On Graph Complexity

S. JUKNA[†]

Universität Frankfurt, Institut für Informatik
Robert-Mayer-Str. 11-15, D-60054 Frankfurt, Germany
jukna@thi.informatik.uni-frankfurt.de
and
Institute of Mathematics and Informatics
Akademijos 4, LT-08663 Vilnius, Lithuania

By the *complexity* of a graph we mean the minimum number of union and intersection operations needed to obtain the whole set of its edges starting from stars. This measure of graphs is related to the circuit complexity of boolean functions.

We prove some lower bounds on the complexity of explicitly given graphs. This yields some new lower bounds for boolean functions, as well as new proofs of some known lower bounds in the graph-theoretic frame. We also formulate several combinatorial problems whose solution would have intriguing consequences in computational complexity.

1. Introduction

A major challenge in computational complexity is to exhibit an *explicit* boolean function $f_m : \{0,1\}^m \to \{0,1\}$ that has high computational complexity, i.e. cannot be computed using a small number of basic boolean operations (gates) such as OR $x \vee y$, AND $x \wedge y$ or Parity $x \oplus y = x + y \pmod 2$; inputs for such a circuit are literals, i.e. variables x_i and their negations $\overline{x}_i$. Though this problem is intensively studied for more than fifty years there is no proof of a lower bound super-linear in the number of variables m. The main difficulty here is that we want the function f_m be *explicitely constructed*—easy counting shows that almost all functions require circuits of size $2^{\Omega(m)}$. The problem of proving super-linear lower bounds is widely open even if we assume the additional restriction of circuit depth be logarithmic in m.

Pudlák, Rödl and Savický [36] observed that in oder to construct boolean functions requiring large circuits it would be enough to construct graphs that require many union and intersection operations to represent them starting from some "simplest" graphs, like stars. Although simple, this observation is important because it allows to translate the lower bounds problems for boolean functions to purely combinatorial problems for

[†] Research supported by a DFG grant SCHN 503/2-2.

graphs. For example, some old problems in computational complexity could be solved by proving non-trivial lower bounds on the following generalisation of the edge clique covering number (we show this in Section 5):

$\mathrm{cov}(G) =_{\mathrm{df}}$ minimal t for which there exist t (spanning) subgraphs of G such that each edge of G is an edge of at least one of the subgraphs, and the non-edges of each of these subgraphs can be covered by at most t independent sets.

It is not difficult to show (by standard counting argument) that $\mathrm{cov}(G) = \Omega(n^{1/2})$ for almost all n-vertex graphs. The problem, however, is to exhibit an *explicit* n-vertex graph with $\mathrm{cov}(G) \geq n^{\epsilon}$ for some constant $\epsilon > 0$. Actually, such a lower bound even with ϵ tending slowly to 0 would resolve some long-standing open problems in computational complexity. If proved with $\epsilon = \omega(1/\sqrt{\log n})$, this would give an explicit boolean function in m variables requiring depth-3 circuits of size $2^{\omega(\sqrt{m})}$. If proved with $\epsilon = (\log\log n)^{\omega(1)}/\log n$, this would give an explicit boolean function outside the second level of the communication complexity hierarchy introduced in [6]. If proved with $\epsilon = \omega(1/\log\log\log n)$ this would give a super-linear lower bound for log-depth circuits.

Motivated by this connection between graphs and boolean functions, in this paper we are trying to better understand what properties of graphs make them hard to represent by different kinds of circuits. One of the aims of this paper is to draw once more readers attention to graph complexity. This concept has already led to interesting results [8, 36, 39, 38, 33, 28], and its potential seems to be far from being exhausted.

Graph complexity. We shall use standard graph theory notation. Of particular interest for us will be bipartite graphs. We shall look at such a graph G with a fixed bipartition $V = U \cup W$ as a set $G \subseteq U \times W$ of its edges. A *non-edge* is a pair uv of non-adjacent vertices. A non-edge in a bipartite graph $G \subseteq U \times W$ is a pair uv of non-adjacent vertices with $u \in U$ and $v \in W$; hence, pairs of vertices within one part of a bipartition are neither edges nor non-edges. If not stated otherwise, by a subgraph we will always mean a *spanning* subgraph.

Given a graph $G = (V, E)$ we associate to each its vertex v a boolean variable x_v, and let $X = \{x_v : v \in V\}$. We say that a boolean function (or a circuit) $f(X)$ accepts/rejects a subset of vertices $S \subseteq V$ if f accepts/rejects the characteristic vector of S, i.e. a binary vector in $\{0,1\}^V$ with 1's in positions u for all $u \in S$, and 0's elsewhere.

We say that a boolean function $f(X)$ *represents* a graph G if it accepts all edges and rejects all non-edges of G.

Hence, $f(X)$ represents the graph G if for every input vector $a \in \{0,1\}^X$ with precisely two 1's in, say, positions u and v, $f(a) = 1$ if uv is an edge, and $f(a) = 0$ if uv is a non-edge of G. Note that if uv is neither an edge nor a non-edge (in the bipartite case) or if a contains more or less that two 1's, then the value $f(a)$ may be arbitrary.

For example, a single variable x_v represents a star around (a set of all edges incident with) the vertex v. An OR $\bigvee_{v\in S} x_v$ represents a union of stars which, in turn, is a complement of a complete graph on $V \setminus S$. The formula $\left(\bigvee_{u\in S} x_u\right) \wedge \left(\bigvee_{v\in T} x_v\right)$ with $S \cap T = \emptyset$ represents a bipartite complete graph $S \times T$.

Every graph $G = (V, E)$ can be represented by a monotone formula $\bigvee_{uv\in E} x_u x_v$ with

$|E|+1$ gates: $|E|$ AND gates of fanin 2 and one OR gate of fanin $|E|$. Hence, no n-vertex graph G of maximal degree d requires circuits of size larger than dn. However, this trivial upper bound may be exponentially far from the truth: a result of Alon [2] on the clique covering number of graphs (see Theorem 5.3 below) implies that every n-vertex graph G of maximal degree d can be represented by a monotone CNF (conjunctive normal form) $\bigwedge_{i=1}^{r} \bigvee_{v \in S_i} x_v$ with $r = O(d^2 \log n)$. In particular, every graph of constant degree can be represented by a monotone depth-2 formula of logarithmic size.

From graphs to boolean functions. The *complexity* of a graph (in a given class of boolean circuits) is the minimum number of gates in a circuit representing this graph. Although estimating the circuit complexity of graphs may be of independent interest, we (just like the authors of [36]) consider the graph complexity mainly as a tool for proving lower bounds for boolean functions.

