On a rank game

Justin Yu
Plano East Senior High

Abstract

We introduce a new game played by two players that generates an (0,1)-matrix of size n. The first player aims to maximize its resulting rank, while the second player aims to minimize it. We show that the first player can force almost full rank given additional power in move possibilities.

1 Introduction

Matrix rigidity is a concept first introduced by Valiant [1] in 1977 to analyze lower bounds in arithmetic circuits for linear transformations. We say a matrix M is rigid if it retains high rank even if several entries are changed. More precisely, we define $R_M(r)$ as the minimum number of entries that must be changed in order to reduce its rank to r. Valiant showed that if $R_M(O(n)) = O(n^{1+\varepsilon})$ for some M and any $\varepsilon > 0$, then the linear transformation $f : \mathbb{F}^n \to \mathbb{F}^n$ that corresponds to M cannot be computed by an arithmetic circuit of minimal size and depth.

However, while Valiant demonstrates that there exist highly rigid matrices, there is no known explicit family of matrices that satisfies these conditions. The best known results are on finite fields and are due to Friedman [2], for which at least $\Omega\left(\frac{n^2}{r}\log\left(\frac{n}{r}\right)\right)$ entries must be changed to reduce its rank to r. However, this is still not rigid enough to establish the desired bounds on arithmetic circuit complexity. More recently, Alman and Williams [3] proved that the Walsh-Hadamard transformation is not rigid, while Dvir and Liu [4] showed that the Toeplitz and Fourier transform families of matrices also do not meet Valiant's criteria. Surprisingly, all of these were previously conjectured to be rigid.

The idea of rigidity can be viewed as constructing a matrix that maintains high rank even if an adversary alters several elements. Our interest is in an extension of this idea, where there is an "ally" of the matrix who attempts to increase its rank instead of decreasing it.

Consider a game with two players, whom we refer to as the rank maker and the rank breaker. For brevity, we abbreviate their titles as RM and RB, respectively. The game is played on a $n \times n$ grid, which is initially empty. On his first turn, RM selects an unfilled entry on the first row and fills it with a 0 or 1. On her turn, RB selects a different unfilled entry on the first row and fills it with a 0 or 1. When every element of the first row is filled, both players then must fill a cell in the second row. Once every cell in the second row has been filled, both players then must fill a cell in the third row, and so on. The game ends when every cell on the grid has been filled with a 0 or 1. At that point, the grid is interpreted as a matrix in $\mathbb{R}$, whose rank is calculated. RM aims to maximize this rank, while RB aims to minimize it.

This game is not kind towards the first player.

Proposition 1.1. *For every n, RB can force the rank to be at most $\lceil \frac{n}{2} \rceil$.*

Proof. We arbitrarily group all of the columns into pairs (if n is odd, then we ignore the last row). Then whenever RM fills a cell in one columns, RB fills an identical value in the same row of that column's dual. This means that at the end, there are at most $\lceil \frac{n}{2} \rceil$ distinct columns, which means the rank is at most $\lceil \frac{n}{2} \rceil$. □

However, we are still interested in conditions that would allow RM to force a higher rank. This can be achieved by filling two entries on each turn instead of one. We define a player's *strength* as the number of entries they fill at each turn. To distinguish between the multitude of possibilities, we define the (C_1, C_2)-game as the game where RM can fill C_1 entries independently each turn and RB can fill C_2 entries independently each turn. If on any given turn their strength is greater than the number of empty cells in the row, then they independently fill all empty cells in the row, and make their remaining moves in the next row. Also, let $f_n(C_1, C_2)$ denote the outcome of the the (C_1, C_2)-game on an $n \times n$ grid under optimal play. For example, our results above pertain to the $(1, 1)$-game and shows that $f_n(1, 1) \leq \lceil \frac{n}{2} \rceil$.

We next explore if RM is able to perform better given a slight advantage in strength, where he is given $C > 1$ entries instead of 1 at each turn. In this case, we show that RM can achieve almost full rank.

