Some structural properties of low rank matrices related to computational complexity

Bruno Codenotti*   Pavel Pudlák [†]   Giovanni Resta*

September 25, 1997

Abstract

We consider the conjecture stating that a matrix with rank $o(n)$ and ones on the main diagonal must contain nonzero entries on a 2×2 submatrix with one entry on the main diagonal. We show that a slightly stronger conjecture implies that an explicit linear transformation cannot be computed by linear size and logarithmic depth circuits. We prove some partial results supporting the conjecture.

1   Introduction

The problem of relating the rank of a matrix to its structural properties given by the pattern of its nonzero entries is a classical problem in mathematics. In complexity theory the most famous instance of this problem is the relation between the communication complexity of a $\{0, 1\}$ matrix and its rank over the field of reals [10, 6]. In this paper we consider a less known problem which, however, might have even more interesting consequences, as it could lead to a nonlinear lower bound on some algebraic circuits [7]. Unlike in the case of communication complexity, it seems that the nature of the problem under investigation here does not depend so much on the field in question, whereas the relevant 'structure' seems to be just provided by the distinction between zero and nonzero elements. This research goes in the direction proposed by Valiant [12], who suggested that some lower bounds on the size of circuits can be proved by constructing matrices with certain properties. Before we explain the connection to Valiant's results, we have to introduce some concepts.

We shall call $[2, 2]$ *configuration* a 2×2 submatrix consisting of nonzero elements and having exactly one entry on the main diagonal. In graph theoretical terms a $[2, 2]$ configuration corresponds to a transitive triangle, thus we shall sometimes call it just *triangle*.

Conjecture 1 *([The Triangle Conjecture]) For every field F, there exists $\varepsilon > 0$ such that every $n \times n$ matrix M with nonzero entries on the main diagonal, and such that $rank(M) \leq \varepsilon n$, contains a $[2, 2]$ configuration.*

---

*Istituto di Matematica Computazionale del CNR, Via S. Maria 46, 56126-Pisa (Italy). e-mail: {codenotti,resta}@imc.pi.cnr.it.

[†]Mathematical Institute, AVČR, Žitná 25, 115 67 Praha 1, Czech Republic. e-mail: pudlak@math.cas.cz.

The Triangle Conjecture implies a nonlinear lower bound on the computation of cyclic shifts by *semilinear circuits* [7]. Proving nontrivial lower bounds on cyclic shifts is an important task, since this class of functions can be reduced to the multiplication function. Another notion to which it is related is that of matrix *rigidity*, a concept introduced by Valiant [12]. The rigidity of a matrix M is the function $R_M(r)$, which for a given r gives the minimum number of entries of M which one has to change in order to reduce its rank to r or less. Valiant proved the following result.

Theorem 1 ([12]) *If for some $\varepsilon > 0$, the $n \times n$ matrix M_n has rigidity $R_{M_n}(\varepsilon n) \geq n^{1+\varepsilon}$, then the transformation $x \to M_n x$ cannot be computed by linear size and logarithmic depth circuits with gates computing linear functions over a given field.*

Although both a random matrix and a matrix whose entries are different indeterminates have rigidity even larger than required by Theorem 1 (close to n^2), very little is known about *explicit* matrices. The best known lower bounds on the rigidity of explicit matrices are of the form $\Omega(\frac{n^2}{r} \log \frac{n}{r})$ [3], which gives only linear lower bounds on $R_M(\varepsilon n)$. It seems that Hadamard matrices have large rigidity over the real field, but the best bound is so far only $\Omega(n^2/r^2)$, due to Alon (unpublished).

Another relation of this type was found by Razborov [8]. He proved that some weaker bounds on rigidity would imply that a $\{0,1\}$ matrix defines a function which is not in the communication complexity version of the polynomial hierarchy. The existing lower bounds on the rigidity are, however, insufficient even for that.

As an intermediate step, Smolensky (and independently one of the authors) proposed to prove a large lower bound on the rigidity of a Toeplitz matrix with indeterminates. He suggested to make the problem of proving a lower bound easier by allowing the changed values to be just linear functions of the indeterminates. The proof in [7] can be easily adapted to show that the 'triangle conjecture' implies a lower bound of the type needed in Theorem 1 for a Toeplitz matrix with indeterminates. (This is true even if one allows the changes to be polynomials.)

In this paper we come up with an explicit matrix which has the rigidity required by Theorem 1, just assuming a slightly stronger conjecture (Section 2). In the current situation where we lack any methods for proving large lower bounds on rigidity, this gives a concrete program for proving such bounds, provided that the Triangle Conjecture is true. Even if the conjecture fails, this matrix may be a good candidate for large rigidity.

We also show a construction of circulant matrices which gives nonlinear bounds on the rigidity $R_M(\varepsilon n)$, for every fixed $\varepsilon > 0$ assuming only the Triangle Conjecture (Section 3). Such a bound is sufficient for proving a nonlinear lower bound on the size of series-parallel circuits computing the transformation M.

In the following sections we present some supporting evidence for and results on the conjecture. We have tested the conjecture for some small symmetric matrices, in the case of rank over $GF[2]$. For the sizes $n \leq 32$ we have verified that every symmetric matrix with ones on the main diagonal and rank $\leq n/4+1$ contains a triangle. There is a unique, up to isomorphisms, family of symmetric matrices of rank $n/4 + 2$ with ones on the main diagonal and without triangles. Such matrices do exist for every n, but we cannot prove that they are extremal. They will be described in Section 5.

We shall also show a decomposition for symmetric matrices over $GF[2]$ with at least one 1 on the main diagonal, which simplifies either the search for counterexamples or a possible proof (Section

4). Namely, every such matrix A can be represented as $UU^\top$. This representation allows us to investigate the conjecture for $GF[2]$ in a purely combinatorial way, since we can treat the rows of the matrix U as a set system. Since in the Triangle Conjecture we assumed that all the diagonal entries are nonzero, we have that all sets have odd cardinalities. A triangle corresponds to three sets, every two of which intersecting in odd sets. In Section 6 we prove a theorem on set systems which implies the conjecture for the special case where the number of ones in each row of U is bounded by a constant. This theorem is a version of the famous Erdös-Rado theorem on *sunflowers* (also called *delta systems*). Accidentally, this type of a theorem has been used at least twice to prove a lower bound on the size of circuits [9, 4]. Especially the last one is very much related to our new result, though the application is completely different.

We also show a more general result that applies to matrices that admit, over $GF[2]$, the factorization UV, where the number of ones in each row of U and each column of V is bounded by a constant.