The translation of lower bounds for graphs to lower bounds for boolean functions is given by the following lemma (we give its proof in the Appendix). With every bipartite $n \times n$ graph $G \subseteq U \times W$, where $n = 2^m$ and $U = W = \{0,1\}^m$, one may associate a boolean function f in $2m$ variables—the *characteristic function* of G—such that $f(uv) = 1$ if and only if $uv \in G$. In the lemma below, by a *circuit* we will mean an arbitrary computational model whose inputs are literals, i.e. boolean variables x_i and their negations $\overline{x}_i$.

Lemma 1.1 (Magnification Lemma). *Given a circuit computing the characteristic function f of a bipartite $n \times n$ graph G, it is possible to replace each its input literal by an OR of at most n variables so that the obtained circuit represents the graph G. The same holds when Parity gates are used instead of OR gates.*

This fact is particularly useful in such circuit models where computing an OR (or a Parity) of input literals is "cheap." For example, if the circuit computing f has unbounded fanin OR gates on the bottom (next to the inputs) level, then the obtained (monotone) circuit represents G and has just the same number of gates! Hence, if we could prove that a bipartite $n \times n$ graph G with $n = 2^m$ cannot be represented using, say, fewer than n^ϵ gates, this would immediately imply that the characteristic function f of G requires at least $n^\epsilon = 2^{\epsilon m}$ gates, which is *exponential* in the number $2m$ of variables of f_m (this is where the term "magnification" comes from).

Note, however, that proving lower bounds for graphs may be even more difficult task than for boolean functions. For example, the Parity function $x_1 \oplus x_2 \oplus \cdots \oplus x_m$ cannot be computed by a constant-depth circuit using a polynomial number of unbounded fanin AND and OR gates ([19]) whereas the corresponding to this function graph is just a union of two vertex-disjoint bipartite complete graphs, and can be represented by a circuit using just seven gates! This also demonstrates that the Magnification Lemma has no inverse: if a graph can be represented by a small circuit, this does not imply that its characteristic function can be computed by a small circuit. Still, studying the graph-theoretic structure of boolean functions may provide new insights into their complexity—unlike for boolean functions, the structure of graphs is much better understood.

As it is, the Magnification Lemma only holds for bipartite graphs. Still, lower bounds for general (non-bipartite) graphs would also yield lower bounds for boolean functions, because every graph $G = (V, E)$ on $2n$ vertices is a union of $O(\log n)$ bipartite $n \times n$ graphs of the form $E \cap (U \times W)$ with $|U| = |W| = n$.

Why even depth-3 circuits are interesting? In this paper we are mainly interested in proving lower bounds on graph complexity in the class of depth-3 circuits. Although this restriction seems to be rather severe, non-trivial lower bounds even in this class of circuits would already resolve some old problems in the circuit complexity of boolean functions.

Of particular interest is the case of Σ_3 circuits. These circuits consist of unbounded fanin AND and OR gates which are organised in three levels: the bottom (next to the inputs) level consists of OR gates, the middle level consists of AND gates, and the top level consists of a single OR gate. Inputs are variables and their negation. If there are no negated inputs then the circuit is *monotone*. Hence, a monotone Σ_3 circuit has the form

$$\bigvee_{i=1}^{s} \bigwedge_{j=1}^{r} \bigvee_{u \in S_{ij}} x_u .$$

Here s is the *top fanin* and r the *middle fanin* of a circuit; by the *size* of a circuit we will mean the maximum $\max\{s, r\}$ of its top and middle fanins.

Our motivation to study representation of graphs by depth-3 circuits comes from the following result due to Valiant [41]: if a boolean function f in m variables can be computed by a log-depth circuit of size $O(m)$ then f can be computed by a Σ_3 circuit of size $2^{O(m/\log\log m)}$; here a log-depth circuit is a circuit of depth $O(\log m)$ using any boolean functions in constant number of variables as gates. Together with the Magnification Lemma, this implies the following: if a bipartite $n \times n$ graph cannot be represented by a monotone Σ_3 circuit of size n^{ϵ} with $\epsilon = O(1/\log\log\log n)$, then its characteristic function cannot be computed by a log-depth circuit of linear size.

In last two decades there was a considerable progress in proving lower bounds on the size of small-depth circuits [1, 14, 42, 19, 37, 40, 15, 21, 31]. However, for Σ_3 circuits these bounds are of the form $2^{\Omega(\sqrt{m})}$, and hence, are too weak to imply lower bounds for log-depth circuits. The only known strongly exponential lower bounds were obtained in [32] under the restriction that the bottom OR gates have fanin 2, that is, when the circuit is just an OR of 2-CNFs. However, Valiant's reduction requires bottom fanin m^{ϵ} and, as noted in [32], their argument fails already when bottom fanin is larger than 2.

Our goal is to obtain higher lower bounds for small depth circuits using the graph-theoretical frame. So far, we have not succeeded to do this for "pure" Σ_3 circuits but are able to do this for some of their variants.

2. Depth-2 circuits

To "warm-up" we start with the simplest model of circuits whose gates are arranged in two levels. Such circuits are easy to deal with, and the only goal of this section is to show that, even in this model, there may be a big discrepancy between the *combinatorial* and *computational* complexity of graphs: some "combinatorially complicated" (or "combinatorially interesting") graphs can be represented by very small circuits and some "combinatorially simple" graphs require large circuits (of the same type).

Example 1. The *Kneser graph* $K(r, k)$ $(r > 2k \geq 4)$ has all k-element subsets v of $\{1, \ldots, r\}$ as vertices, and two vertices are adjacent iff the corresponding k-subsets are disjoint. These graphs were introduced by Lovász [29] in his famous proof of Kneser's conjecture [25] that whenever the k-subsets of a $(2k+s)$-set are divided into $s+1$ classes, then two disjoint subsets end up in the same class. It is not difficult to see that $K(r, k)$ can be represented by the following depth-2 circuit:

$$\bigwedge_{i=1}^{r} \bigvee_{v \in S_i} x_v \tag{2.1}$$

where $S_i = \{v : i \notin v\}$. Indeed, $u \neq v$ are non-adjacent in $K(r, k)$ iff $u \cap v \neq \emptyset$ iff $\{u, v\} \cap S_i = \emptyset$ for some i iff uv is rejected by some OR $\bigvee_{v \in S_i} x_v$. Note that with respect to the total number $n = \binom{r}{k}$ of vertices the representation is quite compact: the circuit has only $1 + r = O(kn^{1/k})$ gates (r OR gates and one AND gate).