Theorem 1.2. *Let $C > 1$ be an integer. Then $f_n(C, 1) \geq n - 1$ for all sufficiently large n. Furthermore, if n is not divisible by $C + 1$ we in fact have $f_n(C, 1) = n$.*

This is proven in the next section. We expect this generalizes as long as RM has the advantage in strength.

Conjecture 1.3. *As long as $C_1 > C_2$ we have that $f_n(C_1, C_2) \geq n - 1$ for all sufficiently large n and some fixed constant c dependent on C_1 and C_2.*

Now, we introduce another variation of our game. The rules are mostly the same, except that each player is allowed to fill any empty entries on their turn. Like before, we denote $g_n(C_1, C_2)$ by the outcome of this variation under optimal play. Considering the existing work on matrix rigidity allows us to establish a basic result that shows the following:

Proposition 1.4. *Let C_2 be an integer and let $\varepsilon > 0$. Then there exists $C = C(\varepsilon, C_2)$ such that $g_n(C, C_2) > n(1 - \varepsilon)$ for all values of n.*

Proof. RB's strategy is to construct a predetermined rigid matrix. Pudlák and Rödl [5] showed that a majority of (0,1)-matrices in $\mathbb{R}$ have $R_M(\varepsilon n) = \Theta(n^2)$, where ε is sufficiently small. However, any matrix can be reduced to rank $r = \Theta(n)$ by changing at most $(n-r)^2 = \Theta(n^2)$, a result which follows by considering a nonsingular $r \times r$ submatrix of M. RB always has $\Theta(n^2)$ moves in total, which means that this approach is limited to the given bound. $\square$

Still, although we cannot demonstrate this at present we expect that our strategy from Theorem 1.2 will work in a modified form.

Conjecture 1.5. *There exists C such that $g_n(C, 1) \geq n - c$ for all sufficiently large n and some fixed constant c.*

2 Proof of Theorem 1.2

Our first step is to prove the following lemma:

Lemma 2.1. *To extend the rank for the k-th row for any k, it suffices for RM to fill the last entry in the k-th row that belongs to a column in R_i for some $k < i \leq n$.*

Proof. Assume that this is possible. We claim that, when the row is full, the $(k+1)\times(k+1)$ submatrix consisting of columns $1, 2, \ldots, k, i$ is nonsingular. Clearly, the first k columns are linearly independent, so we need to show that C_i is not the sum of some subset of the first k columns. For the sake of contradiction, assume that this is possible. Then the first k rows of the sum are zero, which means that this subset must be $R_i \backslash \{i\}$. However, since RM makes the last move in one of these columns, they can just set it so that the sums of this row are not equal. Thus, all $k+1$ columns would be linearly independent, which means the rank is successfully extended. $\square$

Returning to the main problem, our idea is to construct a strategy that ensures there exists a nonsingular $k \times k$ submatrix after the first k rows are completed for each $1 \leq k \leq n-1$, which can be informally described as "extending the rank at each row." Also, we work over $\mathbb{F}_2$, which implies the result for $\mathbb{R}$.

Proof of Theorem 1.2. Let M denote the $n \times n$ playing field. We prove our strategy works by induction. The base case is $k = 1$, which is trivial because RM can just place a 1 in any spot. For our inductive assumption, assume that $1 < k < n-1$ and the first k rows have been completed and a nonsingular $k \times k$ submatrix exists, which we call M_k. Our goal is to prove the existence of a strategy so that the first $(k+1)$ rows, when filled, will contain a $(k+1)\times(k+1)$ invertible submatrix. By rearranging the columns we can assume without loss of generality that M_k is contained by columns 1 through k.

Let C_i denote the ith column of M. As each column is dynamically extended as gameplay progresses, it will always be interpreted as the state of the column at that moment in time. By definition, every $k \times 1$ vector can be uniquely written as a linear combination of columns in M_k. For $k < i \leq n$, we define the *representation* of a C_i as the set $R_i = \{i, a_1, a_2, \ldots, a_j\}$, where $1 \leq a_1, a_2, \ldots, a_j \leq k$ and

$$C_i = C_{a_1} + C_{a_2} + \cdots + C_{a_j}.$$

We also define M'_k as the $(k+1) \times k$ submatrix obtained by taking M_k and the elements immediately below it.

