

Min-Rank Conjecture for Log-Depth Circuits [☆]

Stasys Jukna^{a,*}, Georg Schnitger^{b,1}

^a*Institute of Mathematics and Computer Science, Akademijos 4, LT-80663 Vilnius, Lithuania*

^b*University of Frankfurt, Institut of Computer Science, Robert-Mayer-Str. 11-5, D-60054 Frankfurt, Germany.*

Abstract

A completion of an m -by- n matrix A with entries in $\{0, 1, *\}$ is obtained by setting all $*$ -entries to constants 0 and 1. A system of semi-linear equations over GF_2 has the form $M\mathbf{x} = f(\mathbf{x})$, where M is a completion of A and $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is an operator, the i th coordinate of which can only depend on variables corresponding to $*$ -entries in the i th row of A . We conjecture that no such system can have more than $2^{n-\epsilon \cdot \text{mr}(A)}$ solutions, where $\epsilon > 0$ is an absolute constant and $\text{mr}(A)$ is the smallest rank over GF_2 of a completion of A . The conjecture is related to an old problem of proving super-linear lower bounds on the size of log-depth boolean circuits computing linear operators $\mathbf{x} \mapsto M\mathbf{x}$. The conjecture is also a generalization of a classical question about how much larger can non-linear codes be than linear ones. We prove some special cases of the conjecture and establish some structural properties of solution sets.

Key words: Boolean circuits; Partial matrix; Matrix completion; Min-rank; Matrix rigidity; Sum-Sets; Cayley graphs; Error-correcting codes

1. Introduction

One of the challenges in circuit complexity is to prove a super-linear lower bound for log-depth circuits over $\{\&, \vee, \neg\}$ computing an explicitly given boolean operator $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Attempts to solve it have led to several weaker problems which are often of independent interest. The problem is open even if we impose an additional restriction that the depth of the circuit is $O(\log n)$. It is even open for *linear* log-depth circuits, that is, for log-depth circuits over the basis $\{\oplus, 1\}$, in spite of the apparent simplicity of such circuits. It is clear that the operators computed by linear circuits must also be linear, that is, be matrix-vector products $\mathbf{x} \rightarrow M\mathbf{x}$ over the field $GF_2 = (\{0, 1\}, \oplus, \cdot)$,

An important result of Valiant [27] reduces the lower bounds problem for log-depth circuits over $\{\&, \vee, \neg\}$ to proving lower bounds for certain depth-2 circuits, where we allow *arbitrary* boolean functions as gates.

1.1. Reduction to depth-2 circuits

A depth-2 circuit of *width* w has n boolean variables x_1, \dots, x_n as input nodes, w arbitrary boolean functions h_1, \dots, h_w as gates on the middle layer, and m arbitrary boolean functions g_1, \dots, g_m as

[☆]Research of both authors supported by a DFG grant SCHN 503/4-1.

*Corresponding author

Email addresses: jukna@thi.informatik.uni-frankfurt.de (Stasys Jukna),
georg@thi.informatik.uni-frankfurt.de (Georg Schnitger)

¹Currently at University of Frankfurt, Institut of Computer Science, Frankfurt, Germany.

gates on the output layer. Direct input-output wires, connecting input variables with output gates, are allowed. Such a circuit computes an operator $f = (f_1, \dots, f_m) : \{0, 1\}^n \rightarrow \{0, 1\}^m$ if, for every $i = 1, \dots, m$,

$$f_i(\mathbf{x}) = g_i(\mathbf{x}, h_1(\mathbf{x}), \dots, h_w(\mathbf{x})).$$

The *degree* of such a circuit is the maximum, over all output gates g_i , of the number of wires going directly from input variables x_1, \dots, x_n to the gate g_i . That is, we ignore the wires incident with the gates on the middle layer. Let $\deg_w(f)$ denote the smallest degree of a depth-2 circuit of width w computing f .

It is clear that $\deg_n(f) = 0$ for $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$: just put the functions f_1, \dots, f_n on the middle layer. Hence, this parameter is only nontrivial for $w < n$. Especially interesting is the case when $w = O(n/\ln \ln n)$ (see also Theorem 2.2 in [20] for more details):

Lemma 1.1 (Valiant [27]). *If $\deg_w(f) = n^{\Omega(1)}$ for $w = O(n/\ln \ln n)$, then the operator f cannot be computed by a circuit of depth $O(\ln n)$ using $O(n)$ constant fan-in gates.*

Recently, there was a substantial progress in proving lower bounds on the *size* of (that is, on the total number of wires in) depth-2 circuits. Superlinear lower bounds of the form $\Omega(n \log^2 n)$ were proved using graph-theoretic arguments by analyzing some superconcentration properties of the circuit as a graph [6, 14, 15, 18, 16, 2, 20, 21, 22]. Higher lower bounds of the form $\Omega(n^{3/2})$ were proved using information theoretical arguments [4, 9]. But the highest known lower bound on the *degree* of width w circuits has the form $\Omega((n/w) \ln(n/w))$ [20], and is too weak to have a consequence for log-depth circuits.

A natural question therefore was to improve the lower bound on the degree at least for *linear* circuits, that is, for depth-2 circuits whose middle gates as well as output gates are linear boolean functions (parities of their inputs). Such circuits compute linear operators $\mathbf{x} \mapsto M\mathbf{x}$ for some $(0, 1)$ -matrix M ; we work over GF_2 . By Valiant's reduction, this would give a super-linear lower bound for log-depth circuits over $\{\oplus, 1\}$.

This last question attracted attention of many researchers because of its relation to a purely algebraic characteristic of the underlying matrix M —its rigidity. The *rigidity* $\mathcal{R}_M(r)$ of a $(0, 1)$ -matrix M is the smallest number of entries of M that must be changed in order to reduce its rank over GF_2 to r . It is not difficult to show (see [27]) that any linear depth-2 circuit of width w computing $M\mathbf{x}$ must have degree at least $\mathcal{R}_M(w)/n$: If we set all direct input-output wires to 0, then the resulting degree-0 circuit will compute some linear transformation $M'\mathbf{x}$ where the rank of M' does not exceed the width w . On the other hand, M' differs from M in at most dn entries, where d is the degree of the original circuit. Hence, $\mathcal{R}_M(w) \leq dn$ from which $d \geq \mathcal{R}_M(w)/n$ follows.

Motivated by its connection to proving lower bounds for log-depth circuits, matrix rigidity (over different fields) was considered by many authors, [23, 1, 17, 7, 16, 20, 25, 24, 10, 11, 19, 26] among others. It is therefore somewhat surprising that the highest known lower bounds on $\mathcal{R}_M(r)$ (over the field GF_2), proved in [7, 25] also have the form $\Omega((n^2/r) \ln(n/r))$, resulting to the same lower bound $\Omega((n/w) \ln(n/w))$ on the degree of linear circuits as that for general depth-2 circuits proved in [20]. This phenomenon is particularly surprising, because general circuits may use *arbitrary* (not just linear) boolean functions as gates. We suspect that the absence of higher lower bounds for linear circuits than those for non-linear ones could be not just a coincidence.

Conjecture 1 (Linearization conjecture for depth-2 circuits). *Depth-2 circuits can be linearized. That is, every depth-2 circuit computing a linear operator can be transformed into an equivalent linear depth-2 circuit without substantial increase of its width or its degree.*

If true, the conjecture would have important consequences for log-depth circuits. Assuming this conjecture, any proof that every depth-2 circuit of width $w = O(n/\ln \ln n)$ with unbounded fan-in parity gates for a given linear operator $M\mathbf{x}$ requires degree $n^{\Omega(1)}$ would imply that $M\mathbf{x}$ requires a super-linear number of gates in any log-depth circuit over $\{\&, \vee, \neg\}$. In particular, this would mean that proving high lower bounds on matrix rigidity is a much more difficult task than assumed before: such bounds would yield super-linear lower bounds for log-depth circuits over a general basis $\{\&, \vee, \neg\}$, not just for circuits over $\{\oplus, 1\}$.

As the first step towards Conjecture 1, in this paper we relate it to a purely combinatorial conjecture about partially defined matrices—the *min-rank conjecture*, and prove some results supporting this last conjecture. This turns the problem about the linearization of depth-2 circuits into a problem of Combinatorial Matrix Theory concerned with properties of completions of partially defined matrices (see, e.g., the survey [8]). Hence, the conjecture may also be of independent interest.

Unfortunately, we were not able to prove the conjecture in its full generality. So far, we are only able to prove that some of its special cases are true. This is not very surprising because the conjecture touches a basic problem in circuit complexity: Can non-linear gates help to compute linear operators? This paper is just the first step towards this question.

1.2. The Min-Rank Conjecture

A *completion* of a $(0, 1, *)$ -matrix A is a $(0, 1)$ -matrix M obtained from A by setting all $*$'s to constants 0 and 1. A *canonical completion* of A is obtained by setting all $*$'s in A to 0.

If A is an m -by- n matrix, then each its completion M defines a linear operator mapping each vector $\mathbf{x} \in \{0, 1\}^n$ to a vector $M\mathbf{x} \in \{0, 1\}^m$. Besides such (linear) operators we also consider general ones. Each operator $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ can be looked at as a sequence $G = (g_1, \dots, g_m)$ of m boolean functions $g_i : \{0, 1\}^n \rightarrow \{0, 1\}$.

We say that an operator $G = (g_1, \dots, g_m)$ is *consistent* with an m -by- n $(0, 1, *)$ -matrix $A = (a_{ij})$ if the i th boolean function g_i can only depend on those variables x_j for which $a_{ij} = *$. That is, the i th component g_i of G can only depend on variables on which the i th row of A has stars (see Example 1.6).

Definition 1.2. With some abuse in notation, we call a set $L \subseteq \{0, 1\}^n$ a *solution* for a partial matrix A if there is a completion M of A and an operator G such that G is consistent with A and $M\mathbf{x} = G(\mathbf{x})$ holds for all $\mathbf{x} \in L$. A solution L is *linear* if it forms a linear subspace of $\{0, 1\}^n$ over GF_2 .

That is, a solution for A is a set L of $(0, 1)$ -vectors of the form $L = \{\mathbf{x} : M\mathbf{x} = G(\mathbf{x})\}$, where M is a completion of A , and G is an operator consistent with A . A solution L is linear, if $\mathbf{x} \oplus \mathbf{y} \in L$ for all $\mathbf{x}, \mathbf{y} \in L$.

Since, besides the consistency, there are no other restrictions on the operator G in the definition of the solution L , we can always assume that M is the canonical completion of A (with all stars set to 0).

Observation 1.3 (Canonical completions). *If $L = \{\mathbf{x} : M\mathbf{x} = G(\mathbf{x})\}$ is a solution for A , and M' is the canonical completion of A , then there is an operator G' such that G' is consistent with A and $L = \{\mathbf{x} : M'\mathbf{x} = G'(\mathbf{x})\}$.*

Proof. The i th row \mathbf{m}_i of M must have the form $\mathbf{m}_i = \mathbf{m}'_i + \mathbf{p}_i$, where $\mathbf{m}'_i \in \{0, 1\}^n$ is the i th row of the canonical completion M' of A , and $\mathbf{p}_i \in \{0, 1\}^n$ is a vector with no 1's in positions where the i th row of A has no stars. We can then define an operator $G' = (g'_1, \dots, g'_m)$ by $g'_i(\mathbf{x}) := g_i(\mathbf{x}) \oplus \langle \mathbf{p}_i, \mathbf{x} \rangle$. (As customary, the scalar product of two vectors $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$ over GF_2 is $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i \pmod{2}$.) Since G was consistent with A , the new operator G' is also consistent with A . Moreover, for every vector $\mathbf{x} \in \{0, 1\}^n$, we have that $\langle \mathbf{m}_i, \mathbf{x} \rangle = g_i(\mathbf{x})$ iff $\langle \mathbf{m}'_i, \mathbf{x} \rangle = g'_i(\mathbf{x})$. \square

We are interested in how much the maximum $\text{opt}(A) = \max_L |L|$ over all solutions L for A can exceed the maximum $\text{lin}(A) = \max_L |L|$ over all linear solutions L for A . It can be shown (Corollary 6.3 below) that

$$\text{lin}(A) = 2^{n-\text{mr}(A)},$$

where $\text{mr}(A)$ is the *min-rank* of A defined as the smallest possible rank of its completion:

$$\text{mr}(A) = \min\{\text{rk}(M) : M \text{ is a completion of } A\}.$$

If we only consider *constant* operators G , that is, operators with $G(\mathbf{x}) = \mathbf{b}$ for some $\mathbf{b} \in \{0, 1\}^m$ and all $\mathbf{x} \in \{0, 1\}^n$, then Linear Algebra tells us that no solution for A can have more than 2^{n-r} vectors, where $r = \text{rk}(M)$ is the rank (over GF_2) of the canonical completion M of A , obtained by setting all stars to 0.