Example 2. An *Hadamard matrix* of oder n is an $n \times n$ matrix with entries ± 1 and with row vectors mutually orthogonal. A graph associated with an Hadamard matrix M (or just an Hadamard graph) of oder n is a bipartite $n \times n$ graph H_n where two vertices u and v are adjacent if and only if $M(u, v) = +1$. An example of an Hadamard graph is the *Sylvester* graph $S(n)$. This is a bipartite $n \times n$ graph with $n = 2^r$ vertices on each part identified with subsets of $\{1, \ldots, r\}$; two vertices u and v are adjacent iff $|u \cap v|$ is odd. It is easy to see that (for even r) this graph can be represented by a depth-2 circuit

$$\bigoplus_{i=1}^{r} \bigvee_{v \in S_i} x_v \tag{2.2}$$

with $S_i = \{v : i \notin v\}$. Indeed, u and v are adjacent in $S(n)$ iff $|u \cap v|$ is odd iff $r - |u \cap v|$ is odd iff the number of sets S_i containing at least one of u and v is odd iff the number of clauses $\bigvee_{v \in S_i} x_v$ accepting uv is odd. Again, the representation is quite compact: the circuit has only [1] $r + 1 = \log(2n)$ gates (r OR gates and one Parity gate of fanin $r + 1$). On the other hand, each Hadamard graph (including the graph $S(n)$) is "combinatorially complicated" because, as shown in [36] (see also [35]), it contains an induced subgraphs on $\sqrt{n}$ vertices which is Ramsey, meaning that it does not contain cliques or independent sets of size $\omega(\log n)$. By setting the corresponding variables in the circuit (2.2) to 0, we obtain that this Ramsey graph can be represented by a depth-2 circuit of size $O(\log n)$.

Razborov in [39] developed a general probabilistic machinery allowing to show that a whole string of other "combinatorially complicated" n-vertex graphs can be represented by constant-depth circuits using a small (polynomial in $\log n$) number of unbounded fanin AND and Parity gates (see Lemma 8.1 below).

[1] All logarithms in this paper are to the basis of 2.

On the other hand, some "combinatorially simple" graphs—like an n to n matching M_n (a set on n vertex disjoint edges) or its complement—cannot be represented by depth-2 circuits using fewer than $\Omega(n)$ gates. If a graph G is represented by a depth-2 circuit of the form (2.1) with top fanin r then its complement is just a union of r cliques. Hence, $r = \Omega(n)$ for any such circuit representing $\overline{M_n}$. High lower bounds for depth-2 circuits of the form (2.2) can be obtained via simple rank argument: every such circuit representing a graph G must have size at least $\mathrm{rk}(G)/2$ where $\mathrm{rk}(G)$ is the rank over $GF(2)$ of the adjacency matrix of G (just because each graph represented by an OR gate is a complement of a clique, and hence, has rank at most 2). Hence, $r = \Omega(n)$ for any circuit of the form (2.2) representing M_n.

3. Depth-3 circuits with Parity gates

We have just seen that proving high lower bounds for *depth*-2 circuits is an easy task. However, already the case of *depth*-3 circuits turns out to be much more difficult. And this is not surprising because, as we already mentioned above, lower bounds on the size of monotone Σ_3 circuits of the form n^{ϵ} (where ϵ may even tend slowly to 0) would resolve some old problems in computational complexity.

Pudlák, Rödl and Savický [36] asked whether dense C_4-free bipartite graphs could be good candidates. In a somewhat weaker (but still enough to imply super-linear lower bounds for log-depth circuits) form their question can be stated as follows.

Problem 3.1. Is there a constant $c > 0$ such that every bipartite C_4-free $n \times n$ graph with M edges requires monotone Σ_3 circuits of size at least $(M/n)^c$?

In this section we show that this question has an *affirmative* answer for a modified version of Σ_3 circuits where Parity gates (instead of OR gates) are used on the bottom level; we call such circuits $\Sigma_3^{\oplus}$ *circuits.* Such a circuit of top fanin s and middle fanin r is just an OR of s boolean functions, each of which is a product of r linear forms over $GF(2)$:

$$\bigvee_{i=1}^{s}\bigwedge_{i=1}^{r}\bigoplus_{u\in S_{ij}} x_u \oplus \lambda_{ij}$$

where $\lambda_{ij} \in \{0,1\}$. The *top fanin* is the fanin s of the last OR gate, i.e. the number of AND gates. If all scalars λ_{ij} are equal 0, the circuit is *positive.*

Remark 3.2. If we would require that the top gate of a circuit must also be a Parity gate (not an OR gate), then *truly exponential* lower bounds $2^{\Omega(m)}$ for such version of $\Sigma_3^{\oplus}$ circuits could be obtained using the algebraic (approximation by low-degree polynomials) techniques of [37, 40, 15]. However, these techniques seem incapable of proving truly exponential lower bounds for $\Sigma_3^{\oplus}$ circuits themselves because, in this case, we would be forced to approximate the top OR gate as well, which would invariably result in the square root $2^{\Omega(\sqrt{m})}$ in the final bound.

In what follows, $K_{a,b}$ denotes a *biclique* (bipartite clique, complete bipartite graph) $A \times B$, $A \cap B = \emptyset$ with parts of size $a = |A|$ and $b = |B|$.

Theorem 3.3. *Let $G \subseteq U \times W$ be a bipartite $n \times n$ graph. If G contains no copy of $K_{a,b}$, then any $\Sigma_3^{\oplus}$ circuit representing G (and hence, any $\Sigma_3^{\oplus}$ circuit computing the characteristic function of G) has top fanin at least*

$$\frac{|G|}{(a+b)n}.$$

To prove the theorem, we first give a combinatorial characterisation of the top fanin of $\Sigma_3^{\oplus}$ circuits and then give a general lower bound on this characteristic.

A *fat matching* is a union of vertex-disjoint bipartite cliques (these cliques need not to cover all vertices). A *fat covering* of a graph G is a family of fat matchings such that each of these fat matchings is a subgraph of G and every edge of G is an edge of at least one member of the family. Let $\mathrm{fat}(G)$ denote the minimum number of fat matchings in a fat covering of G. This measure was also considered by several authors, [10, 13, 2, 34] among others (fat matchings are often called "equivalence graphs"). In particular, it is known that $\mathrm{fat}(G) = O(n/\log n)$ for every n-vertex graph [34].

Lemma 3.4. *For every bipartite graph $G \subseteq U \times W$, $\mathrm{fat}(G)$ equals to the minimum top fanin of a $\Sigma_3^{\oplus}$ circuit representing G.*

Proof. First note that in the case of graphs we can safely restrict ourselves to positive circuits, because $\bigoplus_{u \in A \cup \overline{B}} x_u$ with $A \subseteq U$ and $B \subseteq W$ represents the same graph as $1 \oplus \bigoplus_{u \in A \cup B} x_u$.

Let $g = \bigoplus_{v \in A \cup B} x_v$ be a gate on the bottom level of a $\Sigma_3^{\oplus}$ circuit representing G. Then g represents a fat matching $(A \times \overline{B}) \cup (\overline{A} \times B)$ where $\overline{A} = U \setminus A$ and $\overline{B} = W \setminus B$. Since the intersection of any number of fat matchings is a fat matching, each AND gate on the middle level represents a fat matching. Hence, if the circuit has top fanin s, then the OR gate on the top represents a union of these s fat matchings, implying that $s \geq \mathrm{fat}(G)$.