Our strategy is divided into multiple stages. The first stage occurs when $1 < k < \frac{n}{2}$. We provide a Method A and a Method B, which are used under varying circumstances. Method A will be used if there exist columns with only zeroes, and Method B will be used otherwise.

With Method A, RM picks columns with only zeroes and fills them with a 1. In addition to decreasing the number of all zero columns, this guarantees that the rank has successfully been extended, as there is no linear combination of the first k columns that sum to $(0, 0, \ldots, 0, 1)^T$.

Because RM has C moves and RB has 1, Method A also reduces the number of all zero columns to around $\frac{1}{C+1}$ of the previous value, so they should disappear completely after around $\log_{C+1}(n)$ rows of gameplay. By making n sufficiently large, this can be made much less than $\frac{n}{2}$, which ensures this task can be fulfilled before the second stage of the overall strategy.

If such an initial move is not possible (i.e. there is only one column with all zeroes and RB already claimed it), Method B is used. However, this situation will not happen consecutively, so for the next row Method A can be used.

With Method B, RM arbitrarily fills spots in the first k columns. Because $k < \frac{n}{2}$, they will finish this task before all elements of the row have been filled. Then, on their next move in the same row (which exists), they pick an empty column C_i. We claim that this column can be made linearly independent from all columns in M'_k, thus creating the desired submatrix. This is a consequence of Lemma 2.1, because by assumption everything in $R_i \backslash \{i\}$ has already been filled.

Thus, by the time that $k \geq \frac{n}{2}$, our strategy has ensured that full rank is maintained and there are no zero columns remaining. This allows us to proceed with the second stage, which applies for all subsequent k up to $n-2$.

Note that Method B cannot be extended to any k larger than $\frac{Cn}{C+1}$, because the assumption that RM can fill up the first k columns quickly becomes false. However, by Lemma 2.1, we only need to fill up the representation of another column as opposed to the entire space below M_k.

Let

$$R_{k+1} = \{a_1, \ldots, a_i, k+1\}$$

$$R_{k+2} = \{b_1, \ldots, b_j, k+2\}$$

be the representations of columns $k+1$ and $k+2$. These sets have cardinality of at least two because we have established that every column is nonzero, so both columns $k+1, k+2$

are the sum of a nonzero collection of other columns in M_k. Recall that by Lemma 2.1, RM just has to fill the last cell belonging to a column in R_{k+1} to win the row, and likewise for R_{k+2}.

Firstly, if either R_{k+1} or R_{k+2} have size less than $\frac{Cn}{C+1}$, then we can apply Method B with respect to the smaller one. Otherwise, both of them have size of at least $\frac{Cn}{C+1}$, so by the Pigeonhole Principle their intersection is nonempty because all elements are in $\{1, 2, \ldots, n\}$. Now our goal is to show that $R_{k+1} \oplus R_{k+2}$ also has this property, where $\oplus$ denotes the exclusive-or operation. Without loss of generality, set

$$R_{k+1} \cap R_{k+2} = \{a_1, a_2, \ldots, a_t\} = \{b_1, b_2, \ldots, b_t\}$$

Now, we claim that the submatrix obtained by swapping columns $k+1$ and a_1 is also nonsingular. Let N_k be the $k \times k$ submatrix of M that contains the first $k+1$ columns *excluding* C_{a_1}. It suffices to show that every $\{0, 1\}$ column of height k can be written as a nonzero sum of some columns in N_k. However, we have that this property is true for M_k and

$$C_{a_1} = C_{a_2} + C_{a_3} + \cdots + C_{a_i} + C_{k+1}$$

by the definition of a representation, so we can replace C_{a_1} with this equivalent any time it appears (if a column is present twice in a sum we can remove it because we're working in $\mathbf{F}_2$).