If we only consider *affine* operators G , that is, operators of the form $G(\mathbf{x}) = H\mathbf{x} \oplus \mathbf{b}$ where H is an m -by- n $(0, 1)$ -matrix, then no solution for A can have more than $2^{n-\text{mr}(A)}$ vectors, because then the consistency of $G(\mathbf{x})$ with A ensures that, for every completion M of A , the matrix $M \oplus H$ is a completion of A as well.

Remark 1.4. This last observation implies, in particular, that $\text{opt}(A) \leq 2^{n-\text{mr}(A)}$ for all $(0, 1, *)$ -matrices A with at most one $*$ in each row: In this case each g_i can depend on at most one variable, and hence, must be a linear boolean function.

We conjecture that a similar upper bound also holds for *any* operator G , as long as it is consistent with A . That is, we conjecture that linear operators are almost optimal.

Conjecture 2 (Min-Rank Conjecture). *There exists a constant $\epsilon > 0$ such that for every m -by- n $(0, 1, *)$ -matrix A we have that $\text{opt}(A) \leq 2^{n-\epsilon \cdot \text{mr}(A)}$ or, equivalently,*

$$\text{opt}(A) \leq 2^n \left(\frac{\text{lin}(A)}{2^n} \right)^\epsilon. \quad (1)$$

Remark 1.5. To have consequences for log-depth circuits, it would be enough, by Lemma 1.1, that the conjecture holds at least for $\epsilon = o(1/\log \log n)$.

Example 1.6. To illustrate the introduced concepts, let us consider the following system of 3 equations in 6 variables:

$$\begin{aligned} x_1 \oplus x_6 &= x_3 \cdot x_5 \\ x_2 \oplus x_3 \oplus x_4 &= x_1 \cdot (x_5 \oplus x_6) \\ x_4 &= (x_2 \oplus x_5) \cdot (x_3 \oplus x_6). \end{aligned} \quad (2)$$

The corresponding $(0, 1, *)$ -matrix for this system is

$$A = \begin{pmatrix} 1 & 0 & * & 0 & * & 1 \\ * & 1 & 1 & 1 & * & * \\ 0 & * & * & 1 & * & * \end{pmatrix}, \quad (3)$$

and the system itself has the form $M\mathbf{x} = G(\mathbf{x})$, where M is the canonical completion of A :

$$M = \begin{pmatrix} 1 & 0 & \underline{0} & 0 & \underline{0} & 1 \\ \underline{0} & 1 & \underline{1} & 1 & \underline{0} & \underline{0} \\ \underline{0} & \underline{0} & \underline{0} & 1 & \underline{0} & \underline{0} \end{pmatrix},$$

and $G = (g_1, g_2, g_3) : \{0, 1\}^6 \rightarrow \{0, 1\}^3$ is an operator with

$$\begin{aligned} g_1(\mathbf{x}) &= x_3 \cdot x_5; \\ g_2(\mathbf{x}) &= x_1 \cdot (x_5 \oplus x_6); \\ g_3(\mathbf{x}) &= (x_2 \oplus x_5) \cdot (x_3 \oplus x_6). \end{aligned}$$

The min-rank of A is equal 2, and is achieved by the following completion:

$$M' = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & \underline{1} & \underline{1} & 1 & \underline{0} & \underline{0} \end{pmatrix}.$$

1.3. Our results

In Section 2 we prove the main consequence of the min-rank conjecture for boolean circuits: If true, it would imply that non-linear gates are powerless when computing linear operators $M\mathbf{x}$ by depth-2 circuits (Lemmas 2.2 and 2.3).

In Sections 3 and 4 we prove some partial results supporting Conjectures 1 and 2. We first show (Corollary 3.4) that every depth-2 circuit of width w computing a linear operator can be transformed into an equivalent *linear* depth-2 circuit of the same degree and width at most w plus the maximum number of wires in a matching formed by the input-output wires of the original circuit.

We then prove two special cases of Min-Rank Conjecture. A set of $(0, 1, *)$ -vectors is *independent* if they cannot be made linearly dependent over GF_2 by setting stars to constants 0 and 1. If A is a $(0, 1, *)$ -matrix, then the upper bound $\text{opt}(A) \leq 2^{n-r}$ holds if the matrix A contains r independent columns (Theorem 4.4). The same upper bound also holds if A contains r independent rows, and the sets of star positions in these rows form a chain with respect to set-inclusion (Theorem 4.11).

After that we concentrate on the *structure* of solutions. In Section 5 we show that solutions for a $(0, 1, *)$ -matrix A are precisely independent sets in a Cayley graph over the Abelian group $(\{0, 1\}^n, \oplus)$ generated by a special set $K_A \subseteq \{0, 1\}^n$ of vectors defined by the matrix A (Theorem 5.2).

In Section 6 we first show that every linear solution for A lies in the kernel of some completion of A (Theorem 6.2). This, in particular, implies that $\text{lin}(A) = 2^{n-\text{mr}(A)}$ (Corollary 6.3), and gives an alternative definition of the min-rank $\text{mr}(A)$ as the smallest rank of a boolean matrix H such that $H\mathbf{x} \neq \mathbf{0}$ for all $\mathbf{x} \in K_A$ (Corollary 6.4). In Section 7 we show that non-linear solutions must be “very non-linear”: they cannot contain linear subspaces of dimension exceeding the maximum number of $*$ ’s in a row of A (Theorem 7.1).

In Section 8 we consider the relation of the min-rank conjecture with error-correcting codes. We define $(0, 1, *)$ -matrices A , the solutions for which are error-correcting codes, and show that the min-rank conjecture for these matrices is true: In this case the conjecture is implied by well known lower and upper bounds on the size of linear and nonlinear error correcting codes (Lemma 8.3).

For readers convenience, we summarize the introduced concepts at the end of the paper (see Table 1).

2. Min-rank conjecture and depth-2 circuits

Let F be a depth-2 circuit computing a linear operator $\mathbf{x} \rightarrow M\mathbf{x}$, where M is an m -by- n $(0, 1)$ -matrix. Say that the (i, j) th entry of M is *seen* by the circuit, if there is a direct wire from x_j to the i th

output gate. Replace all entries of M seen by the circuit with $*$'s, and let A_F be the resulting $(0, 1, *)$ -matrix. That is, given a depth-2 circuit F computing a linear operator $\mathbf{x} \rightarrow M\mathbf{x}$, we replace by $*$'s all entries of M seen by the circuit, and denote the resulting $(0, 1, *)$ -matrix by A_F . Note that the original matrix M is one of the completions of A_F ; hence, $\text{rk}(M) \geq \text{mr}(A_F)$.

Lemma 2.1. *Every linear depth-2 circuit F has $\text{width}(F) \geq \text{mr}(A_F)$.*

In particular, if F computes a linear operator $\mathbf{x} \mapsto M\mathbf{x}$ and has no direct input-output wires at all, then $A_F = M$ and $\text{width}(F) \geq \text{rk}(M)$.

Proof. Let $M\mathbf{x}$ be a linear operator computed by F . Every assignment of constants to direct input-output wires leads to a depth-2 circuit of degree $d = 0$ computing a linear operator $B\mathbf{x}$, where B is a completion of A_F . This operator takes $2^{\text{rk}(B)}$ different values. Hence, the operator $H : \{0, 1\}^n \rightarrow \{0, 1\}^w$ computed by $w = \text{width}(F)$ boolean functions on the middle layer of F must take at least so many different values, as well. This implies that the width w must be large enough to fulfill $2^w \geq 2^{\text{rk}(B)}$, from which $w \geq \text{rk}(B) \geq \text{mr}(A_F)$ follows. \square

Lemma 2.2. *Every depth-2 circuit F computing a linear operator can be transformed into an equivalent linear depth-2 circuit of the same degree and width at most $\text{mr}(A_F)$.*

Together with Lemma 2.1, this implies that $\text{width}(F) = \text{mr}(A_F)$ for every optimal linear depth-2 circuit F .

Proof. Let $\mathbf{x} \rightarrow M\mathbf{x}$ be the operator computed by F , and let $A = A_F$ be the $(0, 1, *)$ -matrix of F . We can construct the desired *linear* depth-2 circuit computing $M\mathbf{x}$ as follows. Take a completion B of A with $\text{rk}(B) = \text{mr}(A)$. By the definition of completions, the i th row \mathbf{b}_i of B has the form $\mathbf{b}_i = \mathbf{a}_i + \mathbf{p}_i$, where \mathbf{a}_i is the i th row of A with all stars set to 0, and \mathbf{p}_i is a $(0, 1)$ -vector having no 1's in positions, where this row of A has non-stars. The i th row \mathbf{m}_i of the original $(0, 1)$ -matrix M is of the form $\mathbf{m}_i = \mathbf{a}_i + \mathbf{m}'_i$, where \mathbf{m}'_i is a $(0, 1)$ -vector which coincides with \mathbf{m}_i in all positions, where the i th row of A has stars, and has 0's elsewhere.

The matrix B has $r = \text{rk}(B) = \text{mr}(A)$ linearly independent rows. Assume w.l.o.g. that these are the first rows $\mathbf{b}_1, \dots, \mathbf{b}_r$ of B , and add r linear gates computing the scalar products $\langle \mathbf{b}_1, \mathbf{x} \rangle, \dots, \langle \mathbf{b}_r, \mathbf{x} \rangle$ over GF_2 on the middle layer. Connect by wires each of these linear gates with all input and all output nodes. Note that the i th output gate, knowing the vectors \mathbf{p}_i and \mathbf{m}'_i , can compute both scalar products $\langle \mathbf{p}_i, \mathbf{x} \rangle$ and $\langle \mathbf{m}'_i, \mathbf{x} \rangle$ by only using existing direct wires from inputs x_1, \dots, x_n to this gate. Hence, using the r linear gates $\langle \mathbf{b}_1, \mathbf{x} \rangle, \dots, \langle \mathbf{b}_r, \mathbf{x} \rangle$ on the middle layer, the i th output gate, for $i \leq r$, can also compute the whole scalar product $\langle \mathbf{m}_i, \mathbf{x} \rangle$ of the input vector with the i th row of M by:

$$\langle \mathbf{m}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus \langle \mathbf{m}'_i, \mathbf{x} \rangle = \langle \mathbf{b}_i, \mathbf{x} \rangle \oplus \langle \mathbf{p}_i, \mathbf{x} \rangle \oplus \langle \mathbf{m}'_i, \mathbf{x} \rangle.$$

For $i > r$, just replace vector \mathbf{b}_i in this expression by the corresponding linear combination of $\mathbf{b}_1, \dots, \mathbf{b}_r$. We have thus constructed an equivalent linear depth-2 circuit of the same degree and of width $r = \text{mr}(A_F)$. \square

By Lemma 2.2, the main question is: How much the width of a circuit F can be smaller than the min-rank of its matrix A_F ? Ideally, we would like to have that $\text{width}(F) \geq \epsilon \cdot \text{mr}(A_F)$: then the width of the resulting *linear* circuit would be at most $1/\epsilon$ times larger than that of the original circuit F .

Lemma 2.1 lower bounds the width of *linear* circuits F in terms of the min-rank of their $(0, 1, *)$ -matrices A_F . We now show that the Min-Rank Conjecture implies a similar fact also for general (non-linear) circuits.