To show that G can be represented by a $\Sigma_3^{\oplus}$ circuit of top fanin $\mathrm{fat}(G)$, let $M = \bigcup_{i=1}^{r} A_i \times B_i$ be a fat matching. If $A = \bigcup_{i=1}^{r} A_i$ and $B = \bigcup_{i=1}^{r} B_i$, then M is an intersection of $A \times B$ with r fat matchings of the form $H_i = A_i \times B_i \cup \overline{A_i} \times \overline{B_i}$. Since $A \times B$ can be represented by an AND of two Parity gates $\bigoplus_{u \in A} x_u$ and $\bigoplus_{v \in B} x_v$, and each H_i can be represented by the Parity gate $\bigoplus_{v \in A_i \cup \overline{B_i}} x_v$, every fat matching can be represented by an AND of Parity gates. Hence, every graph G can be represented by a $\Sigma_3^{\oplus}$ circuit of top fanin $\mathrm{fat}(G)$. □

Lemma 3.5. *Let $G \subseteq U \times W$ be a bipartite $n \times n$ graph. If G contains no copy of $K_{a,b}$ then*

$$\mathrm{fat}(G) \geq \frac{|G|}{(a+b)n}.$$

Proof. Let $H = \bigcup_{i=1}^{t} A_i \times B_i$ be a fat matching, and suppose that $H \subseteq G$. By the definition of a fat matching, the sets $A_1, \ldots, A_t$, as well as the sets $B_1, \ldots, B_t$ are mutually disjoint. Moreover, since G contains no copy of $K_{a,b}$, we have that $|A_i| < a$ or $|B_i| < b$ for all i. Hence, if we set $I = \{i : |A_i| < a\}$, then

$$|H| = \sum_{i=1}^{t} |A_i \times B_i| = \sum_{i=1}^{t} |A_i| \cdot |B_i| \le \sum_{i \in I} a \cdot |B_i| + \sum_{i \notin I} |A_i| \cdot b \le (a+b)n.$$

Thus, no fat matching $H \subseteq G$ can cover more than $(a+b)n$ edges of G, implying that we need at least $|G|/(a+b)n$ fat matchings to cover all edges of G. □

There are many explicit bipartite $n \times n$ graphs which are dense enough and do not have large bicliques. Theorem 3.3 immediately yields *truly exponential* lower bounds (i.e. lower bounds of the form $2^{\Omega(m)}$) on the top fanin of $\Sigma_3^{\oplus}$ circuits computing the characteristic functions of these graphs; recall that these functions have only $2m = 2\log n$ variables. Here we restrict ourselves with few examples.

The *disjointness function* is a boolean function $DISJ_{2m}$ in $2m$ variables such that

$$DISJ_{2m}(y_1, \ldots, y_m, z_1, \ldots, z_m) = 1 \text{ if and only if } \sum_{i=1}^{m} y_i z_i = 0.$$

Corollary 3.6. *Every $\Sigma_3^{\oplus}$ circuit computing $DISJ_{2m}$ has top fanin $2^{\Omega(m)}$.*

Proof. The function $DISJ_{2m}$ is the characteristic function of the Kneser-type bipartite graph $K(m) \subseteq U \times V$ where U and W consist of all $n = 2^m$ subsets of $[m] = \{1, \ldots, m\}$, and $uv \in K(m)$ iff $u \cap v = \emptyset$. The graph $K(m)$ can contain a complete bipartite $a \times b$ subgraph $\emptyset \neq A \times B \subseteq K$ only if $a \le 2^k$ and $b \le 2^{m-k}$ for some $0 \le k \le m$, because then $\left(\bigcup_{u \in A} x_u\right) \cap \left(\bigcup_{v \in B} x_v\right) = \emptyset$. In particular, $K(m)$ can contain a copy of $K_{a,a}$ only if $a \le 2^{m/2} = \sqrt{n}$. Since this graph has

$$|K(m)| = \sum_{u \in U} d(u) = \sum_{u \in U} 2^{m-|u|} = \sum_{i=0}^{m} \binom{m}{i} 2^{m-i} = 3^m \ge n^{1.58}$$

edges, Theorem 3.3 yields that any $\Sigma_3^{\oplus}$ circuit representing $K(m)$—and hence, any $\Sigma_3^{\oplus}$ circuit computing $DISJ_{2m}$—must have top fanin at least $|K(m)|/(2an) = \Omega(n^{0.08}) = 2^{\Omega(m)}$. □

Remark 3.7. In the context of boolean functions, $\Sigma_3^{\oplus}$ circuits *cannot* be efficiently simulated by Σ_3 circuits: the Parity function $x_1 \oplus x_2 \oplus \cdots \oplus x_m$ has an obvious $\Sigma_3^{\oplus}$ circuit of size 1, whereas (as shown in [19]) this function requires Σ_3 circuits of size $2^{\Omega(\sqrt{m})}$. It may be, therefore, interesting to note that in the context of graphs the situation is entirely different: if a graph can be represented by a $\Sigma_3^{\oplus}$ circuit of size L then G can be represented by a monotone Σ_3 circuit of size at most $2L$. This holds because we can just replace each parity gate $\bigoplus_{u \in S} x_u$ on the bottom level by an AND $\left(\bigvee_{u \in S} x_u\right) \wedge \left(\bigvee_{u \notin S} x_u\right)$ of two OR gates; the obtained monotone Σ_3 circuit will represent the same graph. Moreover, the graph $K(m)$ shows that $\Sigma_3^{\oplus}$ circuits may be even *exponentially weaker*: this graph can be represented by a monotone Σ_3 circuit of size $O(\log n)$ (see (2.1)) but requires $\Sigma_3^{\oplus}$ circuits of size at least $\Omega(n^{\epsilon})$.