In the case of the field of reals it seems that the triangle conjecture holds with $\varepsilon = 1/2$. Our experimental evidence is supported by a result of Rosenfeld [11]. He proved that a symmetric positive definite $n \times n$ matrix of rank $\leq n/2$ with ones on the main diagonal contains a $[2,2]$ configuration. His proof uses properties of the eigenvalues of the matrix, which do not hold in the non definite case. However he pointed out that our conjecture can be reduced to a stronger conjecture on symmetric matrices. Namely, consider the following conjecture, for some fixed field F and $\varepsilon > 0$.

We use the term *principal submatrix* to denote a submatrix which shares the main diagonal with a given matrix.

Conjecture 2 *Every symmetric $n \times n$ matrix of rank $\leq \varepsilon n$ contains a 6×6 principal submatrix of nonzero elements.*

Conjecture 2 implies the Triangle Conjecture for fields of characteristic different from 2, as follows from the argument below. Take a matrix M with rank $\leq \frac{\varepsilon}{2}n$ (ε from Conjecture 2), and let $A = M + M^\top$. By Conjecture 2, A contains a 6×6 principal submatrix of nonzero elements. This submatrix corresponds to a complete graph on 6 elements. Color the edges of this graph blue, if the corresponding entry in the right upper half of M is nonzero, and red otherwise. Since 6 is the Ramsey number $R(3,3)$, we must have either a blue or a red triangle, i.e., a $[2,2]$ configuration either in the right upper half or in the left lower part.

This conjecture seems to be less likely to be true than the previous one. Alon and Szegedy [1] proved that if we replace 6 by a sufficiently large constant, the statement is false. Namely, for every $\delta > 0$ there exists k such that there are matrices with ones on the main diagonal and rank $\leq n^\delta$ with no $k \times k$ principal submatrices of nonzero elements. The minimal k for which one can get sublinear rank from their proof is still rather large, but there is no a priori reason for that.

In general there is only an $\Omega(\sqrt{n \log n})$ lower bound on the rank of $\{0,1\}$ matrices with ones on the main diagonal and without triangles. This bound easily follows from the well-known bound on the Ramsey number $R(3,k) = O(k^2 / \log k)$. Namely, let an $n \times n$ matrix M be given and $n \geq R(3,k)$. Color the edges of the complete graph on n vertices as above, but now using the given matrix M. If M does not contain a $[2,2]$ configuration in the right upper part, the complete graph does not contain a blue triangle. Hence there must be a red complete subgraph on k elements. This corresponds to a $k \times k$ principal submatrix with zeros above the main diagonal, which has rank k.

Although our results do not improve on any current lower bound on circuit complexity, we nevertheless think that we made a visible progress in that area. Fundamental problems in circuit complexity cannot be solved by gradually increasing lower bounds. We need to make progress in associated combinatorial and algebraic problems, and this paper is a step in this direction.

2 A possibly rigid matrix

Call $[2, 3]$ *configuration* a 2×3 submatrix consisting of nonzero elements and having at least one entry on the main diagonal. A $[3, 2]$ configuration is defined in a similar manner. Note that unlike the case of $[2, 2]$, we allow two entries to be on the main diagonal; this is to make the following conjecture weaker.

Conjecture 3 *For every field F there exists $\varepsilon > 0$ such that every $n \times n$ matrix M with $\mathrm{rank}(M) \leq \varepsilon n$ contains either a $[2, 3]$ or a $[3, 2]$ configuration.*

In Lemma 2 below, we prove some properties of an explicit family of *circulant* matrices. We will then use these properties to show that, assuming Conjecture 3, these matrices have high rigidity (see Theorem 3).

Lemma 2 *For each $n > 2$ there exists at least one $\{0, 1\}$ circulant matrix C_n such that*

1. *C_n contains at least $n^{1+1/5}$ nonzero entries;*

2. *C_n does not contain a 2×2 submatrix of nonzero entries, i.e., every two rows and every two columns share at most one nonzero coordinate;*

3. *for every pair i, j of indices of a zero entry of C_n, there exist at most two 2×2 submatrices containing this entry and with the other entries different from zero.*

Proof. For a given n, we consider a subset $A_n = \{a_1, a_2, \ldots, a_k\}$ of $\{0, \ldots, n-1\}$ such that (a) all the two-term sums $S^{(2)} = \{a_i + a_j : 1 \leq i \leq j \leq k\}$ are distinct mod n, and (b) all the three-term sums $S^{(3)} = \{a_h + a_i + a_j : 1 \leq h \leq i \leq j \leq k\}$ are distinct mod n. It is easy to see that one can use a "greedy" algorithm to construct a set A_n satisfying (a) and (b), and such that $a_k \leq k^5$, from which $|A_n| \geq n^{\frac{1}{5}}$.

Let now C_n be the $n \times n$ $\{0, 1\}$ circulant matrix whose first row is given by the characteristic vector of the set A_n. We shall index the rows and the columns by numbers $0, \ldots, n-1$. For $i \in \{0, \ldots, n-1\}$, we shall call the $i-th$ diagonal the set of entries with indices $(0, i), (1, i+1), \ldots, (n-1, i+n-1 (\equiv_n))$.

1. The circulant matrix C_n has at least $|A_n| \geq n^{1/5}$ nonzero diagonals and thus at least $n^{1+1/5}$ nonzero elements.

2. Let us assume that the matrix C_n contains a rectangle, and denote by a, b, c, and d the numbers of the diagonals corresponding to the corners of the rectangle, with a and b on the

first row and c and d on the second row of the 2×2 submatrix. It is easy to see that, since C_n is circulant, we must have

$$b - a \equiv_n d - c. \tag{1}$$

Hence $c + b \equiv_n a + d$, which is in contradiction with property (a) if the two sets $\{c, b\}$ and $\{a, d\}$ do not coincide. Since two entries of the matrix cannot simultaneously lie on the same diagonal and on the same row or column, we have $c \neq a$ and $c \neq d$, from which the thesis follows.

3. Let us assume that the matrix C_n contains a zero entry which completes a 2×2 submatrix and, as above, let us denote by a, b, c and d the numbers of the diagonals on which the corners lie. Since the matrix is circulant, we can assume without loss of generality that the zero entry is on its first row. We further assume that the zero entry belongs to the diagonal a. The other case (zero entry on diagonal b) is similar. The other possible rectangles (a', b', c', d') and (a'', b'', c'', d'') containing the selected zero will have it in either the left or the right upper corner. Let us consider the first case, i.e., $a' = a$. By (1), we have $b + c - d \equiv_n a = a' \equiv_n b' + c' - d'$, and thus

$$b + c + d' \equiv_n b' + c' + d. \tag{2}$$

Since C_n satisfies (b), the two sets $\{b, c, d'\}$ and $\{b', c', d\}$ must be equal.