Lemma 2.3. *For every depth-2 circuit F computing a linear operator in n variables, we have that*

$$\text{width}(F) \geq n - \log_2 \text{opt}(A_F).$$

Hence, the Min-Rank Conjecture (stating that $\text{opt}(A) \leq 2^{n-\epsilon \cdot \text{mr}(A)}$) implies that $\text{width}(F) \geq \epsilon \cdot \text{mr}(A_F)$.

Proof. Let M be an m -by- n $(0, 1)$ -matrix. Take a depth-2 circuit F of width w computing $M\mathbf{x}$, and let A_F be the corresponding $(0, 1, *)$ -matrix. Let $H = (h_1, \dots, h_w)$ be an operator computed at the gates on the middle layer, and $G = (g_1, \dots, g_m)$ an operator computed at the gates on the output layer. Hence, $M\mathbf{x} = G(\mathbf{x}, H(\mathbf{x}))$ for all $\mathbf{x} \in \{0, 1\}^n$. Fix a vector $\mathbf{b} \in \{0, 1\}^w$ for which the set $L = \{\mathbf{x} \in \{0, 1\}^n : M\mathbf{x} = G(\mathbf{x}, \mathbf{b})\}$ is the largest one; hence, $|L| \geq 2^{n-w}$. Note that the operator $G'(\mathbf{x}) := G(\mathbf{x}, \mathbf{b})$ must be consistent with A : its i th component $g'_i(\mathbf{x})$ can only depend on input variables x_j to which the i th output gate g_i is connected. Hence, L is a solution for A_F , implying that $\text{opt}(A_F) \geq |L| \geq 2^{n-w}$ from which the desired lower bound $w \geq n - \log_2 \text{opt}(A_F)$ on the width of F follows. \square

We can now show that the Min-Rank Conjecture (Conjecture 2) indeed implies the Linearization Conjecture (Conjecture 1).

Corollary 2.4. *Conjecture 2 implies Conjecture 1.*

Proof. Let F be a depth-2 circuit computing a linear operator in n variables. Assuming Conjecture 2, Lemma 2.3 implies that $\epsilon \cdot \text{mr}(A_F) \leq n - \log_2 \text{opt}(A_F) \leq \text{width}(F)$. By Lemma 2.2, the circuit F can be transformed into an equivalent linear depth-2 circuit of the same degree and width at most $\text{mr}(A_F) \leq \text{width}(F)/\epsilon$. \square

Hence, together with Valiant's result, the Min-Rank Conjecture implies that a linear operator $M\mathbf{x}$ requires a super-linear number of gates in any log-depth circuit over $\{\&, \vee, \neg\}$, if every depth-2 circuit for $M\mathbf{x}$ over $\{\oplus, 1\}$ of width $w = O(n/\ln \ln n)$ requires degree $n^{\Omega(1)}$.

Finally, let us show that the only ‘‘sorrow’’, when trying to linearize a depth-2 circuit, is the possible non-linearity of *output* gates—non-linearity of gates on the middle layer is no problem.

Lemma 2.5. *Let F be a depth-2 circuit computing a linear operator. If all gates on the output layer are linear boolean functions, then F can be transformed into an equivalent linear depth-2 circuit of the same degree and width.*

Proof. Let M be an m -by- n $(0, 1)$ -matrix, and let F be a depth-2 circuit of width w computing $M\mathbf{x}$. Let $H = (h_1, \dots, h_w)$ be the operator $H : \{0, 1\}^n \rightarrow \{0, 1\}^w$ computed by the gates on the middle layer. Assume that all output gates of F are linear boolean functions. Let B be the m -by- n adjacency $(0, 1)$ -matrix of the bipartite graph formed by the direct input-output wires, and C be the m -by- w adjacency $(0, 1)$ -matrix of the bipartite graph formed by the wires joining the gates on the middle layer with those on the output layer. Then

$$M\mathbf{x} = B\mathbf{x} \oplus C \cdot H(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \{0, 1\}^n,$$

where $C \cdot H(\mathbf{x})$ is the product of the matrix C with the vector $\mathbf{y} = H(\mathbf{x})$. Hence,

$$C \cdot H(\mathbf{x}) = D\mathbf{x} \tag{4}$$

is a linear operator with $D = M \oplus B$. Write each vector $\mathbf{x} = (x_1, \dots, x_n)$ as the linear combination

$$\mathbf{x} = \sum_{i=1}^n x_i \mathbf{e}_i \quad (5)$$

of unit vectors $\mathbf{e}_1, \dots, \mathbf{e}_n \in \{0, 1\}^n$, and replace the operator H computed on the middle layer by a *linear* operator

$$H'(\mathbf{x}) := \sum_{i=1}^n x_i H(\mathbf{e}_i) \pmod{2}. \quad (6)$$

Then, using the linearity of the matrix-vector product, we obtain that (with all sums mod 2):

$$\begin{aligned} C \cdot H(\mathbf{x}) &= D \cdot \left(\sum x_i \mathbf{e}_i \right) && \text{by (4) and (5)} \\ &= \sum x_i D \mathbf{e}_i && \text{linearity} \\ &= \sum x_i C \cdot H(\mathbf{e}_i) && \text{by (4)} \\ &= C \cdot \left(\sum x_i H(\mathbf{e}_i) \right) && \text{linearity} \\ &= C \cdot H'(\mathbf{x}) && \text{by (6)}. \end{aligned}$$

Hence, we again have that $M\mathbf{x} = B\mathbf{x} \oplus C \cdot H'(\mathbf{x})$, meaning that the obtained *linear* circuit computes the same linear operator $M\mathbf{x}$. \square

3. Bounds on $\text{opt}(A)$

Recall that $\text{opt}(A)$ is the largest possible number of vectors in a solution for a given $(0, 1, *)$ -matrix A . The simplest properties of this parameter are summarized in the following

Lemma 3.1. *Let A be an m -by- n $(0, 1, *)$ -matrix. If A' is obtained by removing some rows of A , then $\text{opt}(A') \geq \text{opt}(A)$. If $A = [B, C]$ where B is an m -by- p submatrix of A for some $1 \leq p \leq n$, then*

$$\text{opt}(B) \cdot \text{opt}(C) \leq \text{opt}(A) \leq \text{opt}(B) \cdot 2^{n-p}.$$

Proof. The first claim $\text{opt}(A') \geq \text{opt}(A)$ is obvious, since addition of new equations can only decrease the number of solutions in any system of equations.

To prove $\text{opt}(A) \leq \text{opt}(B) \cdot 2^{n-p}$, take an optimal solution $L_A = \{\mathbf{x} : M\mathbf{x} = G(\mathbf{x})\}$ for A ; hence, $|L_A| = \text{opt}(A)$. Fix a vector $\mathbf{b} \in \{0, 1\}^{n-p}$ for which the set

$$L_B = \{\mathbf{y} \in \{0, 1\}^p : (\mathbf{y}, \mathbf{b}) \in L_A\}$$

is the largest one; hence, $|L_B| \geq \text{opt}(A)/2^{n-p}$. The completion M of A has the form $M = [M', M'']$, where M' is a completion of B and M'' is a completion of C . If we define an operator $G' : \{0, 1\}^p \rightarrow \{0, 1\}^m$ by

$$G'(\mathbf{y}) := G(\mathbf{y}, \mathbf{b}) \oplus M''\mathbf{b},$$

then $M'\mathbf{y} = G'(\mathbf{y})$ for all $\mathbf{y} \in L_B$. Hence, L_B is a solution for B , implying that $\text{opt}(A) \leq |L_B| \cdot 2^{n-p} \leq \text{opt}(B) \cdot 2^{n-p}$.

To prove $\text{opt}(A) \geq \text{opt}(B) \cdot \text{opt}(C)$, let $L_B = \{\mathbf{y} \in \{0, 1\}^p : M'\mathbf{y} = G'(\mathbf{y})\}$ be an optimal solution for B , and let $L_C = \{\mathbf{z} \in \{0, 1\}^{n-p} : M''\mathbf{z} = G''(\mathbf{z})\}$ be an optimal solution for C . For any pair $\mathbf{x} = (\mathbf{y}, \mathbf{z}) \in L_B \times L_C$, we have that $M\mathbf{x} = G(\mathbf{x})$, where $M = [M', M'']$ and $G(\mathbf{y}, \mathbf{z}) := G'(\mathbf{y}) \oplus G''(\mathbf{z})$. Hence, the set $L_B \times L_C \subseteq \{0, 1\}^n$ is a solution for A , implying that $\text{opt}(B) \cdot \text{opt}(C) = |L_B \times L_C| \leq \text{opt}(A)$, as claimed. \square

Let A be an m -by- n $(0, 1, *)$ -matrix. The min-rank conjecture claims that the largest number $\text{opt}(A)$ of vectors in a solution for A can be upper bounded in terms of the min-rank of A as $\text{opt}(A) \leq 2^{n-\epsilon \cdot \text{mr}(A)}$. The claim is true if the min-rank of A is “witnessed” by some $(0, 1)$ -submatrix of A , that is, if A contains a $(0, 1)$ -submatrix of rank equal to the min-rank of A . This is a direct consequence of the following simple

Lemma 3.2. *If A is an m -by- n $(0, 1, *)$ -matrix, then $\text{opt}(A) \leq 2^{n-\text{rk}(B)}$ for every $(0, 1)$ -submatrix B of A .*

Proof. Let B be a p -by- q $(0, 1)$ -submatrix of A . Since B has no stars, only constant operators can be consistent with B . Hence, if $L \subseteq \{0, 1\}^q$ is a solution for B , then there must be a vector $\mathbf{b} \in \{0, 1\}^p$ such that $B\mathbf{x} = \mathbf{b}$ for all $\mathbf{x} \in L$. This implies $|L| \leq 2^{q-\text{rk}(B)}$. Together with Lemma 3.1, this yields $\text{opt}(A) \leq 2^{q-\text{rk}(B)} \cdot 2^{n-q} = 2^{n-\text{rk}(B)}$. \square

The *max-rank* $\text{Mr}(A)$ of a $(0, 1, *)$ -matrix A is a maximal possible rank of its completion. A *line* of A is either its row or its column. A *cover* of A is a set X of its lines covering all stars. Let $\text{cov}(A)$ denote the smallest possible number of lines in a cover of A .

Lemma 3.3. *For every m -by- n $(0, 1, *)$ -matrix A , we have that*

$$\text{opt}(A) \leq 2^{n-\text{Mr}(A)+\text{cov}(A)}.$$

Proof. Given a cover X of the stars in A by lines, remove all these lines, and let A_X be the resulting $(0, 1)$ -submatrix of A . Clearly, we have: $\text{Mr}(A) \leq \text{rk}(A_X) + |X|$. (In fact, it is shown in [5] that $\text{Mr}(A) = \min_X (\text{rk}(A_X) + |X|)$, where the minimum is over all covers X of A .) Take a cover X of A of size $|X| = \text{cov}(A)$. Hence, $\text{Mr}(A) \leq \text{rk}(A_X) + \text{cov}(A)$. Since A_X is a $(0, 1)$ -submatrix of A , Lemma 3.2 yields $\text{opt}(A) \leq 2^{n-\text{rk}(A_X)}$, where $\text{rk}(A_X) \geq \text{Mr}(A) - |X| = \text{Mr}(A) - \text{cov}(A)$. \square

Given a depth-2 circuit F , let $\text{m}(F)$ denote the largest number of wires in a matching formed by direct input-output wires. That is, $\text{m}(F)$ is the largest number of $*$ -entries in the matrix A_F of F , no two on the same line. By the well-known König–Egeváry theorem, stating that the size of a largest matching in a bipartite graph is equal to the smallest set of vertices which together touch every edge, we have that $\text{m}(F) = \text{cov}(A_F)$. This leads to the following

Corollary 3.4. *Every depth-2 circuit F computing a linear operator can be transformed into an equivalent linear depth-2 circuit F' of the same degree and*

$$\text{width}(F') \leq \text{width}(F) + \text{m}(F).$$

Proof. Let A_F be the $(0, 1, *)$ -matrix of F . By Lemmas 2.3 and 3.3, we have that

$$\begin{aligned} \text{width}(F) &\geq n - \log_2 \text{opt}(A_F) \geq n - [n - \text{Mr}(A_F) + \text{cov}(A_F)] \\ &= \text{Mr}(A_F) - \text{cov}(A_F) = \text{Mr}(A_F) - \text{m}(F). \end{aligned}$$

By Lemma 2.2, the circuit F can be transformed into an equivalent linear depth-2 circuit of the same degree and width at most $\text{mr}(A_F) \leq \text{Mr}(A_F) \leq \text{width}(F) + \text{m}(F)$. \square

4. Row and column min-rank

We are now going to show that the min-rank conjecture holds for stronger versions of min-rank—row min-rank and column min-rank.