A prominent example of a dense bipartite graph without $K_{2,2}$ was constructed by Erdős and Rényi [11]: this is the incidence $n \times n$ graph $P_n \subseteq U \times W$ of a projective plane $PG(2,q)$ of order q $(n = q^2 + q + 1)$. This graph is $(q+1)$-regular and has no copies of $K_{2,2}$. According to the well-known construction of $PG(2,q)$ (which can be found in any textbook on finite geometries), the characteristic function π_{2m} of P_n is just the boolean version of the function $f : GF(q)^6 \to \{0,1\}$ defined by: $f(x,y,z,a,b,c) = 1$ if and only if $ax + by + cz = 0$ modulo q. Since the graph P_n has $\Omega(n^{3/2})$ edges, Theorem 3.3 yields

Corollary 3.8. *Every $\Sigma_3^{\oplus}$ circuit computing π_{2m} has top fanin $\Omega(2^{m/2})$.*

For every constant $a > 1$ explicit constructions of $n \times n$ graphs (so-called *norm-graphs*) with $\Omega(n^{2-1/a})$ edges and no copies of $K_{a,a!+1}$ were found by Kollár, Rónyai and Szabó in [27]; explicit graphs without $K_{r,s}$ but for somewhat larger values of r and s were earlier constructed by Andreev [5]. For the characteristic functions f_{2m}^{a} of these graphs, Theorem 3.3 yields

Corollary 3.9. *For every constant $a > 1$, every $\Sigma_3^{\oplus}$ circuit computing f_{2m}^{a} has top fanin $\Omega(2^{m-1/a})$.*

The only previously known truly exponential lower bound for $\Sigma_3^{\oplus}$ circuits we are aware of was proved by Grolmusz [16] for the Inner Product function

$$IP_{2m}(y_1, \ldots, y_m, z_1, \ldots, z_m) = \sum_{i=1}^{m} y_i z_i \pmod 2.$$

Quite recently Pudlák and Rödl [35] have also proved such a lower bound for the characteristic functions of certain pseudorandom sets. Both proofs employ non-trivial facts—the probabilistic communication complexity of IP in [16] and some properties of pseudorandom sets in [35].

Actually, the lower bounds in [16] and [35] were proved for a more general model of $\Sigma_3^{\oplus}$ circuits: instead of an OR gate they allow an arbitrary threshold gate on the top level. (Recall that a threshold-k function accepts an input iff it contains at least k 1's.) Let us show that lower bounds for this extended model can be proved in the context of graphs as well.

At this point, it is worth to note that in some cases it can even make sense to *reprove* known lower bounds for boolean functions in the frame of graphs. For example, reproving known lower bound $2^{\Omega(\sqrt{m})}$ for Σ_3 circuits—or even proving a much weaker lower bound $2^{(\log m)^{\omega(1)}}$—in the graph-theoretic frame would give us a graph outside the second level of the communication complexity hierarchy introduced in [6] (see Problem 8.3 below).

Corollary 3.10. *Any $\Sigma_3^{\oplus}$ circuit which has an arbitrary threshold gate on the top and represents an $n \times n$ Hadamard graph must have top fanin $\Omega(\sqrt{n})$.*

Since the inner product function IP_{2m} is the characteristic function of an Hadamard $n \times n$ graph H_n with $n = 2^m$, Corollary 3.10 and the Magnification Lemma immediately yield a lower bound $\Omega(2^{m/2})$ for IP_{2m} in this (more general than $\Sigma_3^{\oplus}$) class of circuits.

For the proof of Corollary 3.10 we need the so-called "discriminator lemma" for threshold gates (used also in [16, 35]). Let $\mathcal{B}$ be a family of subsets of a finite set X. For a subset $A \subseteq X$, let $\mathrm{thr}_{\mathcal{B}}(A)$ denote the minimum number t for which there exist t members $B_1, \ldots, B_t$ of $\mathcal{B}$ and a number $0 \leq k \leq t$ such that, for every $x \in X$, $x \in A$ if and only if x belongs to at least k of B_i's. A set A is an *ϵ-discriminator* for a set B if

$$\left| \frac{|A \cap B|}{|A|} - \frac{|\overline{A} \cap B|}{|\overline{A}|} \right| \geq \epsilon.$$

Lemma 3.11. ([17]) *If* $\mathrm{thr}_{\mathcal{B}}(A) \leq t$ *then* A *is a* $1/t$*-discriminator for some* $B \in \mathcal{B}$*.*

Proof. Let $B_1, \ldots, B_t \in \mathcal{B}$ be a threshold-k covering of A, i.e. $x \in A$ iff x belongs to at least k of B_i's. Our goal is to show that then A is a $1/t$-discriminator for at least one B_i. Since every element of A belongs to at least k of the sets $A \cap B_i$, the average size of these sets must be at least k. Since no element of $\overline{A}$ belongs to more than $k-1$ of the sets $\overline{A} \cap B_i$, the average size of these sets must be at most $k-1$. Hence,

$$1 \quad \leq \quad \sum_{i=1}^{t} \frac{|A \cap B_i|}{|A|} - \sum_{i=1}^{t} \frac{|\overline{A} \cap B_i|}{|\overline{A}|} \leq t \cdot \max_{1 \leq i \leq t} \left| \frac{|A \cap B_i|}{|A|} - \frac{|\overline{A} \cap B_i|}{|\overline{A}|} \right|.$$

□

Proof of Corollary 3.10. Let A be an $n \times n$ Hadamard graph. Lindsey's lemma (see, e.g. [3] or [6]) says that the absolute value of the difference between the number of $+1'$s and $-1'$s in any $a \times b$ submatrix of A is at most $\sqrt{abn}$. Since both A and $\overline{A}$ have $\Theta(n^2)$ edges, by Lemmas 3.4 and 3.11, it is enough to show that $\big||A \cap B| - |\overline{A} \cap B|\big| = O(n^{3/2})$ for every fat matching $B = \bigcup_{i=1}^{t} S_i \times R_i$. By Lindsey's lemma, the absolute value of the difference between $|A \cap (S_i \times R_i)|$ and $|\overline{A} \cap (S_i \times R_i)|$ does not exceed $\sqrt{s_i r_i n}$ where $s_i = |S_i|$ and $r_i = |R_i|$. Since, $\sum_{i=1}^{t} s_i \leq n$ and $\sum_{i=1}^{t} r_i \leq n$, we obtain

$$\begin{aligned} \Big| |A \cap B| - |\overline{A} \cap B| \Big| \quad &= \quad \left| \sum_{i=1}^{t} |A \cap (S_i \times R_i)| - \sum_{i=1}^{t} |\overline{A} \cap (S_i \times R_i)| \right| \\ &\leq \quad \sum_{i=1}^{t} \sqrt{s_i r_i n} \leq \sqrt{n} \sum_{i=1}^{t} \frac{s_i + r_i}{2} \leq n^{3/2}. \end{aligned}$$

□

4. A tradeoff for Σ_3 circuits

We now use the graph theoretic frame to prove a trade-off between top and middle fanins in Σ_3 circuits, where middle fanin of a circuit is the maximum fanin of a gate in the middle level.

Theorem 4.1. *If IP_{2m} is computed by a Σ_3 circuit with top fanin s and middle fanin r, then both $s2^r$ and r^s must be at least $\Omega(2^m)$.*

A trade-off $sr = \Omega\left(m^3/(\log m)^5\right)$ between these parameters for IP_{2m} was recently proved by Lokam [28] (also using a graph-theoretic frame). The trade-off in Theorem 4.1 is better only if one of the parameters r or s is at most m^ϵ—the second parameter must then be at least $2^{\Omega(m^{1-\epsilon})}$.