Since d' cannot lie on the same diagonal as b' or c', we have that $d' \neq b'$ and $d' \neq c'$. Thus it has to be $d' = d$. Since the two 2×2 submatrices are distinct, then the equalities $b' = b$ and $c' = c$ are ruled out, and thus the only way to satisfy (b) and (2) is given by

$$a' = a, \ b' = c, \ c' = b, \ d' = d. \tag{3}$$

In the other case, starting from $a = b''$, we obtain $b + c - d \equiv_n a = b'' \equiv_n a'' + d'' - c''$ and thus $b + c + c'' \equiv_n a'' + d'' + d$. Proceeding as above, we obtain the set of equalities

$$a'' = c, \ b'' = a, \ c'' = d, \ d'' = b. \tag{4}$$

The two rectangles determined by (3) and (4) are incompatible, since by construction $c = a'' \leq a < b \leq b' = c$. Hence at most two 2×2 submatrices, i.e., (a, b, c, d) and one among (a', b', c', d') and (a'', b'', c'', d''), can contain the zero entry on diagonal a.

$\square$

The results shown in Lemma 2 are the key ingredients in the proof of the following Theorem.

Theorem 3 *Assuming Conjecture 3, for every field F there exists an $\varepsilon > 0$ such that*

$$R_{C_n}(\varepsilon n) \geq n^{1+1/20}.$$

Hence the linear transformation determined by C_n cannot be computed by a linear size logarithmic depth circuit.

Proof. First observe that Conjecture 3 implies that for a sufficiently small constant $\varepsilon > 0$ every matrix with ones on the main diagonal and rank less than εn contains a linear number of $[2,3]$ or $[3,2]$ configurations. Namely, if ε is sufficiently small, we can find either a $[2,3]$ or a $[3,2]$ configuration, omit rows and columns containing this configuration, and repeat the process, until the ratio between rank and matrix size becomes too small.

Fix an ε which is four times smaller. Suppose the rank of C_n has been reduced to εn by m changes, and let $d = m/n$. Omit rows and columns which contain more than $4d$ changes. Thus we get a submatrix, say C', of size at least $\frac{3}{4}n \times \frac{3}{4}n$, which contains at least $1/2$ of each diagonal of ones of C_n. We consider two cases.

Case 1: more than half the elements of C' are changed to 0 on more than half the diagonals of ones of C_n. Then m, the number of changes, is at least $\frac{1}{4}n^{1+1/5}$.

Case 2: at least half the elements of C' remain 1 on at least half the diagonals of C_n. For each of these diagonals, take the square submatrix of C' determined by the intersection of the diagonal with C' and apply the conjecture. There must be $\Omega(n)$ $[2,3]$ or $[3,2]$ configurations for each of these diagonals. As one configuration can be shared by at most 6 diagonals (using the properties of C_n, we can actually show that by at most 4) there must be altogether $\Omega(n^{1+1/5})$ such configurations contained in C'. We shall show that the number of the configurations can be bounded from above by $O(nd^4)$, whence $m = \Omega(n^{1+1/20})$, which will complete the proof.

We shall consider several cases according to which elements of the $[2,3]$ configuration are original from C_n and which are new, introduced by the changes. The bound on $[3,2]$ configurations follows by symmetry.

1. Suppose there is a 2×2 submatrix of the $[2,3]$ configuration which has three old and one new elements. Then it is uniquely determined by the new element, by the second property of C_n. Hence there are at most m possible ways to choose such a submatrix. By the first property, there can be at most one old element among the remaining two of the $[2,3]$ configuration. The new element is on one of the rows already determined, so there are at most $8d - 1$ ways of choosing it. Thus we get the bound $(8d - 1)m \leq 8d^2 m$ on such $[2,3]$ configurations.

2. Consider $[2,3]$ configurations which do not fall under the first case and which contain a 2×2 submatrix with a row of old elements and the other row of new elements or a column of old elements and the other column of new elements. By the first property of C_n, such a submatrix is determined by one of the new elements (m possibilities) and the other new one ($4d$ possibilities, as the row or column is already determined). At least one of the remaining two from the $[2,3]$ configuration is new, either by the first property of C_n or because the $[2,3]$ configuration does not contain a submatrix which has three old and one new elements. There are at most $4d$ choices for that element, which gives the bound $16d^3 n$.

3. Consider $[2,3]$ configurations which do not fall under any of the above two cases. Then each row and each column of it contains at most one old element. Thus it contains a 2×2 submatrix with at least three new elements. It is determined by choosing one of them (m possibilities) and then two others (each $\leq 4d - 1$ possibilities). Again, among the last two there must be at least one new ($\leq 4d - 1$ possibilities). Altogether it gives $O(d^m) = O(d^4 n)$. $\qquad\square$

3 Another Construction

We describe an explicit construction of circulant matrices which have rigidity of the order of $n(\log n)^{1/3}$, provided that the Triangle Conjecture is true. For technical reasons, in this section we number rows and columns of the matrices starting from 0, rather than 1. We construct a circulant matrix C'_n whose first row has nonzero entries in columns $1, b, b^2, \ldots, b^k$, where the choices of b and k are described below.

Lemma 4 Let $n = 2^{2m} - 1$, and define $a = 2^{2m-1} + 2^{m-1}$, and $b = a + 1$. The following relations hold over $\mathbf{Z}_n$ for $1 \leq h \leq m$:

$$a^2 \equiv_n a \tag{5}$$

$$b^h \equiv_n 2^{2m-1} - 2^{m-1} + 2^{h-1} + 2^{h+m-1}. \tag{6}$$

Proof. From $2^{2m} \equiv_n 1$, we easily obtain (5), since

$$a^2 = 2^{4m-2} + 2 \cdot 2^{2m-1} 2^{m-1} + 2^{2m-2} \equiv_n 2 \cdot 2^{2m-2} + 2^{m-1} \equiv_n a \,.$$

Hence we also have that $a^h \equiv_n a$, for $h > 0$. Relation (6) is obtained as follows

$$
\begin{aligned}
b^h &= (a+1)^h = 1 + \sum_{i=1}^{h} \binom{h}{i} a^i = (2^h - 1)a + 1 \\
&\equiv_n 2^{h-1} + 2^{h+m-1} - 2^{2m-1} - 2^{m-1} + 1 \equiv_n 2^{2m-1} - 2^{m-1} + 2^{h-1} + 2^{h+m-1} \,.
\end{aligned}
$$

where we used (5) and $2^{2m} \equiv_n 1$ to simplify the expressions. $\square$

Corollary 5 The set $\{1, b, b^2, \ldots, b^{m-1}\}$, with the elements taken modulo n, has size m and it is a subgroup of the multiplicative group $\mathbf{Z}_n^*$.