Now, by definition we have that

$$\begin{aligned}
C_{k+2} &= C_{b_1} + C_{b_2} + \cdots + C_{b_j} \\
&= C_{a_2} + C_{a_3} + \cdots + C_{a_t} + \cdots + C_{a_i} + C_{k+1} \\
&+ C_{b_2} + C_{b_3} + \cdots + C_{b_t} + \cdots + C_{b_j} \\
&= C_{a_{t+1}} + \cdots + C_{a_i} + C_{b_{t+1}} + \cdots + C_{n_j} + C_{k+1}
\end{aligned}$$

which implies that $R_{k+1} \oplus R_{k+2}$ is the representation set of C_{k+2} with respect to N_k, as desired. Thus, being the last to fill any of $R_{k+1}, R_{k+2}, R_{k+1} \oplus R_{k+2}$ will allow RM to extend M_k's rank. By similar logic to the above, we can assume that $|R_{k+1} \oplus R_{k+2}| \geq \frac{Cn}{C+1}$

Now, define pairwise disjoint sets X, Y, Z such that

$$\begin{aligned}
X &= R_{k+1} \cap R_{k+2} \\
Y &= R_{k+1} \backslash R_{k+2} \\
Z &= R_{k+2} \backslash R_{k+1}.
\end{aligned}$$

Clearly,

$$\begin{aligned}
X \cup Y &= R_{k+1} \\
X \cup Z &= R_{k+2} \\
Y \cup Z &= R_{k+1} \oplus R_{k+2}.
\end{aligned}$$

At this point we can ignore R_{k+1} and R_{k+2}, as RM's goal changes to being the last person to fill out two of X, Y, Z. Without loss of generality, assume that $|X| \leq |Y| \leq |Z|$. Note that their union is large because A, B are by assumption. Furthermore, at most one of X, Y, Z can have size less than 4, because $|X \cup Y| \geq \frac{Cn}{C+1}$ and vice versa. RM's strategy is to start by filling X with arbitrary values. Because X is the smallest set and $|Y|, |Z| \geq 4$, neither Y nor Z will have been filled by RB once he is finished.

From here, RM's strategy is to fill whichever of Y or Z has more unclaimed columns. The exception, of course, is if one of Y, Z has one or two elements remaining, in which case he takes them for the immediate win.

To show this works, assume that RB is the first to take the last element of one of Y, Z under optimal play, say Y. For RM's turn before, Y could have one, two, or three elements left. But having one or two elements means RM wins, so there must have been three left, and he took two from Y. The only case where this is allowed in our strategy is if Z also has 3 unclaimed elements. In this case, RM can just take 1 from each set, which is clearly a winning position. Thus, this never happens, as desired.

This process allows RM to maintain full rank throughout the first $n-1$ rows, which establishes the inequality. To prove the equality, we note that in those cases (where $C+1 \nmid n$), RM moves last. When filling out the nth row, it is not hard to show that there is only one representation set among all possible choices of M_{n-1} (they all overlap), so making the last move in the row guarantees RM has an easy strategy to claim its last element. On the other hand, if RB has the last move and the representation set is large, then the rank stays at $n-1$.

□

3 Acknowledgements

I would like to thank my mentor Dr. Asaf Ferber for his mentorship and problem suggestions. Additionally, I would like to thank the PRIMES-USA program for providing me with this opportunity and Dr. Tanya Khovanova for her stylistic feedback and advice.

References

[1] L. G. Valiant. Graph-theoretic arguments in low-level complexity. *Mathematical Foundations of Computer Science*, pages 162–176, 1977.

[2] Joel Friedman. A note on matrix rigidity. *Combinatorica*, 13(2):235–239, 1993.

[3] Josh Alman and Ryan Williams. Probabilistic rank and matrix rigidity. *49th ACM Symposium on Theory of Computing*, 2017.

[4] Zeev Dvir and Allen Liu. Fourier and circulant matrices are not rigid. *34th Computational Complexity Conference (CCC 2019)*, 137:17:1–17:23, 2019.

[5] Satyanarayana V. Lokam. Exercises on matrix rigidity. Available at `https://simons.berkeley.edu/sites/default/files/docs/1738/exercises.pdf`, 2014.