A *clique covering* of a graph G is a family of complete subgraphs of G such that every edge of G is an edge of at least one member of the family; if the graph G is bipartite, then we take complete bipartite subgraphs. The minimum number of such subgraphs is known as the (bipartite) *clique covering number*, and is denoted by $\mathrm{cc}(G)$. This measure was first studied by Erdős, Goodman and Pósa [12], and now is the subject of extensive literature. A complement of a graph G is the graph $\overline{G}$ on the same set of vertices whose edges are non-edges of G and vice versa.

Since, by Lindsey's lemma, the clique covering number of an $n \times n$ Hadamard graph and of its complement is $\Omega(n)$, Theorem 4.1 follows directly from the following lemma.

Lemma 4.2. *If a bipartite graph G can be represented by a monotone Σ_3 circuit of middle fanin r and top fanin s, then $\mathrm{cc}(G) \le s2^r$ and $\mathrm{cc}(\overline{G}) \le r^s$.*

Proof. Take a monotone Σ_3 circuit of middle fanin at most r and top fanin s, and let $G \subseteq U \times W$ be the bipartite graph represented by this circuit. Each gate $g = \bigvee_{i \in S} x_i$ on the bottom level represents a (bipartite) complement of a bipartite clique $A \times B$, where $A = U \setminus S$ and $B = W \setminus S$. Each such complement is a union of two bipartite cliques $A \times \overline{B}$ and $\overline{A} \times W$. Since the intersection of any number of bipartite cliques is a (possibly empty) bipartite clique, each AND gate on the middle level represents a union of at most 2^r bipartite cliques. Since G is a union of s such graphs, we have $\mathrm{cc}(G) \le s2^r$.

To prove $\mathrm{cc}(\overline{G}) \le r^s$, observe that $\overline{G}$ is an intersection of s graphs $H_1, \ldots, H_s$, each of which is a union of r bipartite cliques. Since the intersection of any number of bipartite cliques is a bipartite clique, we have $\mathrm{cc}(\overline{G}) \le \prod_{i=1}^{s} \mathrm{cc}(H_i) \le r^s$. □

5. Σ_3 circuits and the clique covering number

In this section we give a combinatorial characterisation of graphs represented by monotone Σ_3 circuits. Recall that each such circuit of size t is an OR of at most t monotone CNFs of length t, where a monotone CNF conjunctive normal form) of length t is an AND $\left(\bigvee_{u \in S_1} x_u\right) \wedge \cdots \wedge \left(\bigvee_{u \in S_t} x_u\right)$ of t clauses, each of which is an OR of variables.

Let $\mathrm{cnf}(G)$ denote the minimum length of a monotone CNF representing G. This measure can be described combinatorially in terms of the clique covering number as well as in terms of set-intersections.

A graph G admits an *intersection representation of size* t if each vertex u can be associated with a subset $A_u \subseteq \{1, \ldots, t\}$ such that $A_u \cap A_v = \emptyset$ if uv is an edge, and $A_u \cap A_v \neq \emptyset$ if uv is a non-edge of G. Let $\mathrm{int}(G)$ denote the smallest t for which G admits such a representation.

Proposition 5.1. *For every graph $G = (V, E)$ we have* $\mathrm{cnf}(G) = \mathrm{cc}(\overline{G}) = \mathrm{int}(G)$.

Proof. An OR of variables $\bigvee_{u \in S} x_u$ with $S \subseteq V$ represents a complement of a clique (complete graph) on $V \setminus S$, and each such complement can be represented by an OR gate. Hence, a graph G can be represented by a CNF of length t iff G is an intersection of complements of t complete graphs, or equivalently, iff the complement $\overline{G}$ can be represented as a union of t cliques, implying that $\mathrm{cnf}(G) = \mathrm{cc}(\overline{G})$.

The equality $\mathrm{cc}(\overline{G}) = \mathrm{int}(G)$ is also easy to show; it was observed already in [12]. Given an intersection representation of G by subsets A_u of $\{1, \ldots, t\}$, the t sets $I_i = \{u : i \in A_u\}$ are independent and cover all non-edges of G. On the other hand, given a covering of the non-edges of G by independent sets $I_1, \ldots, I_t$, one can take $A_u = \{i : u \in I_i\}$. □

Using the intersection representation of graphs, it may be easily shown that some simple graphs (like an n to n matching M_n) have short CNFs. For example, to show that $\mathrm{cnf}(M_n) = O(\log n)$, let $t = 2 \log n$ and associate with each vertex u_i on the left side its *own* $(t/2)$-element subset A_i of $\{1, \ldots, t\}$, and assign to the unique matched vertex v_i on the right side the complement $B_i = \overline{A_i}$ of this subset. It is clear that then $A_i \cap B_j = \emptyset$ iff $i = j$. Hence, $\mathrm{cnf}(M_n) = \mathrm{int}(M_n) \leq t = 2 \log n$.

Remark 5.2. The same argument as for M_n yields an upper bound $\mathrm{cnf}(H) = O(\log n)$ for every $n \times n$ fat matching H. Hence, in the case of graphs, $\Sigma_3^{\oplus}$ circuits are no more powerful than monotone Σ_3 circuits with logarithmic middle fanin.

By Proposition 5.1, the number $\mathrm{cov}(G)$ (defined in the introduction) is precisely the size (the maximum of the top and middle fanins) of a monotone Σ_3 circuit representing G. The equality $\mathrm{cnf}(G) = \mathrm{int}(G)$ gives an equivalent algebraic characterisation of this number:

$\mathrm{cov}(G) =$ smallest number t such that each vertex u can be associated with a $t \times t$ 0-1 matrix A_u so that $uv \in G$ iff the product $A_u A_v^\top$ has at least one 0 on the diagonal.

Since we have only 2^{t^2} such matrices, we can encode at most 2^{2nt^2} of all 2^{n^2} bipartite $n \times n$ graphs; hence, $\mathrm{cov}(G) = \Omega(n^{1/2})$ for almost all graphs. This also yields that $\mathrm{cov}(G) = \Omega\left((\log n)^{1/2}\right)$ for every graph, in which no two vertices have the same set of neighbours (we need different matrices for different vertices). An *explicit* graph G with $\mathrm{cov}(G) \geq n^{\epsilon}$ for some constant $\epsilon > 0$ would resolve some old problems in computational complexity. However, so far we do not know of any lower bound substantially larger than $\log n$. Logarithmic lower bounds $\Omega(\log n)$ are easy to obtain (take, e.g. an n to n matching), and the best what we know so far is a slightly larger lower bound of the order $(\log n)^{3/2 - o(1)}$ proved by Lokam in [28] for an Hadamard graph.