Proof. The size is immediate from (6). To see that it is a subgroup, just check that $b^m \equiv_n 1$. $\square$

Let us consider, for an integer α invertible over $\mathbf{Z}_n$, a matrix C''_n defined by

$$c''_{i,j} = c'_{\alpha i, \alpha j} \,, \tag{7}$$

where indices run from zero and are computed over $\mathbf{Z}_n$.

It is easy to see that the effect of (7) is to permute the diagonals in such a way that C''_n is still circulant. In particular, if $\alpha = b^{-j}$, with $1 \leq j \leq k$, the elements of the diagonal corresponding to b^j are moved to diagonal 1, and, since $\{1, b, b^2, \ldots, b^k\}$ and $\{b^{-j}, b^{-j+1}, \ldots, b^{k-j}\}$ coincide (by Corollary 5), we have $C''_n = C'_n$.

We summarize relevant properties of C'_n in the following observations.

Observation 1 Let $n = 2^{2m-1}$. There are $m-1$ permutation matrices Q_k such that the automorphism $Q_k C'_n Q_k^T = C''_n = C'_n$ corresponds to the transformation (7). In particular the permutation matrix Q_h, defined as $q_{ij} = 1$ iff $j = b^{-h}i$ and 0 elsewhere, takes the elements of diagonal b^h onto diagonal 1.

Observation 2 *Let M be the matrix obtained from C'_n by deleting its first column and last row. M has a principal submatrix of order $\frac{n}{4}$ which is an identity matrix, since it is easy to verify, from (6), that $n/2 < b^j \pmod{n} < 3n/4$, for $1 \le j \le k$.*

The above two observations can be used to prove the following Theorem.

Theorem 6 *Assuming the Triangle Conjecture, for every field F there exists an $\varepsilon > 0$ such that*

$$R_{C'_n}(\varepsilon n) = \Omega(n(\log n)^{1/3}).$$

Proof. By Observation 2, we have that the submatrix M (associated to the first diagonal of C'_n) contains an $\frac{n}{4} \times \frac{n}{4}$ identity matrix.

Let us assume that the Triangle conjecture is true. Then, in order to decrease the rank of M below εn, for a suitable constant ε, we must introduce a linear number of triangles or change a linear number of the diagonal entries to 0 (as in the proof of Theorem 3 for $[2, 3]$ configurations). By Observation 2, we actually a have linear number of triangles which do not contain entries from other diagonals of C'_n.

By Observation 1, we can rearrange C'_n by means of permutations so that the elements of each diagonal can in turn be moved to the first diagonal. This implies that we can repeat the previous argument for all the $m-1$ *diagonals* of C'_n. Thus either more than half of the elements on more than half of the diagonals are changed to 0, in which case we are done, as this gives $\Omega(nm) = \Omega(n \log n)$ changes, or there are $\Omega(nm)$ triangles. To get a lower bound on the number of changes in the latter case, we shall use similar counting as in the proof of Theorem 3. Let d be the average number of changes in a row. As in that proof we may assume that each row and each column contains at most $4d$ changes. Each triangle is determined by choosing a row (n choices), two elements in the row ($\le \binom{4d}{2}$ choices) and an element in one of the two columns ($\le 8d$ choices). Thus we get $nd^3 = \Omega(nm)$, whence the number of changes must be $\Omega(n(\log n)^{1/3})$. $\qquad\square$

4 A construction over $GF[2]$ for symmetric matrices

We show a decomposition for symmetric matrices in the field $GF[2]$ which is useful for studying our conjecture for such matrices and field. For an $n \times n$ symmetric matrix A of rank r, with at least one 1 on the main diagonal, we show how to construct the factorization $A = UU^T$, where U is an $n \times r$ matrix. Note that the diagonal of UU^T contains all ones if and only if each row of U contains an odd number of ones. Moreover A is triangle-free if and only if for every submatrix V of U, of size $3 \times r$, we have $VV^T \ne J_3$, where J_3 is the 3×3 matrix whose entries are all equal to 1.

Lemma 7 *Given an $n \times n$ symmetric matrix A whose rank over $GF[2]$ is r, there exists an invertible $n \times n$ matrix S such that*

$$SAS^\top = B = \begin{pmatrix} P & \mathbf{0} \\ \mathbf{0}^\top & \mathbf{0} \end{pmatrix}, \tag{8}$$

where P is an $r \times r$ symmetric permutation matrix.

Proof. Let us denote with E_{ij} an $n \times n$ matrix whose only nonzero element is $e_{i,j} = 1$, and let us define the matrices $M_{ij} = I + E_{ij}$ for $i \neq j$. We now consider

$$A' = M_{ij} A M_{ij}^\top = A + E_{ij} A + A E_{ji} + E_{ij} A E_{ji} \, ,$$

i.e., A' is obtained from A adding the j-th row to the i-th row, and the j-th column to the i-th column. In particular $a'_{i,i}$ is equal to $a_{i,i} + a_{j,j} + a_{i,j} + a_{j,i} = a_{i,i} + a_{j,j}$. Note that matrices M_{ij} are invertible and idempotent. We will show a procedure that constructs a matrix S that satisfies (8). S will be obtained as a product of matrices which are either of the M_{ij} type or permutation matrices.

In the following we describe the k-th step of the algorithm. It starts assuming that each of the rows and columns indexed by $1, ..., k - 1$ contains exactly one 1 and that if $a_{i,j} = 1$ for $j < k$, then $a_{i,j}$ and $a_{j,i}$ are the only nonzero entries of i-th row and column. In other words, we assume that the first k rows and columns of A have a structure 'compatible' with the construction of a permutation matrix. The algorithm starts with $k = 1$, so that these conditions are trivially satisfied, and we also let $S = I_n$.

Let us consider the k-th row and column of A. If they do not contain nonzero entries, then we apply a symmetric permutation Π to A in such a way that $A \leftarrow \Pi A \Pi^\top$ has ones in the k-th row and column, and we let $S \leftarrow \Pi S$. If no such permutation exists, then all the rows and columns from the k-th onward are null, and we are done. After the possible application of Π we have two cases:

- $a_{k,k} = 1$. For each other $a_{k,j} = 1$, $j > k$, we let $A \leftarrow M_{jk} A M_{jk}^T$ and $S \leftarrow M_{jk} S$. The effect of this operation is to annihilate all the $a_{k,j}$ (and $a_{j,k}$) for $j > k$. Rows and columns $h < k$ are not modified and $a_{k,k}$ is the only nonzero left on row and column k. This is compatible with the conditions needed to proceed in the next step.