If A is a $(0, 1, *)$ -matrix of min-rank r then, for every assignment of constants to stars, the resulting $(0, 1)$ -matrix will have r linearly independent columns as well as r linearly independent rows. However, for different assignments these columns/rows may be different. It is natural to ask whether the min-rank conjecture is true if the matrix A has r columns (or r rows) that remain linearly independent under any assignment of constants to stars?

Namely, say that $(0, 1, *)$ -vectors are *dependent* if they can be made linearly dependent over GF_2 by setting their $*$ -entries to a constants 0 and 1; otherwise, the vectors are *independent*.

Remark4.1. The dependence of $(0, 1, *)$ -vectors can be defined by adding to $\{0, 1\}$ a new element $*$ satisfying $\alpha \oplus * = * \oplus \alpha = *$ for $\alpha \in \{0, 1, *\}$. Then a set of $(0, 1, *)$ -vectors is dependent iff some its subset sums up to a $(0, *)$ -vector. Indeed, if *some* subset sums up to a $(0, *)$ -vector, then we can set the $*$ -entries to constants so that the corresponding subset of $(0, 1)$ -vectors will sum up (over GF_2) to an all-0 vector. On the other hand, if *no* subset sums up to a $(0, *)$ -vector, for every subset, there must be a position in which all vectors in this subset have no stars, and the sum of these positions over GF_2 is 1.

Remark4.2. A basic fact of Linear Algebra, leading to the Gauss-Algorithm, is that linear independence of vectors $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$ implies that the vectors $\mathbf{x} + \mathbf{y}$ and \mathbf{y} are linear independent as well. For $(0, 1, *)$ -vectors this does not hold anymore. Take, for example, $\mathbf{x} = (0, 1)$ and $\mathbf{y} = (1, *)$. Then $\mathbf{x} \oplus \mathbf{y} = (1, *) = \mathbf{y}$.

For a $(0, 1, *)$ -matrix A , define its *column min-rank*, $\text{mr}_{\text{col}}(A)$, as the maximum number of independent columns, and its *row min-rank*, $\text{mr}_{\text{row}}(A)$, as the maximum number of independent rows. In particular, both $\text{mr}_{\text{row}}(A)$ and $\text{mr}_{\text{col}}(A)$ are at least r if A contains an $r \times r$ “triangular” submatrix, that is, a submatrix with zeroes below (or above) the diagonal and ones on the diagonal:

$$\Delta = \begin{pmatrix} 1 & \circledast & \circledast & \circledast \\ 0 & 1 & \circledast & \circledast \\ 0 & 0 & 1 & \circledast \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

where $\circledast \in \{0, 1, *\}$. It is clear that neither $\text{mr}_{\text{col}}(A)$ nor $\text{mr}_{\text{row}}(A)$ can exceed the min-rank of A . Later (Lemma 8.4 below) we will give an example of a matrix A where both $\text{mr}_{\text{col}}(A)$ and $\text{mr}_{\text{row}}(A)$ are by a logarithmic factor smaller than $\text{mr}(A)$. The question about a more precise relation between these parameters remains open (see Problem 9.3).

Albeit for $(0, 1)$ -matrices we always have that their row-rank coincides with column-rank, for $(0, 1, *)$ -matrices this is no more true. In particular, for some $(0, 1, *)$ -matrices A , we have that $\text{mr}_{\text{row}}(A) \neq \text{mr}_{\text{col}}(A)$.

Example4.3. Consider the following $(0, 1, *)$ -matrix:

$$A = \begin{pmatrix} 1 & 1 & * & 1 \\ 1 & 0 & 1 & * \\ 1 & * & 0 & 0 \end{pmatrix}.$$

Then $\text{mr}_{\text{row}}(A) = \text{mr}(A) = 3$ but $\text{mr}_{\text{col}}(A) = 2$. To see that $\text{mr}_{\text{row}}(A) = 3$, just observe that the rows cannot be made linearly dependent by setting the stars to 0 or 1: the sum of all three vectors is not a $\{0, *\}$ -vector because of the 1st column, and the pairwise sums are not $\{0, *\}$ -vectors because, for each pair of rows there is a column containing 0 and 1. To see that $\text{mr}_{\text{col}}(A) = 2$, observe that the last three columns are dependent (each row has a star). Moreover, for every pair of these columns, there is an assignment of constants to stars such that either the resulting $(0, 1)$ -columns are equal or their sum equals the first column.

We first show that the min-rank conjecture holds with “min-rank” replaced by “column min-rank”.

Theorem 4.4 (Column min-rank). *Let A be a $(0, 1, *)$ -matrix with n columns and of column min-rank r . Then $\text{opt}(A) \leq 2^{n-r}$.*

Proof. Any m -by- n $(0, 1, *)$ -matrix B of column min-rank r must contain an $m \times r$ submatrix A of min-rank r . Since $\text{opt}(B) \leq \text{opt}(A) \cdot 2^{n-r}$ (Lemma 3.1), it is enough to show that $\text{opt}(A) \leq 1$ for all m -by- r $(0, 1, *)$ -matrices A of min-rank r .

To do this, let L be a solution for A . Then there is an operator $G = (g_1, \dots, g_m) : \{0, 1\}^r \rightarrow \{0, 1\}^m$ such that G is consistent with A and $\langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(\mathbf{x})$ holds for all $\mathbf{x} \in L$ and all $i = 1, \dots, m$. Here $\mathbf{a}_1, \dots, \mathbf{a}_m$ are the rows of A with all stars set to 0.

For the sake of contradiction, assume that $|L| \geq 2$ and fix any two vectors $\mathbf{x} \neq \mathbf{y} \in L$. Our goal is to construct a vector $\mathbf{c} \in \{0, 1\}^m$ and a completion M of A such that $M\mathbf{x} = M\mathbf{y} = \mathbf{c}$. Since M must have rank r , this will give the desired contradiction, because at most $2^{r-\text{rk}(M)} = 2^0 = 1$ vectors \mathbf{z} can satisfy $M\mathbf{z} = \mathbf{c}$.

If M is a completion of $A = (a_{ij})$, then its i th row must have the form $\mathbf{m}_i = \mathbf{a}_i \oplus \mathbf{p}_i$ where $\mathbf{p}_i \in \{0, 1\}^n$ is some vector with no 1's in positions where the i th row of A has no stars. To construct the desired vector \mathbf{p}_i for each $i \in [m]$, we consider two possible cases. (Recall that the vectors \mathbf{x} and \mathbf{y} are fixed.)

Case 1: $\langle \mathbf{a}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{y} \rangle$. In this case we can take $\mathbf{p}_i = \mathbf{0}$ and $c_i = \langle \mathbf{a}_i, \mathbf{x} \rangle$. Then $\langle \mathbf{m}_i, \mathbf{x} \rangle = \langle \mathbf{m}_i, \mathbf{y} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle = c_i$, as desired.

Case 2: $\langle \mathbf{a}_i, \mathbf{x} \rangle \neq \langle \mathbf{a}_i, \mathbf{y} \rangle$. In this case we have that $g_i(\mathbf{x}) \neq g_i(\mathbf{y})$, that is, the vectors \mathbf{x} and \mathbf{y} must differ in some position j where the i th row of A has a star. Then we can take $\mathbf{p}_i := \mathbf{e}_j$ (the j th unit vector) and $c_i := \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus x_j$. With this choice of \mathbf{p}_i , we again have

$$\langle \mathbf{m}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus \langle \mathbf{p}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus \langle \mathbf{e}_j, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus x_j = c_i$$

and, since $\langle \mathbf{a}_i, \mathbf{x} \rangle \neq \langle \mathbf{a}_i, \mathbf{y} \rangle$ and $x_j \neq y_j$,

$$\langle \mathbf{m}_i, \mathbf{y} \rangle = \langle \mathbf{a}_i, \mathbf{y} \rangle \oplus \langle \mathbf{p}_i, \mathbf{y} \rangle = \langle \mathbf{a}_i, \mathbf{y} \rangle \oplus \langle \mathbf{e}_j, \mathbf{y} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle \oplus x_j = c_i. \quad \square$$

Example 4.5. It is not difficult to verify that, for the $(0, 1, *)$ -matrix A given by (3), we have that $\text{mr}_{\text{col}}(A) = \text{mr}(A) = 2$. Hence, no linear solution of the system of semi-linear equations (2) can have more than $\text{lin}(A) = 2^{6-2} = 32$ vectors. Theorem 4.4 implies that, in fact, *no* solution can have more than this number of vectors.

The situation with *row* min-rank is more complicated. In this case we are only able to prove an upper bound $\text{opt}(A) \leq 2^{n-r}$ under an additional restriction that the star-positions in the rows of A form a chain under set-inclusion.

Recall that $(0, 1, *)$ -vectors are *independent* if they cannot be made linearly dependent over GF_2 by setting stars to constants. The row min-rank of a $(0, 1, *)$ -matrix is the largest number r of its

independent rows. Since adding new rows can only decrease $\text{opt}(A)$, it is enough to consider r -by- n $(0, 1, *)$ -matrices A with $\text{mr}(A) = r$.

If $r = 1$, that is, if A consists of just one row, then $\text{opt}(A) \leq 2^{n-1} = 2^{n-r}$ holds. Indeed, since $\text{mr}(A) = 1$, this row cannot be a $(0, *)$ -row. So, there must be at least one 1 in, say, the 1st position. Let $L_A = \{\mathbf{x} : \langle \mathbf{a}_1, \mathbf{x} \rangle = g_1(\mathbf{x})\}$ be a solution for A , where \mathbf{a}_1 is the row of A with all stars set to 0. Take the unit vector $\mathbf{e}_1 = (1, 0, \dots, 0)$ and split the vectors in $\{0, 1\}^n$ into 2^{n-1} pairs $\{\mathbf{x}, \mathbf{x} \oplus \mathbf{e}_1\}$. Since the boolean function g_1 cannot depend on the first variable x_1 , we have that $g_1(\mathbf{x} \oplus \mathbf{e}_1) = g_1(\mathbf{x})$. But $\langle \mathbf{a}_1, \mathbf{x} \oplus \mathbf{e}_1 \rangle = \langle \mathbf{a}_1, \mathbf{x} \rangle \oplus 1 \neq \langle \mathbf{a}_1, \mathbf{x} \rangle$. Hence, at most one of the two vectors \mathbf{x} and $\mathbf{x} \oplus \mathbf{e}_1$ from each pair $\{\mathbf{x}, \mathbf{x} \oplus \mathbf{e}_1\}$ can lie in L_A , implying that $|L_A| \leq 2^{n-1}$.

To extend this argument for matrices with more rows, we need the following definition. Let $A = (a_{ij})$ be an r -by- n $(0, 1, *)$ -matrix, and $\mathbf{a}_1, \dots, \mathbf{a}_r$ be the rows of A with all stars set to 0. Let $S_i = \{j : a_{ij} = *\}$ be the set of star-positions in the i th row of A . It will be convenient to describe the star-positions by diagonal matrices. Namely, let D_i be the incidence matrix of stars in the i th row of A . That is, D_i is a diagonal n -by- n $(0, 1)$ -matrix whose j th diagonal entry is 1 iff $j \in S_i$. In particular, $D_i \mathbf{x} = \mathbf{0}$ means that $x_j = 0$ for all $j \in S_i$.