Proposition 5.1, together with an obvious observation that every bipartite clique $A \times B$ can be represented by a CNF consisting of two clauses $\bigvee_{u \in A} x_u$ and $\bigvee_{v \in B} x_v$, gives an upper bound:

$$\mathrm{cov}(G) \leq \min\left\{\mathrm{cc}(G), \mathrm{cc}(\overline{G})\right\}. \tag{5.1}$$

A general upper bound on $\mathrm{cov}(G)$ for graphs of small degree can be obtained from Proposition 5.1 and the following result.

Theorem 5.3 (Alon [2]). *For every n-vertex graph G of maximal degree d, $\mathrm{cc}(\overline{G}) = O(d^2 \log n)$.*

Hence, if G has maximum degree d, then $\mathrm{cnf}(G) = O(d^2 \log n)$. In particular, every graph of constant degree can be represented by a monotone CNF of logarithmic length. This also implies an upper bound $\mathrm{cov}(G) = O(d^{2/3} \log n)$: simply break G into $d^{2/3}$ subgraphs of maximal degree $d^{1/3}$ each.

6. Σ_3 versus Π_3 circuits

As mentioned above, no explicit n-vertex graphs requiring monotone Σ_3 circuits of size $(\log n)^{\omega(1)}$ are known. On the other hand, if we replace the ANDs by ORs and vice versa, then the situation is much easier. The obtained "dual" circuits are known as Π_3 circuits and have the form:

$$\bigwedge_{i=1}^{s} \bigvee_{j=1}^{r} \bigwedge_{v \in S_{ij}} x_v;$$

by the size of such a circuit we again mean $\max\{s, r\}$.

It is worth to mention that, in the context of boolean functions, proving lower bounds for Σ_3 circuits is the same as proving lower bounds for the dual model of Π_3 circuits: if a function is hard in the former model then its negation is hard in the later. However, the following theorem shows that in the context of graphs the situation is different: if a graph is hard for (monotone) Π_3 circuits, then we cannot conclude that its complement must be also hard for (monotone) Σ_3 circuits.

Theorem 6.1. *Let M_n be an n to n matching. Then both the graph M_n and its complement $\overline{M}_n$ can be represented by monotone Σ_3 circuits of size $O(\log n)$, but every monotone Π_3 circuit representing $\overline{M}_n$ must have size at least $\Omega(\sqrt{n})$.*

A larger lower bound on the size of monotone Π_3 circuits can be obtained for Hadamard graphs.

Theorem 6.2. *Every monotone Π_3 circuit representing an Hadamard graph of oder n must have size at least $\Omega(n^{2/3})$.*

We derive both theorems from the following property of graphs represented by monotone Π_3 circuits.

Lemma 6.3. *Suppose that a bipartite graph G can be represented by a monotone Π_3 circuit of size t. Then it is possible to add to $\overline{G}$ a set E of $|E| \leq t^2$ edges so that $\mathrm{cc}(\overline{G} \cup E) \leq t$.*

Proof. Suppose that a graph $G \subseteq U \times W$ can be represented by a monotone Π_3 circuit of size t. Such a circuit is an AND of at most t monotone DNFs $D_1, \ldots, D_t$, each containing at most t monomials (ANDs of variables). Since we are interested in the behaviour of the circuit only on arcs (edges and non-edges), we may assume that none of these monomials contains more than two variables. Hence, each of the DNFs

$$D_i = \bigvee_{u \in S_i} x_u \vee \bigvee_{uv \in F_i} x_u x_v$$

accepts some set $S_i \subseteq U \cup W$ of vertices and some set F_i of $|F_i| \leq t$ arcs. Let $E = \bigcup_{i=1}^t E_i$ where $E_i = F_i \cap G$ is the set of edges of G accepted by the i-th DNF; hence, $|E| \leq t^2$. We may assume that the set $G \setminus E$ of remaining edges is non-empty, since otherwise we would have $E = G$, meaning that $\overline{G} \cup E$ is just a complete graph. By what was said, the CNF $\big(\bigvee_{u \in S_1} x_u\big) \wedge \cdots \wedge \big(\bigvee_{u \in S_t} x_u\big)$ must represent the graph $G \setminus E$. Hence, by Proposition 5.1, $\mathrm{cc}(\overline{G} \cup E) = \mathrm{cc}(\overline{G \setminus E}) = \mathrm{cnf}(G \setminus E) \leq t$. □

Proof of Theorem 6.1 We already know (see Section 5) that both M_n and $\overline{M_n}$ have monotone Σ_3 circuits of size $O(\log n)$. So, it remains to show that $\overline{M_n}$ requires large monotone Π_3 circuits.

Let t be the minimum size of a monotone Π_3 circuit representing $\overline{M}_n$. Then, by Lemma 6.3, it must be possible to add a set E of $|E| \leq t^2$ edges to the matching M_n so that the resulting graph $M_n \cup E$ can be covered by at most t bicliques. At least one of these bicliques, say $A \times B$, must contain at least $|M_n|/t = n/t$ edges of the matching M_n. But this means that $|E \cap (A \times B)| \geq (n/t)^2 - (n/t)$. Together with $|E| \leq t^2$ this implies that t must satisfy the inequality $(n/t)^2 - (n/t) \leq t^2$, that is, $t^4 \geq n^2 - tn$, which implies $t = \Omega(\sqrt{n})$. □

Proof of Theorem 6.2 Let t be the minimum size of a monotone Π_3 circuit representing a bipartite $n \times n$ Hadamard graph $H = H_n$. We may assume that $t \leq n/16$, for otherwise there is nothing to prove. We will use the known fact that any Hadamard graph contains about the same number of edges and non-edges; in particular, both $|H|$ and $|\overline{H}|$ are at least $n^2/4$.

By Lemma 6.3, there is a set E of $|E| \leq t^2$ edges such that the graph $\overline{H} \cup E$ can be covered by at most t bicliques $R_1, \ldots, R_t$, that is, $\overline{H} \cup E = R_1 \cup \cdots \cup R_t$. Let $N = |\overline{H}|$ be the total number of non-edges in H (hence, $N \geq n^2/4$) and take a biclique $R \in \{R_1, \ldots, R_t\}$ containing the largest number of non-edges of H; hence, $N_0 := |R \cap \overline{H}| \geq N/t$. Let $N_1 := |R \cap H|$ be the number of edges of H lying in R. Since $R \cap H$ can contain only edges from E, we have that $N_1 \leq |E| \leq t^2$. On the other hand, by Lindsey's lemma, $|N_1 - N_0| \leq \sqrt{n|R|}$, implying that $N_1 \geq N_0 - \sqrt{n|R|}$. Remembering that $N_1 + N_0 = |R| \geq N/t \geq n^2/4t \geq 4n$, we obtain

$$2N_1 \geq |R| - \sqrt{n|R|} = |R|\left(1 - \sqrt{\frac{n}{|R|}}\right) \geq \frac{N}{2t},$$

that is, $N_1 \geq N/(4t)$. Together with $N_1 \leq t^2$, this implies that $t^3 \geq N/4$. Thus, t must be at least $(N/4)^{1/3} \geq (n^2/16)^{1/3} = \Omega(n^{2/3})$. □

7. Boolean formulas

In this section we consider circuits of *arbitrary* depth with unbounded fanin AND and OR gates; as before, inputs are literals (variables and their negations). A *formula* is a circuit with all gates having fanout 1, i.e. the underlying graph in this case is just a tree. The *length* of a formula is the number of input literals.