- $a_{k,k} = 0$. If there is an index $h > k$ such that $a_{h,h} = 1$, then we apply a symmetric permutation Π to A that exchanges rows (and columns) h and k, we let $S \leftarrow \Pi S$, and we apply the case $a_{k,k} = 1$ above. Otherwise let h be the smallest index such that $a_{k,h} = 1$. Then for each other $a_{k,j} = 1$, $j > h$, we let $A \leftarrow M_{jh} A M_{jh}^T$ and $S \leftarrow M_{jh} S$. The effect of this operation is to annihilate all the $a_{k,j}$ (and $a_{j,k}$) for $j > h$, while rows and columns with indices less than h are not affected. (In fact, by hypothesis, there can not be nonzeros in the entries $a_{i,j}$ for $i < k$, because otherwise we should have $a_{k,j} = 0$). The only nonzero entry of row k is thus $a_{k,h}$. Then we apply the same procedure to row and column h. At the end, $a_{k,h}$ is the only nonzero of row k and column h. The same holds for $a_{h,k}$, so that we can proceed with the next step.

It is clear that the above algorithm ends in at most n steps, and reduces the original matrix A to a matrix B that satisfies (8). In particular, since the matrix S constructed during the algorithm is nonsingular, then B and A have the same rank r. Thus P has size r, and hence the algorithm above stops after at most r steps. $\qquad\square$

Lemma 8 *Given an $n \times n$ symmetric permutation matrix P, there exists an $n \times n$ matrix U such that $P = UU^\top$ over $GF[2]$, if and only if P has at least one nonzero entry on the main diagonal.*

Proof. Let us assume first that P has at least one nonzero entry on the main diagonal. Let $\alpha_1, \alpha_2, \ldots, \alpha_{2k}$ be the indices i such that $p_{i,i} = 0$. These indices are ordered so that $p_{\alpha_{2j-1}, \alpha_{2j}} = 1$ for $j = 1, \ldots, k$. If $k = 0$ then $P = I$ and $U = I$ and we are done. Otherwise let $\beta_1, \beta_2, \ldots, \beta_{n-2k}$ be the indices i such that $p_{i,i} = 1$. By hypothesis we have $n - 2k \geq 1$. The rows of U, denoted by $U_1, \ldots, U_n$, can be described as follows:

$$
\begin{aligned}
U_{\alpha_{2j-1}} &= [1_{2j-1}\, 10\, 0_{n-2j-1}] &&\text{for } j = 1, \ldots, k \\
U_{\alpha_{2j}} &= [1_{2j-1}\, 01\, 0_{n-2j-1}] &&\text{for } j = 1, \ldots, k \\
U_{\beta_1} &= [1_{2k+1}\, 0_{n-2k-1}] && \\
U_{\beta_j} &= [0_{2k+j-1}\, 1\, 0_{n-2k-j}] &&\text{for } j = 2, \ldots, n - 2k \,,
\end{aligned}
$$

where 1_h (resp. 0_h) denotes a string of h ones (resp. zeros). Is is easy to see that $UU^\top = P$. In fact we have

- $U_{\alpha_{2j-1}} U_{\alpha_{2j}}^\top = 2j - 1 \equiv_2 1$, for $j = 1, \ldots, k$.

- $U_{\alpha_{2j-1}} U_{\alpha_{2h-1}}^\top = U_{\alpha_{2j}} U_{\alpha_{2h}}^\top = \min(2j, 2h) \equiv_2 0$, for $j, h = 1, \ldots, k$.

- $U_{\alpha_{2j-1}} U_{\alpha_{2h}}^\top = \min(2j, 2h) \equiv_2 0$, for $j, h = 1, \ldots, k$, and $j \neq h$.

- $U_{\alpha_{2j-1}} U_{\beta_1}^\top = U_{\alpha_{2j}} U_{\beta_1}^\top = 2j \equiv_2 0$, for $j = 1, \ldots, k$.

- $U_{\alpha_{2j-1}} U_{\beta_h}^\top = U_{\alpha_{2j}} U_{\beta_h}^\top = 0$, for $j = 1, \ldots, k$ and $k = 2, \ldots, n - 2k$.

- $U_{\beta_1} U_{\beta_1}^\top = 2k + 1 \equiv_2 1$.

- $U_{\beta_j} U_{\beta_j}^\top = 1$, for $j = 2, \ldots, n - 2k$.

- $U_{\beta_j} U_{\beta_h}^\top = 0$, for $j, h = 1, \ldots, n - 2k$ and $j \neq h$.

On the other hand, let us assume that P has only zeros on the main diagonal. Then, in order to have $P = UU^\top$, each row of matrix U should have an even number of nonzeros. Hence the sum of the rows of U is null over $GF[2]$ and so U would not have full rank. This leads to a contradiction since P is a permutation matrix and has full rank. $\qquad\square$

Theorem 9 *Given an $n \times n$ symmetric matrix A whose rank over $GF[2]$ is r, and with at least one nonzero entry on the main diagonal, there exists an $n \times r$ matrix U such that $A = UU^\top$.*

Proof. By Lemma 7 there exists an $n \times n$ matrix S such that $SAS^\top = B$, where B contains an $r \times r$ symmetric permutation matrix P. If A has at least a 1 on the main diagonal, then, by construction, also P has at least a 1 on the main diagonal. Indeed the algorithm described in the proof of Lemma 7, reduces A to B by means of symmetrical multiplication by either permutation or M_{ij} matrices. Symmetric permutation does not change the overall number of nonzero entries on the main diagonal, while the product by M_{ij} matrices produces the effect $a_{i,i} \leftarrow a_{i,i} + a_{j,j}$. Thus if at least one of $a_{i,i}$ and $a_{j,j}$ is different from zero, then the same will hold after the multiplication by M_{ij} matrices.