Definition 4.6. A matrix A is *isolated* if there exist vectors $\mathbf{z}_1, \dots, \mathbf{z}_r \in \{0, 1\}^n$ such that, for all $1 \leq i \leq r$, we have $D_i \mathbf{z}_i = \mathbf{0}$ and

$$\langle \mathbf{a}_j, \mathbf{z}_i \rangle = \begin{cases} 1 & \text{if } j = i; \\ 0 & \text{if } j < i. \end{cases}$$

If $D_1 \mathbf{z}_1 = \dots = D_r \mathbf{z}_r = \mathbf{0}$, then the matrix is *strongly isolated*.

Lemma 4.7. *If A is a strongly isolated r -by- n $(0, 1, *)$ -matrix, then $\text{opt}(A) \leq 2^{n-r}$.*

Proof. Let $\mathbf{a}_1, \dots, \mathbf{a}_r$ be the rows of A with all stars set to 0. We prove the lemma by induction on r . The basis case $r = 1$ is already proved above. For the induction step $r - 1 \mapsto r$, let

$$L_A = \{\mathbf{x} \in \{0, 1\}^n : \langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(\mathbf{x}) \text{ for all } i = 1, \dots, r\}$$

be an optimal solution for A , and let B be a submatrix of A consisting of its first $r - 1$ rows. Then

$$L_B = \{\mathbf{x} \in \{0, 1\}^n : \langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(\mathbf{x}) \text{ for all } i = 1, \dots, r - 1\}$$

is a solution for B . Since A is strongly isolated, the matrix B is strongly isolated as well. The induction hypothesis implies that $|L_B| \leq 2^{n-(r-1)}$.

Let $\mathbf{z} = \mathbf{z}_r$ be the r -th isolating vector. For each row $i = 1, \dots, r - 1$, the conditions $\langle \mathbf{z}, \mathbf{a}_i \rangle = 0$ and $D_i \mathbf{z} = \mathbf{0}$ imply that $\langle (\mathbf{x} \oplus \mathbf{z}), \mathbf{a}_i \rangle = \langle \mathbf{x}, \mathbf{a}_i \rangle$ and $g_i(\mathbf{x} \oplus \mathbf{z}) = g_i(\mathbf{x})$. That is,

$$\mathbf{x} \in L_B \text{ iff } \mathbf{x} \oplus \mathbf{z} \in L_B.$$

For the r th row, the conditions $\langle \mathbf{z}, \mathbf{a}_r \rangle = 1$ and $D_r \mathbf{z} = \mathbf{0}$ imply that $\langle (\mathbf{x} \oplus \mathbf{z}), \mathbf{a}_r \rangle \neq \langle \mathbf{x}, \mathbf{a}_r \rangle$ whereas $g_r(\mathbf{x} \oplus \mathbf{z}) = g_r(\mathbf{x})$. That is,

$$\mathbf{x} \in L_A \text{ iff } \mathbf{x} \oplus \mathbf{z} \notin L_A.$$

Hence, for every vector $\mathbf{x} \in L_B$, only one of the vectors \mathbf{x} and $\mathbf{x} \oplus \mathbf{z}$ can belong to L_A , implying that

$$\text{opt}(A) = |L_A| \leq |L_B|/2 \leq 2^{n-r}. \quad \square$$

We are now going to show that $(0, 1, *)$ -matrices with some conditions on the distribution of stars in them are strongly isolated. For this, we need the following two facts. A *projection* of a vector $\mathbf{x} = (x_1, \dots, x_n)$ onto a set of positions $I = \{i_1, \dots, i_k\}$ is the vector

$$\mathbf{x} \upharpoonright_I = (x_{i_1}, \dots, x_{i_k}).$$

A $(0, 1, *)$ -vector \mathbf{x} is *independent* of $(0, 1, *)$ -vectors $\mathbf{y}_1, \dots, \mathbf{y}_k$ if no completion of \mathbf{x} can be written as a linear combination of some completions of these vectors.

Lemma 4.8. *Let $\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_k$ be $(0, 1, *)$ -vectors, and $I = \{i : x_i \neq *\}$. If \mathbf{x} is independent of $\mathbf{y}_1, \dots, \mathbf{y}_k$, then $\mathbf{x} \upharpoonright_I$ is also independent of $\mathbf{y}_1 \upharpoonright_I, \dots, \mathbf{y}_k \upharpoonright_I$.*

Proof. Assume that $\mathbf{x} \upharpoonright_I$ is dependent on the projections $\mathbf{y}_1 \upharpoonright_I, \dots, \mathbf{y}_k \upharpoonright_I$. Then there is an assignment of stars to constants in the vectors \mathbf{y}_i such that $\mathbf{x} \upharpoonright_I$ can be written as a linear combination of the projections $\mathbf{y}'_1 \upharpoonright_I, \dots, \mathbf{y}'_k \upharpoonright_I$ on I of the resulting $(0, 1)$ -vectors $\mathbf{y}'_1, \dots, \mathbf{y}'_k$. But since \mathbf{x} has stars in all positions outside I , these stars can be set to appropriate constants so that the resulting $(0, 1)$ -vector \mathbf{x}' will be a linear combination of $\mathbf{y}'_1, \dots, \mathbf{y}'_k$, a contradiction. \square

Lemma 4.9. *Let $\mathbf{a} \in \{0, 1\}^n$ be a vector and M be an m -by- n $(0, 1)$ -matrix of rank $r \leq n - 1$. If \mathbf{a} is linearly independent of the rows of M , then there exists a set $Z \subseteq \{0, 1\}^n$ of $|Z| \geq 2^{n-r-1}$ vectors such that, for all $\mathbf{z} \in Z$, we have $\langle \mathbf{z}, \mathbf{a} \rangle = 1$ and $M\mathbf{z} = \mathbf{0}$.*

Proof. Let $Z = \{\mathbf{z} : M\mathbf{z} = \mathbf{0}, \langle \mathbf{a}, \mathbf{z} \rangle = 1\}$, and let M' be the matrix M with an additional row \mathbf{a} . Note that $Z = \ker(M) \setminus \ker(M')$, where $\ker(M) = \{\mathbf{z} : M\mathbf{z} = \mathbf{0}\}$ is the kernel of M . Since $\text{rk}(M') = \text{rk}(M) + 1 \leq n$, we have that $|\ker(M')| = |\ker(M)|/2$, implying that

$$|Z| = |\ker(M) \setminus \ker(M')| = |\ker(M)|/2 \geq 2^{n-r-1}. \quad \square$$

Lemma 4.10. *If A is an r -by- n $(0, 1, *)$ -matrix with $\text{mr}(A) = r$, then A is isolated.*

Proof. Let $\mathbf{a}_1, \dots, \mathbf{a}_r$ be the rows of A with all stars set to 0. Let $I \subseteq \{1, \dots, n\}$ be the set of all star-free positions in the i th row of A , and consider an $(r - 1)$ -by- $|I|$ $(0, 1)$ -matrix M_i whose rows are the projections $\mathbf{a}'_j = \mathbf{a}_j \upharpoonright_I$ of vectors \mathbf{a}_j with $j \neq i$ onto the set I . By Lemma 4.8, the projection $\mathbf{a}'_i = \mathbf{a}_i \upharpoonright_I$ of the i th vector \mathbf{a}_i onto I cannot be written as a linear combination of the rows of M_i ; hence, $\text{rk}(M_i) \leq |I| - 1$. Since $2^{|I| - \text{rk}(M_i) - 1} \geq 2^0 = 1$, Lemma 4.9 gives us a vector $\mathbf{z}'_i \in \{0, 1\}^{|I|}$ such that $\langle \mathbf{z}'_i, \mathbf{a}'_i \rangle = 1$ and $\langle \mathbf{z}'_i, \mathbf{a}'_j \rangle = 0$ for all $j \neq i$. But then $\mathbf{z}_i := (\mathbf{z}'_i, \mathbf{0})$ is the desired $(0, 1)$ -vector: $D_i \mathbf{z}_i = D_i \cdot \mathbf{0} = \mathbf{0}$, $\langle \mathbf{z}_i, \mathbf{a}_i \rangle = \langle \mathbf{z}'_i, \mathbf{a}'_i \rangle = 1$, and $\langle \mathbf{z}_i, \mathbf{a}_j \rangle = \langle \mathbf{z}'_i, \mathbf{a}'_j \rangle = 0$ for all rows $j \neq i$. \square

Say that an r -by- n $(0, 1, *)$ -matrix A is *star-monotone* if the sets S_1, \dots, S_r of star-positions in its rows form a chain, that is, if $S_1 \subseteq S_2 \subseteq \dots \subseteq S_r$.

Theorem 4.11 (Star-monotone matrices). *Let A be a $(0, 1, *)$ -matrix with n columns. If A contains an r -by- n star-monotone submatrix of min-rank r , then $\text{opt}(A) \leq 2^{n-r}$.*

Proof. Since addition of new rows can only decrease the size of a solution, we can assume that A itself is an r -by- n star-monotone matrix of min-rank r . Let $\mathbf{a}_1, \dots, \mathbf{a}_r$ be the rows of A with all stars set to 0. By Lemma 4.10, the matrix A is isolated. That is, there exist vectors $\mathbf{z}_1, \dots, \mathbf{z}_r \in \{0, 1\}^n$ such that $\langle \mathbf{a}_i, \mathbf{z}_j \rangle = 1$ iff $i = j$, and $D_i \mathbf{z}_i = \mathbf{0}$ for all $1 \leq i \leq r$. Since $S_j \subseteq S_i$ for all $j < i$, this last condition implies that $D_j \mathbf{z}_i = \mathbf{0}$ for all $1 \leq j < i \leq r$, that is, A is strongly isolated. Hence, we can apply Lemma 4.7. \square

5. Solutions as independent sets in Cayley graphs

Let $A = (a_{ij})$ be an m -by- n $(0, 1, *)$ -matrix. In the definition of solutions L for A we take a completion M of A and an operator $G(\mathbf{x})$, and require that $M\mathbf{x} = G(\mathbf{x})$ for all $\mathbf{x} \in L$. The operator $G = (g_1, \dots, g_m)$ can be arbitrary—the only restriction is that its i th component g_i can only depend on variables corresponding to stars in the i th row of A . In this section we show that the actual *form* of operators G can be ignored—only star-positions are important. To do this, we associate with A the following set of “forbidden” vectors:

$$K_A = \{\mathbf{x} \in \{0, 1\}^n : \exists i \in [m] \ D_i\mathbf{x} = \mathbf{0} \text{ and } \langle \mathbf{a}_i, \mathbf{x} \rangle = 1\},$$

where D_i is the incidence n -by- n $(0, 1)$ -matrix of stars in the i th row of A , and \mathbf{a}_i is the i th row of A with all stars set to 0. Hence, K_A is a union $K_A = \bigcup_{i=1}^m K_i$ of m affine spaces

$$K_i = \left\{ \mathbf{x} : \begin{pmatrix} D_i \\ \mathbf{a}_i \end{pmatrix} \mathbf{x} = \begin{pmatrix} \mathbf{0} \\ 1 \end{pmatrix} \right\}.$$

Lemma 5.1. *For every vector $\mathbf{x} \in \{0, 1\}^n$, $\mathbf{x} \in K_A$ if and only if $M\mathbf{x} \neq \mathbf{0}$ for all completions M of A .*

Proof. (\Rightarrow): Take a vector $\mathbf{x} \in K_A$. Then there exists an $i \in [m]$ such that vector \mathbf{x} has zeroes in all positions, where the i th row of A has stars, and $\langle \mathbf{a}_i, \mathbf{x} \rangle = 1$, where \mathbf{a}_i is obtained by setting all stars in this row to 0. So, if \mathbf{b}_i is any completion of the i th row of A then $\langle \mathbf{b}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle = 1$. Thus, the scalar product of \mathbf{x} with the i th row of any completion of A must be equal to 1.

