^j - 3}{3 \cdot 2^j - 3} v, \frac{2^j}{3 \cdot 2^j - 3} v \right),$$

where $v = \frac{1}{3 \cdot 2^j - 3}(2^{j-1}, 2^{j-1}, 2^{j-1}, 2^{j-2}, 2^{j-2}, 2^{j-2}, \dots, 2, 2, 2, 1, 1, 1)$.

Proof: Consider the game $(G', 1 - G')$. This will have a Nash equilibrium (x, y) with full support only if all entries of $G'y^T$ and all entries of $x(1 - G')$ are equal.

We show by induction that these constraints are satisfied iff both x and y are scaled versions of v . Suppose we know the first $3i$ entries of y are in the proportions of v . Then in the $i + 1$ st block row of G' , we have three rows, which must all have equal incentives for the row player since the row player's strategy has full support, by hypothesis. For each of those rows, the incentives from the A -blocks are the same by the inductive hypotheses, so we must have that the payoffs from the B -block are also the same, i.e.

$$y_{3i+1} + y_{3i+2} = y_{3i+1} + y_{3i+3} = y_{3i+2} + y_{3i+3},$$

which implies $y_{3i+1} = y_{3i+2} = y_{3i+3}$. To show the ratio of 2 : 1 between adjacent blocks, we note that the payoff for this block of rows must be the same as that for the previous one; the payoffs of the two blocks differ by $2y_{3i} - (y_{3i} + 2y_{3i+3})$; setting that to zero proves the inductive step. Thus we have shown that the entries of y are in the ratio specified by v .

Clearly, the same argument applies to x . Thus since both x and y must sum to 1, $x = y = v$ is the only Nash equilibrium with full support.

We now note that the game $(G', 1 - G')$ is in fact a constant-sum game, so its Nash equilibria are the solutions to a linear program. This implies that the set of Nash equilibria is convex. If we suppose for the sake of contradiction that there is another Nash equilibrium in addition to the one $x = y = v$, then all linear

combinations of these two equilibria must also be equilibria, and hence by standard topology arguments there are a continuum of full support equilibria, which contradicts the uniqueness argument of the previous paragraph.

Thus $x = y = v$ is the unique equilibrium for the game $(G', 1 - G')$, and by lemma 10 the game (G, I) has the unique equilibrium described above. ■

Thus we can express integers t in our game by representing them in binary as $(t_{j-1} \dots t_1 t_0)$, and putting digit t_i in column $3(j - i)$ of a matrix based on G .

We now show how to embed the matrix G in larger games so as to allow the binary representation described at the beginning of this section.

Given a game (M', I) , with M' an $n \times n$ matrix with integral entries, construct the $kn \times kn$ $\{0, 1\}$ -matrix M'' as follows. Construct the $k - 1 \times k - 1$ matrix G defined in the above lemma, and append a column of $k - 1$ ones to the right end to create a $k - 1 \times k$ matrix \bar{G} . Place this matrix along the main diagonal of M'' , (with the upper-left corner on the diagonal) filling the rest of these $(k - 1)n$ rows with zeros. Note that this leaves n rows unaccounted for. Since M'' may be considered as a $n \times n$ block matrix, we fill in each block's k unspecified entries with the binary expression of the corresponding entry in the $n \times n$ matrix M' , putting the i th digit of entry $M'_{r,s}$ in $M''_{kr, k(s-1)+3(j-i)}$ as described above. We make the slight modification of making first two entries in these rows one, $M''_{kr, k(s-1)+1} = M''_{kr, k(s-1)+2} = 1$, so that the number represented is actually $M'_{r,s} + 2^j$. The rest of the entries are 0.

We prove the following lemma, which implies our main result.

Lemma 12 *Given any game (M, I) with M rational, construct the game (M'', I') by first rescaling M to an integral matrix M' , and then shifting M' so as to make its entries integers in the range $0 < M'_{r,s} < 2^j$ for some j . Let $k = 6j + 1$, and construct the matrix M'' as above. Then, up to scaling, the Nash equilibrium strategies for the column player of the game (M, I) are identical to the elements $(1, k + 1, \dots, (n - 1)k + 1)$ of the Nash equilibrium strategies for the column player of the $\{0, 1\}$ -game (M'', I') .*

We note that since the size of M'' is polynomial in the number of bits used to express M , and lemma 10 proves that a mimicking game $M = \begin{pmatrix} 0 & C^T \\ R & 0 \end{pmatrix}$ may be constructed so that the Nash equilibrium strategies for its column player correspond exactly to the Nash equilibrium strategies for both players in the game (R, C) , this lemma trivially implies the following theorem.