Since S is invertible we can write $A = S^{-1} B S^{-T}$ and then $A = V P V^T$, where V is the principal $n \times r$ submatrix of S^{-1}. Since P has at least a nonzero diagonal element we can apply Lemma 8 and claim that there exists an $r \times r$ matrix W such that $P = WW^T$. Thus we obtain $A = VPV^T = VWW^TV^T = (VW)(VW)^T$, where $U = VW$ is an $n \times r$ matrix. $\qquad\square$

5 A family of graphs with extremal properties

Let I_k, J_k, and P_k denote the identity matrix, the matrix with all the entries equal to 1, and the matrix with the $(i, k-i)$-th entries equal to 1, respectively, all of size k. Let us consider the following $n \times n$ matrix, for $n = 4k$, written in block form.

$$A_n = I_n + B_n = \begin{pmatrix} I_k & I_k & I_k & J_k - P_k \\ I_k & I_k & J_k - I_k & P_k \\ I_k & J_k - I_k & I_k & P_k \\ J_k - P_k & P_k & P_k & I_k \end{pmatrix}$$

This is a family of symmetric matrices with the following properties:

- A_n is triangle-free. This property can be easily verified computing the trace of B_n^3. We have

$$B_n^2 = \begin{pmatrix} (k-2)J_k + 3I_k & 2(J_k - I_k) & 2(J_k - I_k) & 2P_k \\ 2(J_k - I_k) & (k-2)J_k + 3I_k & J_k - I_k & 2(J_k - P_k) \\ 2(J_k - I_k) & J_k - I_k & (k-2)J_k + 3I_k & 2(J_k - P_k) \\ 2P_k & 2(J_k - P_k) & 2(J_k - P_k) & (k-2)J_k + 3I_k \end{pmatrix},$$

  from which we readily see that $\mathrm{Tr}(B_n^3) = 4\mathrm{Tr}(6J_k - 6I_k) = 0$.

- $\mathrm{Rank}_2(A) = r = n/4 + 2$. Indeed the matrices A_n can be obtained as $A_n = UU^T$, where U^T is the following $r \times n$ matrix.

$$\left(\begin{array}{cccc|c|cccc|c|cccc|cccc|c} 1 & 0 & \cdots & 0 & 0 & 1 & \cdots & 1 & 0 & 1 & \cdots & 1 & 1 & \cdots & 1 & 0 \\ 1 & 0 & \cdots & 0 & 1 & 0 & \cdots & 0 & 0 & 1 & \cdots & 1 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 & 1 & \cdots & 1 & 1 & 0 & \cdots & 0 & 0 & \cdots & 0 & 1 \\ \hline 0 & & & & 0 & & & & 0 & & & & & & & 1 \\ \vdots & & I_{r-3} & & \vdots & & I_{r-3} & & \vdots & & I_{r-3} & & & J_{r-3} - P_{r-3} & & \vdots \\ 0 & & & & 0 & & & & 0 & & & & & & & 1 \end{array} \right). \qquad (9)$$

- A_n is regular of degree r.

- B_n has independent sets of size $k = r - 2$.

- $\mathrm{Rank}_{\mathbf{R}}(A_n) = n$, and A_n has exactly 5 distinct integer eigenvalues. More precisely A_n has the following eigensystem.

  - $\lambda_1 = r$ with multiplicity 1 and eigenvector $(1, \cdots, 1)^T$.
  - $\lambda_2 = r - 4$ with multiplicity 1 and eigenvector

$$(\overbrace{-1, \cdots, -1}^{k}, \overbrace{1, \cdots, 1}^{k}, \overbrace{1, \cdots, 1}^{k}, \overbrace{-1, \cdots, -1}^{k})^T.$$

– $\lambda_3 = 4 - r$ with multiplicity 2 and eigenvectors

$$(\overbrace{-1,\cdots,-1}^{k},\overbrace{0,\cdots,0}^{k},\overbrace{0,\cdots,0}^{k},\overbrace{1,\cdots,1}^{k})^T$$

and

$$(\overbrace{-1,\cdots,-1}^{k},\overbrace{1,\cdots,1}^{k},\overbrace{-1,\cdots,-1}^{k},\overbrace{1,\cdots,1}^{k})^T.$$

– $\lambda_4 = -2$ with multiplicity $k-1$ and eigenvectors

$$\left(\begin{array}{c|c|c|c|c|c|c}\begin{matrix}1\\\vdots\\1\end{matrix} & -P_{k-1} & \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & -I_{k-1} & \begin{matrix}1\\\vdots\\1\end{matrix}\end{array}\right)^T.$$

– $\lambda_5 = 2$ with multiplicity $n - r - 1 = 3(k-1)$ and eigenvectors

$$\left(\begin{array}{c|c|c|c|c|c}-I_{k-1} & \begin{matrix}1\\\vdots\\1\end{matrix} & 0 & 0 & \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1}\\\hline \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & 0 & \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & 0\\\hline \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & \begin{matrix}-1\\\vdots\\-1\end{matrix} & P_{k-1} & 0 & 0\end{array}\right)^T.$$

The verification of the above properties can now be done by direct inspection.

6   A sunflower theorem related to the conjecture

The decomposition of a symmetric matrix A over GF_2 given in Theorem 9 can be interpreted as representing A as an intersection matrix as follows. The rows and columns are indexed by sets of some family of subsets of $\{1,\ldots,r\}$, where r is the rank of A. The (i,j)-th entry of A is 1 iff the intersection of the index sets corresponding to row i and column j is odd. The rows of U are the characteristic vectors of the sets. We shall call set systems also *hypergraphs*. If all sets have size k, then we speak of k-*hypergraphs*.

Let us state our conjecture for symmetric matrices over GF_2 in terms of set systems. The entries on the main diagonal are all equal to one, and this implies that the sizes of the sets must be odd. A triangle corresponds to a triple of sets in the family such that every two intersect in an odd set. Thus the conjecture can be rephrased as: *there exists a constant K such that for any family of odd subsets of $\{1,\ldots,r\}$ of size $\geq Kr$ there exists a triple of sets in the family such that every two intersect in an odd set.*

One of the special properties of the matrices considered in the previous section is that the sets in this representation have only four sizes, i.e., $1, 3, r-3, r-1$. This raises the question of whether it

is possible to prove the conjecture under such restriction on sizes of the representing sets. Notice that it is actually sufficient to prove it only for size equal to 3. We prove below a more general result which implies that the conjecture holds for any constant size (Theorem 10).

Definition 1 *A sunflower with l petals and core Y is a family of sets $X_1, \ldots, X_l$ such that $X_i \cap X_j = Y$ for every $i \neq j$.*

Note that the assumption that $X_1, \ldots, X_l$ form a sunflower with an odd core is stronger than the assumption that every two sets intersect in an odd set. A classical result of Erdös and Radó [2] states that for a given k and l, every sufficiently large k-hypergraph contain a sunflower with l petals. It has been observed in [4] that a k-hypergraph, k odd, on an n element set of vertices with at least $n^{(k-1)/2} l^{(k+1)/2} \frac{1 \cdot 3 \cdots k}{2 \cdot 4 \cdots (k-1)}$ edges contains a sunflower with l petals and with an *even* core. A k-hypergraph can be easily constructed which shows that this bound cannot be essentially improved. Our result shows that there is a big difference between the cases with *even* and *odd* cores, i.e., for the latter a linear number of edges suffices to guarantee the presence of a sunflower.