(\Leftarrow): Take a vector $\mathbf{x} \notin K_A$. We have to show that then $M\mathbf{x} = \mathbf{0}$ for at least one completion M of A . The fact that \mathbf{x} does not belong to K_A means that for each $i \in [m]$ either (i) $\langle \mathbf{a}_i, \mathbf{x} \rangle = 0$, or (ii) $\langle \mathbf{a}_i, \mathbf{x} \rangle = 1$ but vector \mathbf{x} has a 1 in some position j , where the i th row of A has a star. We can therefore construct the i th row \mathbf{m}_i of the desired completion M of A with $M\mathbf{x} = \mathbf{0}$ by taking $\mathbf{m}_i = \mathbf{a}_i$, if (i), and $\mathbf{m}_i = \mathbf{a}_i + \mathbf{e}_j$, if (ii). In both cases we have $\langle \mathbf{m}_i, \mathbf{x} \rangle = 0$, as desired. \square

The *sum-set* of two sets of vectors $S, T \subseteq \{0, 1\}^n$ is the set of vectors

$$S + T = \{\mathbf{x} \oplus \mathbf{y} : \mathbf{x} \in S \text{ and } \mathbf{y} \in T\}.$$

Theorem 5.2. *A set $L \subseteq \{0, 1\}^n$ is a solution for A if and only if $(L + L) \cap K_A = \emptyset$.*

Proof. Observe that the sum $\mathbf{x} \oplus \mathbf{y}$ of two vectors belongs to K_A iff these vectors coincide on all stars of at least one row of A such that $\langle \mathbf{a}_i, \mathbf{x} \rangle \neq \langle \mathbf{a}_i, \mathbf{y} \rangle$. By this observation, we see that the condition $(L + L) \cap K_A = \emptyset$ is equivalent to:

$$\forall \mathbf{x}, \mathbf{y} \in L \ \forall i \in [m] : D_i\mathbf{x} = D_i\mathbf{y} \text{ implies } \langle \mathbf{a}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{y} \rangle. \quad (7)$$

Having made this observation, we now turn to the actual proof of Theorem 5.2.

(\Rightarrow) Let L be a solution for A . Hence, there is an operator $G = (g_1, \dots, g_m)$ consistent with A such that $\langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(\mathbf{x})$ for all $\mathbf{x} \in L$ and all rows $i \in [m]$. To show that then L must satisfy (7), take any two vectors $\mathbf{x}, \mathbf{y} \in L$ and assume that $D_i\mathbf{x} = D_i\mathbf{y}$. This means that vectors \mathbf{x} and \mathbf{y} must coincide in all positions where the i th row of A has stars. Since g_i can only depend on these positions, this implies $g_i(\mathbf{x}) = g_i(\mathbf{y})$, and hence, $\langle \mathbf{a}_i, \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{y} \rangle$.

(\Leftarrow) Assume that $L \subseteq \{0, 1\}^n$ satisfies (7). We have to show that then there exists an operator $G = (g_1, \dots, g_m)$ consistent with A such that $\langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(\mathbf{x})$ for all $\mathbf{x} \in L$ and $i \in [m]$; here, as before, \mathbf{a}_i is the i th row of A with all stars set to 0. The i th row of A splits the set L into two subsets

$$L_i^0 = \{\mathbf{x} \in L : \langle \mathbf{a}_i, \mathbf{x} \rangle = 0\} \quad \text{and} \quad L_i^1 = \{\mathbf{x} \in L : \langle \mathbf{a}_i, \mathbf{x} \rangle = 1\}.$$

Condition (7) implies that $D_i \mathbf{x} \neq D_i \mathbf{y}$ for all $(\mathbf{x}, \mathbf{y}) \in L_i^0 \times L_i^1$. That is, if S_i is the set of star-positions in the i th row of A , then the projections $\mathbf{x}|_{S_i}$ of vectors \mathbf{x} in L_i^0 onto these positions must be different from all the projections $\mathbf{y}|_{S_i}$ of vectors \mathbf{y} in L_i^1 . Hence, we can find a boolean function $g_i : \{0, 1\}^{S_i} \rightarrow \{0, 1\}$ taking different values on these two sets of projections. This function will then satisfy $g_i(\mathbf{x}) = \langle \mathbf{a}_i, \mathbf{x} \rangle$ for all $\mathbf{x} \in L$. \square

A coset of a set of vectors $L \subseteq \{0, 1\}^n$ is a set $\mathbf{v} + L = \{\mathbf{v} \oplus \mathbf{x} : \mathbf{x} \in L\}$ with $\mathbf{v} \in \{0, 1\}^n$. Since $(\mathbf{v} + L) + (\mathbf{v} + L) = L + L$, Theorem 5.2 implies:

Corollary 5.3. *Every coset of a solution for a $(0, 1, *)$ -matrix A is also a solution for A .*

Remark 5.4. A Cayley graph over the Abelian group $(\{0, 1\}^n, \oplus)$ generated by a set $K \subseteq \{0, 1\}^n$ of vectors has all vectors in $\{0, 1\}^n$ as vertices, and two vectors \mathbf{x} and \mathbf{y} are joined by an edge iff $\mathbf{x} \oplus \mathbf{y} \in K$. Theorem 5.2 shows that solutions for a $(0, 1, *)$ -matrix A are precisely the independent sets in a Cayley graph generated by a special set K_A .

Remark 5.5. If A is an m -by- n $(0, 1)$ -matrix, that is, has no stars at all, then $K_A = \{\mathbf{x} : A\mathbf{x} \neq \mathbf{0}\}$. Hence, in this case, a set $L \subseteq \{0, 1\}^n$ is a solution for A iff there is a vector $\mathbf{b} \in \{0, 1\}^m$ such that $A\mathbf{x} = \mathbf{b}$ for all $\mathbf{x} \in L$. That is, in this case, $\ker(A) = \{\mathbf{x} : A\mathbf{x} = \mathbf{0}\}$ is an optimal solution.

6. Structure of linear solutions

By Theorem 5.2, a set of vectors $L \subseteq \{0, 1\}^n$ is a solution for an m -by- n $(0, 1, *)$ -matrix A if and only if $(L + L) \cap K_A = \emptyset$, where $K_A \subseteq \{0, 1\}^n$ is the set of “forbidden” vectors for A . Thus, *linear* solutions are precisely vector subspaces of $\{0, 1\}^n$ avoiding the set K_A . Which subspaces these are? We will show (Theorem 6.2) that these are precisely the subspaces lying entirely in the kernel of some completion of A .

Each vector subspace of $\{0, 1\}^n$ is a kernel $\ker(H) = \{\mathbf{x} : H\mathbf{x} = \mathbf{0}\}$ of some $(0, 1)$ -matrix H . Hence, linear solutions for A are given by matrices H such that $H\mathbf{x} \neq \mathbf{0}$ for all $\mathbf{x} \in K_A$; in this case we also say that the matrix H *separates* K_A from zero. By the *span-matrix* of a $(0, 1)$ -matrix H we will mean the matrix \widehat{H} whose rows are all linear combinations of the rows of H .

Lemma 6.1. *Let A be a $(0, 1, *)$ -matrix and H be $(0, 1)$ -matrix. Then $\ker(H)$ is a solution for A iff \widehat{H} contains a completion of A .*

Proof. To prove (\Leftarrow), suppose that some completion M of A is a submatrix of \widehat{H} . Let $\mathbf{x} \in K_A$. By Lemma 5.1, we know that then $M\mathbf{x} \neq \mathbf{0}$, and hence, also $\widehat{H}\mathbf{x} \neq \mathbf{0}$. Since $H\mathbf{x} = \mathbf{0}$ would imply $\widehat{H}\mathbf{x} = \mathbf{0}$, we also have that $H\mathbf{x} \neq \mathbf{0}$.

To prove (\Rightarrow), suppose that $\ker(H)$ is a solution for A , that is, $H\mathbf{x} \neq \mathbf{0}$ for all $\mathbf{x} \in K_A$. Then, for every row $i \in [m]$ and every vector $\mathbf{x} \in \{0, 1\}^n$, $H\mathbf{x} = \mathbf{0}$ and $D_i \mathbf{x} = \mathbf{0}$ imply that $\langle \mathbf{a}_i, \mathbf{x} \rangle = 0$. This means that \mathbf{a}_i must be a linear combination of rows of H and D_i . Hence, for each i , the vector \mathbf{a}_i must lie in the vector space spanned by the rows of H and D_i , that is, $\mathbf{a}_i = \boldsymbol{\alpha}_i^\top H \oplus \boldsymbol{\beta}_i^\top D_i$ for some vectors $\boldsymbol{\alpha}_i$ and $\boldsymbol{\beta}_i$. In other words, the i th linear combination $\boldsymbol{\alpha}_i^\top H$ of the rows of H is the i th row $\mathbf{a}_i \oplus \boldsymbol{\beta}_i^\top D_i$ of a particular completion M of A , implying that M is a submatrix of \widehat{H} , as desired. \square

Theorem 6.2. *Let A be a $(0, 1, *)$ -matrix. A linear subspace is a solution for A if and only if it is contained in a kernel of some completion of A .*

Proof. (\Leftarrow): If a linear subspace $L \subseteq \{0, 1\}^n$ lies in a kernel of some completion of A then $L \cap K_A = \emptyset$, by Lemma 5.1. Since $L + L = L$, the set L must be a solution for A , by Theorem 5.2.

(\Rightarrow): Let $L \subseteq \{0, 1\}^n$ be an arbitrary linear solution for A . Then $L + L = L$ and $L \cap K_A = \emptyset$. Take a $(0, 1)$ -matrix H with $L = \ker(H)$. Since $\ker(H) \cap K_A = \emptyset$, the matrix H separates K_A from zero. Lemma 6.1 implies that then \widehat{H} must contain some completion M of A . But then $L = \ker(H) = \ker(\widehat{H}) \subseteq \ker(M)$, as claimed. \square

Corollary 6.3. *For any $(0, 1, *)$ -matrix A we have that $\text{lin}(A) = 2^{n-\text{mr}(A)}$.*

Proof. By Theorem 6.2, $\text{lin}(A)$ is the maximum of $|\ker(M)| = 2^{n-\text{rk}(M)}$ over all completions M of A . Since $\text{mr}(A)$ is the minimum of $\text{rk}(M)$ over all completions M of A , we are done. \square

Corollary 6.4 (Alternative definition of min-rank). *For every $(0, 1, *)$ -matrix A we have*

$$\text{mr}(A) = \min\{\text{rk}(H) : Hx \neq \mathbf{0} \text{ for all } x \in K_A\}.$$

Proof. Let R be the smallest possible rank of a $(0, 1)$ -matrix separating K_A from zero. To prove $\text{mr}(A) \geq R$, let M be a completion of A with $\text{rk}(M) = \text{mr}(A)$. By Lemma 5.1, the matrix M separates K_A from zero. Hence, $R \leq \text{rk}(M) = \text{mr}(A)$.

To prove $\text{mr}(A) \leq R$, let H be a $(0, 1)$ -matrix such that H separates K_A from zero and $\text{rk}(H) = R$. By Lemma 6.1, the matrix \widehat{H} must contain a completion M of A . Hence, $\text{mr}(A) \leq \text{rk}(M) \leq \text{rk}(\widehat{H}) = \text{rk}(H) = R$. \square

By Lemma 5.1, the complement of K_A is the union of kernels $\ker(M)$ of all completions M of A . So, Theorems 5.2 and 6.2 imply that a subset $L \subseteq \{0, 1\}^n$ is:

- a solution for A iff $L + L \subseteq \bigcup \{\ker(M) : M \text{ is a completion of } A\}$;
- a linear solution for A iff $L \subseteq \ker(M)$ for some completion M of A .

7. Structure of general solutions

The following theorem says that non-linear solutions must be “very non-linear”: they cannot contain large linear subspaces. Recall that in Valiant’s setting (cf. Lemma 1.1) we may assume that each row of a $(0, 1, *)$ -matrix contains at most $s = n^\delta$ stars, where $\delta > 0$ is an arbitrary small constant.

Define the *co-distance* of a vector space as the smallest weight of a non-zero vector in its orthogonal complement.