Theorem 13 $NASH = NASH_{\{0,1\}}$.

We now prove the lemma.

Proof: Consider any Nash equilibrium (x', y') of the game (M'', I') . Motivated by the block decomposition of M'' , we consider y' in blocks of k . Recall from the construction of M'' that each occurrence of \bar{G} stands alone in its corresponding rows, and that the corresponding entries in I' form a $k - 1 \times k - 1$ identity matrix with a column of zeros added. A straightforward application of the definition of a Nash equilibrium reveals that the $k - 1$ corresponding weights in y' are either all zero, or are a scaled Nash equilibrium for the game (G, I) . Thus these $k - 1$ weights $\{y'_{(r-1)k+i}\}_{i=1}^{k-1}$ equal $y''_r(2^{j-1}, 2^{j-1}, 2^{j-1}, \dots, 1, 1, 1, \dots)$ for some $y''_r \geq 0$ as shown in lemma 11. Thus we have found a block representation for any Nash equilibrium strategy y' . Note, however, that we have not yet discussed the n entries $y'_{(r-1)k+k} = y'_{rk}$.

Recall from lemma 9 that $y'_{(r-1)k+i} > 0$ implies that $x'_{(r-1)k+i} > 0$, which implies from the definition of a Nash equilibrium that the $((r - 1)k + i)$ th entry of $M''y'^T$ is at least as big as any other entry. We apply this technique to prove a sequence of useful results.

Suppose for the sake of contradiction that for every r , $y_r'' = 0$. This means that the only nonzero entries are those of the form y_{rk}' . Note, however, that this implies that the (rk) th entries of $M''y''^T$ are all 0, since for any r, s , $M''_{rk,sk} = 0$. However, the (rk) -th columns have ones everywhere else, so every other row gets positive payoff. Thus this is not a Nash equilibrium. We conclude that some $y_r'' > 0$.

For some $y_r'' > 0$, consider the payoffs of the k rows $(r-1)k+1, \dots, rk$. From the construction of the matrix G , we conclude that the first $k-1$ of these payoffs equal $2^j y_r'' + r'_{rk}$ and that the last payoff is at least the sum of the entries in the k corresponding columns: $(M'_{r,r} + 2^j)y_r''$, where by construction, $M'_{r,r} > 0$. Thus this may be a Nash equilibrium only if $r'_{rk} > 0$. Thus whenever $y_r'' > 0$, we must have $y'_{rk} > 0$. The crucial consequence of this is that the payoff of the (rk) th row must now be optimal by the mimicking argument.

Note that the incentive of the (rk) th row is just

$$\sum_s (M'_{r,s} + 2^j)y_s'', \quad (1)$$

which equals the incentive in the game $(M' + 2^j, I)$ when the second player plays strategy y_s'' (up to scaling). Since as noted above, for every nonzero y_r'' , the corresponding row must have optimal incentive, we conclude that the strategy y_s'' , properly scaled, is in fact a Nash equilibrium of the game (M, I) . We have proven one direction of the correspondence.

The other direction is fairly straightforward. Given a Nash equilibrium (x, y) of the game (M, I) , let $y'' = \alpha y$ for some scaling constant α , and let

$$\{y'_{(r-1)k+i}\}_{i=1}^{k-1} = y_r''(2^{j-1}, 2^{j-1}, 2^{j-1}, \dots, 1, 1, 1, \dots),$$

as above. From equation 1 we see that all the optimal incentives in (M, I) will remain optimal in (M'', I') when we restrict our attention to rows rk . Further, each $k-1$ block of rows will have equal payoffs since their corresponding columns have probabilities proportional to the full-support Nash equilibrium of the game (G, I) . To make all these blocks have equal payoffs, we need only pick the additive constants y'_{rk} so that the total payoffs are equal. We then scale these y'_{rk} and α so as to make $\sum_i y'_i = 1$, and we have a Nash equilibrium, as desired.

Thus we have constructed the desired correspondence between the column player strategies in Nash equilibria of the games (M, I) and (M'', I') . ■

References

- [CS1] V. Conitzer and T. Sandholm. Complexity results about nash equilibria. In Proc. of IJCAI, pages 765–771, 2003.
- [CS2] Bruno Codenotti, Daniel Stefankovic. On the computational complexity of Nash equilibria for (0,1)-bimatrix games. Submitted.
- [Pap1] C. H. Papadimitriou, Algorithms, Games, and the Internet. STOC 2001.