Theorem 10 *For every positive integers k, l, with k odd and $l \geq 3$, there exists an integer K such that for every $n \geq 1$ and every k-hypergraph H on n vertices with at least Kn edges, there exists a sunflower in H with l petals and an odd core.*

Proof. Let H be a k-hypergraph on V, $|V| = n$, i.e., H is a set of k element subsets of V. We shall suppose that H does not contain a sunflower with l petals and an odd core and bound the number of edges by a linear function of n.

Claim. Let d be a constant depending only on k and l. There exists a directed graph G on V with the following properties:

1. G has no self-loops;

2. the outdegree of G is bounded by d;

3. for every $h \in H$ and $v \in h$, there exists $u \in h$ such that the edge $v \to u$ is in G.

Proof: Take $d = (l - 1)(k - 1)$. For each $v \in V$, we define outgoing edges as follows. Let $H_v = \{h - \{v\}; \ h \in H, \ v \in h\}$, so that H_v is a $(k - 1)$-hypergraph on $V - \{v\}$. As H does not contain a sunflower with l petals and a one-element core, H_v cannot contain more than $l - 1$ disjoint edges. Thus there exists a set of size at most d such that every edge of H_v intersects the set. We put in G an arc $v \to u$ for every element u of this set. The three conditions are clearly satisfied.

We shall prove the theorem by induction on k, where k ranges over the odd numbers. For $k = 1$ the statement is trivial, as there are only n one element sets. So suppose $k \geq 3$. Fix a graph G with the properties stated in the above claim. Consider the induced graph on $h \in H$. It contains at least one *terminal component*, which is a transitive subset (i.e., each point is reachable from another one by an oriented path) and no point outside is reachable from the set. Such a set has at least two elements, since there are no self-loops in G. We shall choose one such terminal component for each edge and call it *the nucleus* of the edge.

13

We now define another auxiliary graph F. It is a symmetric graph with vertex set U consisting of all transitive subsets in G of size at most k, including one-element sets. Two vertices are connected by an edge, if they are different and have nonempty intersection.

Claim. The maximal degree of F can be bounded by a constant d' depending only on k and l.

Proof: Clearly, each $v \in V$ is contained in at most $2^{1+d+\dots+d^k}$ transitive components of size $\leq k$. Thus we can bound the degree of F by $k2^{1+d+\dots+d^k}$.

Now we construct a hypergraph J on U. For each edge $e \in H$ we put an edge e' in J consisting of the nucleus of e and one element sets $\{v\}$ for each $v \in e$ which is not in the nucleus. Note that there is a one-to-one correspondence between the edges of H and J.

Claim. There exists $\varepsilon > 0$, depending only on k and l such that one can choose $J' \subseteq J$ such that $|J'| \geq \varepsilon |J|$ and $\bigcup J'$ is an independent set of F.

Proof: For each edge of F, choose at random and independently one of the two possible orientations. Let X consist of the vertices v all of whose incident edges are oriented into them. Clearly, X is an independent set of F. It suffices to show that the expected number of edges of J which are contained in X is a positive fraction of all the edges of F. For a vertex v, the probability that $v \in X$ is at least $2^{-d'}$. An edge $e \in J$ does not contain any edge of F, thus the probabilities for the vertices in e are independent. Hence e is contained in X with probability at least $2^{-d'|e|}$. By linearity of expectation we get that at least this fraction of edges of J is in X on the average.

Each edge of J' consists of a nucleus, whose size is at least two, and one element sets. We can assume that all terminal components have the same size, say r, as there are only constantly many sizes.

Let $H' \subseteq H$ be the hypergraph consisting of the edges corresponding to the edges J'. The construction of J' gives us the following property of H'. For every $e \in H'$ and the nucleus C of e, if $f \in H'$ is another edge, then either C is also the nucleus of f or C is disjoint from f.

We now consider two cases.

1. r is odd. Replace each of the nuclei by a single vertex, thus obtaining a $(k - r + 1)$-hypergraph H'' with the same number of edges. A sunflower with a core of size s in H'' corresponds to a sunflower with the same number of petals and a core of size either s or $s + r - 1$ in H' (depending on whether or not a nucleus is in the core). Thus H'' does not contain a sunflower with l petals and an odd core. By the induction assumption, $|H''|$ must be bounded by a linear function of n, hence so is $|H|$.

2. r is even. We construct H'' by deleting the chosen terminal components. Thus H'' is a $(k - r)$-hypergraph. The number of edges in H'' can be smaller now, but at most by the constant factor $l - 1$, since an edge resulting from m edges of H' is the core of a sunflower with m petals in H'. We shall show that H'' does not contain a sunflower with $(l - 1)^2 + 1$ petals and an odd core. This will complete the proof, as by the induction assumption it implies a linear upper bound on the size of H''.

Suppose that H'' does contain a sunflower with $(l - 1)^2 + 1$ petals and a core of an odd size s. Consider the corresponding $(l - 1)^2 + 1$ edges in H'. Some edges may share the nuclei. There are two possibilities, either at least l edges share the same nucleus, or there are l edges with different nuclei. In the first case we get a sunflower with l petals and a core of size $s + r$; in the second case

a sunflower with l petals and a core of size r. This is a contradiction with the assumption on H. □

Note that the version with even k and even cores follows immediately.

We now present a more general result, which applies to factorizations of the form UV, and which implies the Triangle Conjecture for matrices that can be decomposed in the form UV, where both U and V satisfy some *sparsity* requirements. More precisely we have the following Theorem.

Theorem 11 $\forall k, l, \exists \varepsilon > 0$ *such that* $\forall n$, *if* $M = AB$ *(over $GF[2]$), where M is an $n \times n$ matrix with ones on the main diagonal, A is an $n \times r$ matrix, B is an $r \times n$ matrix, $r \leq \varepsilon n$, A (resp. B) has at most k ones in each column (resp. row), then there exists in M an $l \times l$ principal submatrix of ones.*

Remark. Note that the matrix M above does not need to be symmetric.

Theorem 11 can be reformulated in an equivalent way, in set intersection terms. Indeed we have the following.