Theorem 7.1. *Let $L \subseteq \{0, 1\}^n$ be a solution for an m -by- n $(0, 1, *)$ -matrix A , and let s be the maximum number of stars in a row of A . If L contains a linear subspace of co-distance at least $s + 1$, then L lies in a linear solution for A .*

Proof. Since L is a solution for A , W is a linear solution for A as well. Hence, by Theorem 6.2, W is contained in a kernel of some completion M of A . Our goal is to show that then the entire solution L must be contained in $\ker(M)$. To show this, we will use the following simple fact.

Claim 7.2. Let $W \subseteq \{0, 1\}^n$ be a linear subspace of co-distance at least $k + 1$. Then, for every k -element subset $S \subseteq [n]$ and for every vector $\mathbf{y} \in \{0, 1\}^n$, there is a vector $\mathbf{x} \in W$ such that $\mathbf{x} \neq \mathbf{0}$ and $\mathbf{y}|_S = \mathbf{x}|_S$.

Proof of Claim. The set of all projections of vectors in W onto S forms a linear subspace. If this subspace would be proper, then some non-zero vector, whose support lies in S , would belong to the orthogonal complement of W , a contradiction. \square

Assume now that $L \not\subseteq \ker(M)$, and take a vector $\mathbf{y} \in L \setminus \ker(M)$. Since $\mathbf{y} \notin \ker(M)$, we have that $\langle \mathbf{m}_i, \mathbf{y} \rangle = 1$ for at least one row \mathbf{m}_i of M . Let S be the set of star-positions in the i th row of A (hence, $|S| \leq s$), and let \mathbf{a}_i be this row of A with all stars set to 0. By Claim 7.2, there must be a vector $\mathbf{x} \in W \subseteq L \cap \ker(M)$ with $\mathbf{y}|_S = \mathbf{x}|_S$, that is, $D_i(\mathbf{x} \oplus \mathbf{y}) = \mathbf{0}$. But $\mathbf{x} \in \ker(M)$ implies that $\langle \mathbf{m}_i, \mathbf{x} \rangle = 0$. Hence, $\langle \mathbf{m}_i, \mathbf{x} \oplus \mathbf{y} \rangle = \langle \mathbf{m}_i, \mathbf{x} \rangle \oplus \langle \mathbf{m}_i, \mathbf{y} \rangle = \langle \mathbf{m}_i, \mathbf{y} \rangle = 1$. Since the vector \mathbf{a}_i can only differ from \mathbf{m}_i in star-positions of the i th row of A and, due to $D_i(\mathbf{x} \oplus \mathbf{y}) = \mathbf{0}$, the vector $\mathbf{x} \oplus \mathbf{y}$ has no 1's in these positions, we obtain that $\langle \mathbf{a}_i, \mathbf{x} \oplus \mathbf{y} \rangle = 1$. Hence, the vector $\mathbf{x} \oplus \mathbf{y}$ belongs to K_A , a contradiction with $\mathbf{x}, \mathbf{y} \in L$.

This completes the proof of Theorem 7.1. \square

8. Relation to codes

Let $1 \leq r < n$ be integers. A (binary) error-correcting code of minimal distance $r + 1$ is a set $C \subseteq \{0, 1\}^n$ of vectors, any two of which differ in at least $r + 1$ coordinates. A code is *linear* if it forms a linear subspace over GF_2 . The question on how good linear codes are, when compared to non-linear ones, is a classical problem in Coding Theory. We now will show that this is just a special case of a more general “opt(A) versus lin(A)” problem for $(0, 1, *)$ -matrices, and that Min-Rank Conjecture in this special case holds true.

An (n, r) -code matrix, or just an r -code matrix if the number n of columns is not important, is a $(0, 1, *)$ -matrix with n columns and $m = (r+1)\binom{n}{r}$ rows, each of which consists of $n-r$ stars and at most one 0. The matrix is constructed as follows. For every r -element subset S of $[n] = \{1, \dots, n\}$ include in A a block of $r+1$ rows \mathbf{a} with $a_i = *$ for all $i \notin S$, $a_i \in \{0, 1\}$ for all $i \in S$, and $|\{i \in S : a_i = 0\}| \leq 1$. That is, each of these rows has stars outside S and has at most one 0 within S . For $r = 3$ and $S = \{1, 2, 3\}$ such a block looks like

$$\begin{pmatrix} 1 & 1 & 1 & * & \cdots & * \\ 0 & 1 & 1 & * & \cdots & * \\ 1 & 0 & 1 & * & \cdots & * \\ 1 & 1 & 0 & * & \cdots & * \end{pmatrix}.$$

A Hamming ball around the all-0 vector $\mathbf{0}$ is defined by

$$\text{Ball}(r) = \{\mathbf{x} \in \{0, 1\}^n : 0 \leq |\mathbf{x}| \leq r\},$$

where $|\mathbf{x}| = x_1 + \dots + x_n$ is the number of 1's in \mathbf{x} .

Observation 8.1. If A is an r -code matrix, then $K_A = \text{Ball}(r) \setminus \{\mathbf{0}\}$.

Proof. Observe that no vector $\mathbf{x} \in \{0, 1\}^r$, $\mathbf{x} \neq \mathbf{0}$ can be orthogonal to all $r+1$ vectors $\mathbf{1}, \mathbf{1} \oplus \mathbf{e}_1, \dots, \mathbf{1} \oplus \mathbf{e}_r$ in $\{0, 1\}^r$ with at most one 0. Indeed, if $\langle \mathbf{x}, \mathbf{1} \rangle = 0$ then $\langle \mathbf{x}, \mathbf{1} \oplus \mathbf{e}_i \rangle = x_i$ for all $i = 1, \dots, r$. By this observation, a vector \mathbf{x} belongs to K_A iff there is an r -element set $S \subseteq [n]$ of positions such that $\mathbf{x}|_S \neq \mathbf{0}$ and $\mathbf{x}|_{\bar{S}} = \mathbf{0}$, that is, iff $\mathbf{x} \neq \mathbf{0}$ and $\mathbf{x} \in \text{Ball}(r)$. \square

Observation 8.2. *If A is an (n, r) -code matrix, then the solutions for A are error-correcting codes of minimal distance $r + 1$, and linear solutions for A are linear codes.*

Proof. We have $(L + L) \cap (\text{Ball}(r) \setminus \{\mathbf{0}\}) = \emptyset$ iff $|\mathbf{x} \oplus \mathbf{y}| \geq r + 1$ for all $\mathbf{x} \neq \mathbf{y} \in L$, that is, iff every two vectors $\mathbf{x} \neq \mathbf{y} \in L$ differ in at least $r + 1$ positions. Hence, every solution for an r -code matrix A is a code of minimal distance at least $r + 1$, and linear solutions are linear codes. \square

Lemma 8.3. *For code matrices, the min-rank conjecture holds with a constant $\epsilon > 0$.*

Proof. Let A be an (n, r) -code matrix; hence, $K_A = \text{Ball}(r) \setminus \{\mathbf{0}\}$. Set $t := \lfloor (r - 1)/2 \rfloor$. Since $|\mathbf{x} \oplus \mathbf{y}| \leq 2t < r$ for all $\mathbf{x}, \mathbf{y} \in \text{Ball}(t)$, the sum of any two vectors $\mathbf{x} \neq \mathbf{y}$ from $\text{Ball}(t)$ lies in K_A , implying that $\text{Ball}(t)$ is a clique in the Cayley graph generated by K_A . Since, by Remark 5.4, solutions for A are independent sets in this graph, and since in any graph the number of its vertices divided by the clique number is an upper bound on the size of any independent set, we obtain:

$$\text{opt}(A) \leq 2^n / |\text{Ball}(t)| = 2^n / \sum_{i=0}^t \binom{n}{i}, \quad (8)$$

which is the well-known Hamming bound for codes. On the other hand, Gilbert-Varshamov bound says that linear codes in $\{0, 1\}^n$ of dimension k and minimum distance d exist, if

$$\sum_{i=0}^{d-2} \binom{n-1}{i} < 2^{n-k}.$$

Hence,

$$\text{lin}(A) \geq 2^n / \sum_{i=0}^r \binom{n}{i}. \quad (9)$$

Together with (8), this implies that the inequality (1) holds with ϵ about $1/2$. \square

The example of code matrices also shows that the gap between min-rank and row/column min-rank may be at least logarithmic in n .

Lemma 8.4. *If A is an (n, r) -code matrix, then $\text{mr}(A) = \Omega(r \ln(n/r))$ but $\text{mr}_{\text{col}}(A) \leq r + 1$ and $\text{mr}_{\text{row}}(A) \leq 2r$.*

Proof. To prove $\text{mr}(A) = \Omega(r \ln(n/r))$, recall that $K_A = \text{Ball}(r) \setminus \{\mathbf{0}\}$. Hence, Corollary 6.4 implies that $\text{mr}(A)$ is the smallest possible rank of a $(0, 1)$ -matrix H such that $\ker(H) \cap \text{Ball}(r) \subseteq \{\mathbf{0}\}$. On the other hand, for any such matrix H , its kernel $L = \ker(H)$ is a (linear) code of minimal distance at least $r + 1$ containing $|L| = 2^{n-\text{rk}(H)}$ vectors. Since, by Hamming bound (8), no code L of distance at least $r + 1$ can have more than $N = 2^n / (n/r)^{O(r)}$ vectors, we have that

$$\text{rk}(H) = n - \log_2 |L| \geq n - \log_2 N = \Omega(r \ln(n/r)).$$

To prove that $\text{mr}_{\text{col}}(A) \leq r + 1$, suppose that A contains some $m \times k$ submatrix B of min-rank k . Since all k columns must be independent, at least one row \mathbf{b} of B must be $*$ -free and contain an odd number $|\mathbf{b}|$ of 1's. But every row of A (and hence, also \mathbf{b}) can contain at most one 0, implying that $|\mathbf{b}| \geq k - 1$. Together with $|\mathbf{b}| \leq r$, this implies that $k \leq r + 1$.

To prove that $\text{mr}_{\text{row}}(A) \leq 2r$, recall that each row of A consists of $n - r$ stars and at most one 0; the remaining r (or $r - 1$) entries are 1's. Suppose now that A contains some set X of $|X| = k + 1$ independent rows. That is, no subset of these rows can be made linearly dependent by setting *'s to 0 or 1. The rows in X must be, in particular, *pairwise* independent. This, in particular, means that the set X can contain at most one row without 0-entries. So, let $Y \subseteq X$ be a set of $|Y| = k$ rows containing 0-entries. Take any two rows $\mathbf{x} \neq \mathbf{y} \in Y$ with $x_i = 0$ and $y_j = 0$. Since \mathbf{x} and \mathbf{y} are independent and have only *'s or 1's outside their 0-entries, we have that: $i \neq j$ and either $x_j = 1$ or $y_i = 1$. This implies that the total number of 1's in the rows of Y must be at least the number $\binom{k}{2}$ of pairs of vectors in Y . So, there must exist a row $\mathbf{x} \in Y$ with $|\mathbf{x}| \geq \binom{k}{2}/|Y| = (k - 1)/2$. Together with $|\mathbf{x}| \leq r - 1$, this implies that $k \leq 2r - 1$, and thus, that $|X| = k + 1 \leq 2r$. \square

9. Conclusion and open problems

In this paper we pose a conjecture about systems of semi-linear equations and show its relation to proving super-linear lower bounds for log-depth circuits. We then give a support for the conjecture by proving that some its weaker versions are true. We also show that solutions are independent sets in particular Cayley graphs, thus turning the conjecture in a more general (combinatorial) setting. Using this, we prove several structural properties of sets of solutions that might be useful when tackling the original conjecture.

We defined solutions for a given m -by- n $(0, 1, *)$ -matrix A as sets $L \subseteq \{0, 1\}^n$ of vectors \mathbf{x} satisfying a system of equations

$$\langle \mathbf{a}_i, \mathbf{x} \rangle = g_i(D_i \mathbf{x}) \quad i = 1, \dots, m, \quad (10)$$

where \mathbf{a}_i is the i th row of A with all stars replaced by 0, g_i is an arbitrary boolean function, and D_i is a diagonal n -by- n $(0, 1)$ -matrix corresponding to stars in the i th row of A . We have also shown (see Remark 5.4) that solutions for A are precisely the independent sets in a Cayley graph over the Abelian group $(\{0, 1\}^n, \oplus)$ generated by a special set of vectors

$$K_A = \{\mathbf{x} : \exists i D_i \mathbf{x} = \mathbf{0} \text{ and } \langle \mathbf{a}_i, \mathbf{x} \rangle = 1\}. \quad (11)$$

The following two questions about possible generalizations of the min-rank conjecture naturally arise:

1. What if instead of diagonal matrices D_i in (10) we would allow other $(0, 1)$ -matrices?
2. What if instead of special generating sets K_A , defined by (11), we would allow other generating sets?

The following two examples show that the min-rank conjecture cannot be carried too far: its generalized versions are false.

Example 9.1 (Bad generating sets K). Let G be a Cayley graph generated by the set $K \subseteq \{0, 1\}^n$ of all vectors with more than $n - 2\sqrt{n}$ ones. If $L \subseteq \{0, 1\}^n$ consist of all vectors with at most $n/2 - \sqrt{n}$ ones, then $(L + L) \cap K = \emptyset$, that is, L is an independent set in G of size $|L| \geq 2^{n - O(\log n)}$. But any *linear* independent set L' in G is a vector space of dimension at most $n - 2\sqrt{n}$. Hence, $|L'| \leq 2^{n - 2\sqrt{n}}$, and the gap $|L|/|L'|$ can be as large as $2^{\Omega(\sqrt{n})}$.

Note, however, that there is a big difference between the set K we constructed and the sets K_A arising from $(0, 1, *)$ -matrices A : generating sets K_A must be almost "closed downwards". In particular, if $\mathbf{x} \in K_A$ then *all* nonzero vectors, obtained from \mathbf{x} by flipping some even number of its 1's to 0's, must also belong to K_A . Hence, this example does not refute the min-rank conjecture as such.

Example 9.2 (Bad matrices D_i). Let us now look what happens if we allow the matrices D_1, \dots, D_m in the definition of a system of semi-linear equations (10) be *arbitrary* $n \times n$ $(0, 1)$ -matrices. A completion M of A can then be defined as a $(0, 1)$ -matrix with rows $\mathbf{m}_i = \mathbf{a}_i + \mathbf{a}_i^\top D_i$. Now define $\text{mr}(A|D_1, \dots, D_r)$ as the minimal rank of such a completion of A . Observe that this definition coincides with the “old” min-rank, if we take the D_i ’s to be the diagonal matrices corresponding the stars in the i th row of A .

However, Example 9.1 shows that the min-rank conjecture is false in this generalized setting. To see why, we can define appropriate matrices A, D_1, \dots, D_m such that the corresponding set K_A defined by (11) consists of vectors with more than $n - 2\sqrt{n}$ ones: for an arbitrary vector \mathbf{v} with more than $n - 2\sqrt{n}$ ones just define \mathbf{a}_i and D_i such that the system $D_i \mathbf{x} = \mathbf{0}, \langle \mathbf{a}_i, \mathbf{x} \rangle = 1$ has \mathbf{v} as its only solution.

Except of the obvious open problem to prove or disprove the linearization conjecture (Conjecture 1) or the min-rank conjecture (Conjecture 2), there are several more concrete problems.

We have shown (Lemma 8.4) that the gap between min-rank and row/column min-ranks may be as large as $\ln n$. It would be interesting to find $(0, 1, *)$ -matrices A with larger gap.

Problem 9.3. *How large can the gap $\text{mr}(A) / \max\{\text{mr}_{\text{col}}(A), \text{mr}_{\text{row}}(A)\}$ be?*

The next question concerns the clique number $\omega(G_A)$ of (that is, the largest number of vertices in) Cayley graphs G_A generated by the sets of the sets $K_A \subseteq \{0, 1\}^n$ of the form (11). By Remark 5.4, solutions for A are independent sets in this graph. Hence, $\text{opt}(A)$ is just the independence number $\alpha(G_A)$ of this graph. Since in any N -vertex graph G we have that $\omega(G) \cdot \alpha(G) \leq N$, this yields $\text{opt}(A) \leq 2^n / \omega(G_A)$. On the other hand, it is easy to see that $\omega(G_A) \leq 2^{\text{rk}(M)}$, where M is a canonical completion of A obtained by setting all $*$ ’s to 0: If $C \subseteq \{0, 1\}^n$ is a clique in G_A , then we must have $M\mathbf{x} \neq M\mathbf{y}$ for all $\mathbf{x} \neq \mathbf{y} \in C$, because otherwise the vector $\mathbf{x} \oplus \mathbf{y}$ would not belong to K_A .

Problem 9.4. *Give a lower bound on $\omega(G_A)$ in terms of min-rank $\text{mr}(A)$ of A .*

Finally, it would be interesting to eliminate an annoying requirement in Theorem 4.11 that the matrix A must be star-monotone.

Problem 9.5. *If A is an r -by- n $(0, 1, *)$ -matrix of min-rank r , is then $\text{opt}(A) \leq 2^{n-r}$?*

Acknowledgment

We would like to thank the referee for a careful reading of the manuscript.

References

- [1] N. Alon, On the rigidity of an Hadamard matrix, manuscript, 1990.
- [2] N. Alon, P. Pudlák, Superconcentrators of depth 2 and 3; odd levels help (rarely), J. Comp. Sys. Sci. 48 (1994) 194–202.
- [3] N. Alon, M. Karchmer, A. Wigderson, Linear circuits over $\text{GF}(2)$, SIAM. J. Comput. 19(6) (1990) 1064–1067.
- [4] D. Y. Cherukhin, The lower estimate of complexity in the class of schemes of depth 2 without restrictions on a basis, Moscow Univ. Math. Bull. 60(4) (2005) 42–44.
- [5] N. Cohen, C. R. Johnson, L. Rodman, H. J. Woederman, Ranks of completions of partial matrices, Operator Theory: Adv. Appl., 40 (1989) 165–185.
- [6] D. Dolev, C. Dwork, N. Pippenger, A. Wigderson, Superconcentrators, generalizer and generalized connectors with limited depth, in: Proc. 15th Ann. ACM Symp. on Theory of Computing (STOC), 1983, pp. 42–51.
- [7] J. Friedman, A note on matrix rigidity, Combinatorica 13 (1993) 235–239.
- [8] C. R. Johnson, Matrix completion problems: a survey, in: C. R. Johnson, ed., Matrix Theory and Applications, Proc. of AMS Symp. in Applied Math., vol. 40 (1990), pp. 171–198.
- [9] S. Jukna, Entropy of operators or why matrix multiplication is hard for depth-two circuits, Theory of Comput. Syst. (2008), doi 10.1007/s00224-008-9133-y.

- [10] S. Lokam, On the rigidity of Vandermonde matrices, *Theoret. Comput. Sci.* 237(1-2) (2000) 477-483.
- [11] S. Lokam, Spectral methods for matrix rigidity with applications to size-depth trade-offs and communication complexity, *J. Comput. Syst. Sci.* 63(3) (2001) 449-473.
- [12] J. Morgenstern, Note on a lower bound on the linear complexity of fast Fourier transform, *J. ACM* 20(2) (1973) 305–306.
- [13] J. Morgenstern, The linear complexity of computation, *J. ACM* 22(2) (1975) 184–194.
- [14] N. Pippenger, Superconcentrators, *SIAM J. Comput.* 6 (1977) 298–304.
- [15] N. Pippenger, Superconcentrators of depth 2, *J. Comput. Syst. Sci.* 24 (1982) 82–90.
- [16] P. Pudlák, Communication in bounded depth circuits, *Combinatorica* 14(2) (1994) 203–216.
- [17] P. Pudlák, Z. Vavrník, Computation of rigidity of order n^2/r for one simple matrix, *Comment. Math. Univ. Carolinae* 32 (1991) 213–218.
- [18] P. Pudlák and P. Savický. On shifting networks. *Theoret. Comput. Sci.* 116 (1993) 415–419.
- [19] R. Paturi, P. Pudlák, Circuit lower bounds and linear codes, *J. Math. Sciences*, 134(5) (2006) 2425–2434.
- [20] P. Pudlák, V. Rödl, J. Sgall, Boolean circuits, tensor ranks, and communication complexity, *SIAM J. Comput.* 26(3) (1997) 605–633.
- [21] J. Radhakrishnan, A. Ta-Shma, Bounds for dispersers, extractors, and depth-two superconcentrators, *SIAM J. Discrete Math.* 13(1) (2000) 2–24.
- [22] R. Raz, A. Shpilka, Lower bounds for matrix product in bounded depth circuits with arbitrary gates, *SIAM J. Comput.* 32(2) (2003) 488–513.
- [23] A. A. Razborov, On rigid matrices, manuscript, 1989 (in Russian).
- [24] A. A. Razborov, B. Khasin, Improved lower bounds on the rigidity of Hadamard matrices, *Mat. Zametki* 63(4) (1998) 534–540 (in Russian).
- [25] M. A. Shokrollahi, D. A. Spielman, V. Stetmann, A remark on matrix rigidity, *Inf. Process. Letters* 64(6): 283–285, 1997.
- [26] R. de Wolf, Lower bounds on matrix rigidity via a quantum argument, in: *Proc. 33rd Int. Colloq. on Automata, Languages and Programming (ICALP'06)*, in: Springer Lect. Notes in Comput. Sci., vol. 4051 (2006), pp. 62–71.
- [27] L. Valiant, Graph-theoretic methods in low-level complexity, in: J. Gruska (Ed.), *Proc. 6th Symp. on Math. Foundations of Comput. Sci. (MFCS'77)*, in: Springer Lect. Notes in Comput. Sci., vol. 53 (1977), pp. 162–176.
- [28] L. Valiant, Why is boolean complexity theory difficult?, in: M. S. Paterson (Ed.), *Boolean Function Complexity*, Cambridge Univ. Press (1992) pp. 84-94.

Table 1: This table summarizes the concepts introduced in this paper. Here A is a partially defined $m \times n$ matrix with entries from $\{0, 1, *\}$.

Concept	Notation	Meaning
Completion of A		A $(0, 1)$ -matrix obtained from A by setting its $*$ -entries to 0 and 1.
Canonical completion of A		All $*$ -entries of A set to 0.
Min-rank	$\text{mr}(A)$	Minimal rank over GF_2 of a completion of A .
Max-rank	$\text{Mr}(A)$	Maximal rank over GF_2 of a completion of A .
Operator G consistent with A		The i th coordinate of $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ can only depend on variables corresponding to $*$ -entries in the i th row of A .
Solution for A		A set $L \subseteq \{0, 1\}^n$ of the form $L = \{\mathbf{x} : M\mathbf{x} = G(\mathbf{x})\}$, where M is a completion of A , and G is an operator consistent with A .
Linear solution for A		A solution for A forming a linear subspace of $\{0, 1\}^n$.
	$\text{opt}(A)$	Maximum size of a solution for A .
	$\text{lin}(A)$	Maximum size of a linear solution for A ; $\text{lin}(A) = 2^{n-\text{mr}(A)}$.
Min-Rank Conjecture		$\text{opt}(A) \leq 2^{n-\epsilon \cdot \text{mr}(A)}$ for a constant $\epsilon > 0$.
Independence of $(0, 1, *)$ -vectors		Cannot be made linear dependent by setting $*$'s to constants.
Row min-rank	$\text{mr}_{\text{row}}(A)$	Maximal number of independent rows.
Column min-rank	$\text{mr}_{\text{col}}(A)$	Maximal number of independent columns.
Incidence matrix of $*$'s	D_i	Diagonal $(0, 1)$ -matrix with $D_i[j, j] = 1$ iff $A[i, j] = *$.
Set of forbidden vectors	K_A	All vectors $\mathbf{x} \in \{0, 1\}^n$ such that $D_i\mathbf{x} = 0$ and $\langle \mathbf{a}_i, \mathbf{x} \rangle = 1$, where \mathbf{a}_i is the i th row of A with all stars set to 0. Main property: L is a solution for A iff $(L + L) \cap K_A = \emptyset$.