Theorem 12 $\forall k, l, \exists \varepsilon > 0$, $\forall n$, *and for all sets* $(A_1, B_1), \ldots, (A_n, B_n)$, $|A_i|, |B_i| \leq k$, $A_i, B_i \subseteq X$, $|X| = r \leq \varepsilon n$, *where for all i, if $|A_i \cap B_i|$ is odd, then there exist $i_1, \ldots, i_l$, such that for all $1 \leq \alpha, \beta \leq l$, with $\alpha \neq \beta$, we have that $|A_{i_\alpha} \cap B_{i_\beta}|$ is odd.*

Theorem 12 follows from a stronger result, which we prove below.

Theorem 13 $\forall k, l, \exists \varepsilon > 0$, $\forall n$, *and for all sets* $(A_1, B_1), \ldots, (A_n, B_n)$, $|A_i|, |B_i| \leq k$, $A_i, B_i \subseteq X$, $|X| = r \leq \varepsilon n$, *if for all i, $|A_i \cap B_i|$ is odd, then there exists a set D of odd cardinality such that for all $1 \leq \alpha, \beta \leq l$, with $\alpha \neq \beta$, we have that $A_{i_\alpha} \cap B_{i_\beta} = D$.*

Proof. Let $C_i = A_i \cap B_i$, i.e., $|C_i|$ is odd. By Theorem 10, we have that there exist $j_1, \ldots, j_m$ and D such that $\forall \alpha \neq \beta$, $C_{j_\alpha} \cap C_{j_\beta} = D$.

Now, for every i, choose the mappings

$f_i : P(A_i) \rightarrow \{0, \ldots, 2^k - 1\}$,

$g_i : P(B_i) \rightarrow \{0, \ldots, 2^k - 1\}$,

and assign the colour $(f_i(A_i \cap B_j), g_j(A_i \cap B_j))$ to the pair (i, j), for $i < j$. By Ramsey Theorem, there exists $\{i_1, \ldots, i_l\} \subseteq \{j_1, \ldots, j_m\}$ such that all pairs have the same colour.

Claim: $\forall i, i', j, j' \in \{i_1, \ldots, i_l\}, i < j, i' < j'$, we have $A_i \cap B_j = A_{i'} \cap B_{j'}$.

Proof: $A_i \cap B_j = A_i \cap B_{j'} = A_{i'} \cap B_{j'}$.

Thus there exists a set D' such that $\forall i, j \in \{i_1, \ldots, i_l\}$, $i < j$, $A_i \cap B_j = D'$. Symmetrically, there exists D'' such that $\forall i, j \in \{i_1, \ldots, i_l\}$, $i < j$, $A_i \cap B_j = D''$. Since $\forall i, D' \subseteq A_i$ and $D' \subseteq B_j$, we have $D' \subseteq D$. But also $\forall i, D \subseteq A_i$, and $D \subseteq B_j$, and hence $D \subseteq A_i \cap B_j = D'$ (for $i < j$). Therefore $D = D'$. By symmetry $D = D''$.

□

7 Conclusions and more open problems

Problem 1 *Let M be a matrix with ones on the main diagonal and without $[2,2]$ configurations. How many rows do we have to remove to reduce its rank by one?*

If, for a given field, this number can be bounded by a constant, then the Triangle Conjecture is true.

What if the Triangle Conjecture is false? We think that the approach of relating local properties of the graph of nonzero elements can still be used, if weaker properties are used. Let us call an *odd alternating cycle* an oriented graph which is a cycle, when the orientation is forgotten, and the orientation of the arrows on the cycle alternates with one exception (put otherwise, there is a vertex v on the cycle such that if we go around the cycle from v to v, the orientation of the edges alternates). Thus a transitive triangle is an alternating cycle of length 3. For some applications (e.g., the lower bound on the rigidity of Toeplitz matrices with indeterminates) it would suffice to use odd alternating cycles of length bounded by a constant, instead of just transitive triangles. There may be other subgraphs having similar properties.

We do not know if Theorem 3 can be proved only using the Triangle Conjecture. It is even conceivable that the symmetric case of the conjecture could suffice.

Theorem 10 gives only a restricted version of the Triangle Conjecture, namely for matrices which can be decomposed into $UU^\top$, where U has only a constant number of nonzero elements in each row. The restriction on the number of nonzero elements seems to be too severe for the theorem to be useful for intended applications. That may be true for proving lower bounds on rigidity. For proving lower bounds on the circuit size, however, we may use some additional restrictions on the number of nonzero elements in the matrices involved. Low rigidity of a matrix M means that M can be decomposed into $A + B$ where A is sparse and B is of low rank. B being of rank less than r means that B can be further decomposed into $B = C \times D$ where C, D are some matrices of size $n \times r$ and $r \times n$, respectively. A closer look at Valiant's reduction reveals that, in the case of a transformation computed by a small circuit, one can find a decomposition such that C is sparse. In the case of series-parallel circuits of linear size and logarithmic depth, the restriction on C is particularly strong, namely there is a constant bound on the number of nonzero elements in each row of C. This type of restriction is similar to ours.

References

[1] N. Alon, M. Szegedy, *Large sets of nearly orthogonal vectors,* preprint.

[2] P. Erdös and R. Rado, *Intersection theorems for systems of sets,* J. London Math. Soc. 35, 1960, 85-90

[3] J. Friedman, *A note on matrix rigidity,* Combinatorica 13(2), (1993), 235-239

[4] J. Håstad, S. Jukna and P. Pudlák, *Top-down lower bounds for depth-three circuits,* Computational Complexity 5, 1995, 99-112

[5] A. Kotlov, L. Lovász, *The rank and size of graphs,* J. of Graph Theory 23(1) (1996), 185-189.

[6] N. Nisan and A. Wigderson, *On Rank vs Communication Complexity*, Proc. 35th IEEE FOCS (1994), 831-836.

[7] P. Pudlák, V. Rödl, J. Sgall, *Boolean circuits, tensor ranks and communication complexity*, SIAM J. on Computing, to appear.

[8] A.A. Razborov, *On rigid matrices,* preprint (in Russian)

[9] A.A. Razborov, *Lower bounds on the monotone complexity of some boolean functions,* Soviet Math. Doklady 31, (1985) 354-357.

[10] R. Raz, B. Spiker, *On the log-rank conjecture in communication complexity*, Proc. 34th IEEE FOCS (1993), 168-176.

[11] M. Rosenfeld, *Almost orthogonal lines in E^d*, DIMACS Series in Discrete Math. 4 (1991), 489-492.

[12] L.G. Valiant, *Graph-theoretic arguments in low level complexity,* Proc. 6th MFCS, Springer-Verlag LNCS 53 (1977), 162-176.