Given a boolean function f and a graph G, let $L(f)$ (resp., $L_+(f)$) be the minimum length of a formula (resp., monotone formula) *computing* f, and $L_+(G)$ the minimum length of a monotone formula *representing* G.

If F is a formula computing the characteristic function f (in $2m$ variables) of a bipartite $n\times n$ graph G (with $n = 2^m$) then, by the Magnification Lemma, we can replace each input literal in F by a monotone formula of length at most n (computing the corresponding OR of variables) so that the resulting monotone formula recognises G. Thus,

$$L(f) \geq L_+(G)/n.$$

Easy counting shows that $L_+(G) = \Omega(n^2/\log n)$ for most $n \times n$ graphs G. Pudlák, Rödl and Savický have proved in [36] that $L_+(G) = \Omega(n\log(n/a))$ for any $n \times n$ graph G such that neither G nor its complement contains a copy of $K_{a,a}$. But, so far, no *explicit* graph with $L_+(G) = \Omega(n\log^3 n)$ is known. Such a graph would improve the strongest currently known lower bound $\Omega(m^{3-o(1)})$ on the (non-monotone) formula length of an explicit boolean function in m variables [20].

The reason, why it is difficult to show that a given graph $G = (V, E)$ cannot be represented by a short (monotone!) formula F, is that we only know that the formula must behave correctly on the 2-*element* subsets of vertices: for all $S \subseteq V$ with $|S| \leq 2$

$$F(S) = 0 \ \text{ if and only if } S \text{ is an independent set in } G. \tag{7.1}$$

On larger sets the formula may output arbitrary values. In particular, it can accept independent sets of size $k \geq 3$.

In this section we look what happens if we require that the formula F must reject independent sets only up to some size $k \geq 2$. That is, this time we require that (7.1) must hold for all subsets $S \subseteq V$ of size $|S| \leq k$. Note that the quadratic function

$$f_G(X) = \bigvee_{uv\in E} x_u x_v$$

rejects *all* independent sets of G, but the corresponding formula has length $2|E|$. Can we essentially decrease the length of the formula by relaxing this condition and requiring that it must reject only independent sets up to some size k for $k < n$? Using a rank-argument it can be shown that, for some graphs, this is *not* possible unless k is smaller than two times the maximal degree of G.

Theorem 7.1. *Let $G = (V, E)$ be a triangle-free graph without 4-cycles and of maximal degree d. Let f be a monotone boolean function which accepts all edges and rejects all independent sets of G of size at most $2d$. Then $L_+(f) \geq |E|/2$. In particular, $L_+(f_G) = \Theta(|E|)$.*

Proof. We look at vertices as one element and edges as two element sets. For a vertex $y \in V$, let I_y be the set of its neighbours. For an edge $y \in E$, let I_y be the set of all its *proper* neighbours; that is, $v \in I_y$ precisely when $v \notin y$ and v is adjacent with an endpoint of y. Since G has no triangles and no 4-cycles, the sets I_y are independent sets of size at most $2d$, and must be rejected by f. We will concentrate only on these independent sets.

Let M be a matrix whose rows correspond to edges $x \in E$, columns to vertices and edges $y \in V \cup E$, and

$$M_{x,y} = x \setminus I_y.$$

A *rectangle* in M is a submatrix $A \times B \subseteq M$ with the property that there is a vertex v such that

$$v \in x \setminus I_y \text{ for all } x \in A \text{ and } y \in B;$$

we call v a common element of the rectangle. Let $\mathcal{R}$ be a smallest possible set of mutually disjoint rectangles covering the whole matrix M. It is well known that every monotone formula computing f has length at least $|\mathcal{R}|$ (see [38, 22]). Hence, it remains to prove that $|\mathcal{R}| \geq |E|/2$.

To do this, re-fill the entries of M with constants 0 and 1 by the following rule:

$$M_{x,y} = 1 \text{ if and only if } x \cap y \neq \emptyset \tag{7.2}$$

Let $R = A \times B$ be a rectangle in $\mathcal{R}$, and let v be its common element. Then $v \in x$ for all edges $x \in A$ and $v \notin I_y$ for all $y \in B$. Hence, for each $y \in B$, the corresponding column in R is either the all-1 column (if $v \in y$) or the all-0 column (if $v \notin y$) because in this last case the second endpoint of x cannot belong to y (for otherwise, the first endpoint v would belong to I_y). Thus, either the rectangle R is monochromatic or we can split it into two monochromatic rectangles. In this way we obtain a covering $\mathcal{R}'$ of M by at most $2|\mathcal{R}|$ mutually disjoint monochromatic rectangles. To estimate their number we use the rank argument. Let $\mathrm{rk}(M)$ stand for the rank of M over $GF(2)$. Since the rectangles in $\mathcal{R}'$ are mutually disjoint and have rank 1, it follows that $|\mathcal{R}'| \geq \mathrm{rk}(M)$. Hence, it remains to prove that M has full row-rank over $GF(2)$.

Take an arbitrary subset $\emptyset \neq F \subseteq E$ of edges. We have to show that the rows of the submatrix M_F of M corresponding to the edges in F cannot sum up to the all-0 row over $GF(2)$. If F is not an even factor, that is, if the number of edges in F containing some vertex v is odd, then the column of v in M_F has an odd number of 1's, and we are done. Hence, we may assume that F *is* an even factor. Take an arbitrary edge $y = uv \in F$, and let $H \subseteq F$ be the set of edges in F incident to at least one endpoint of y. Since both vertices u and v have even degree (in F), the edge y has a nonempty intersection with an *odd* number of edges in F: one intersection with itself and an even number of intersections with the edges in $H \setminus \{y\}$. Hence, the colum of y in M_F contains an odd number of 1's, as desired. □

For the incidence $n \times n$ graph P_n of a projective plane $PG(2, q)$, Theorem 7.1 yields

Corollary 7.2. $L_+(f_{P_n}) = \Theta(n^{3/2})$.

Note that if we would only know that the formula must reject non-edges (independent sets of size 2)—the case interesting in the context of boolean functions—then the same rank argument with the matrix M defined by the rule (7.2) would not work. In this case we would have that $M_{x,y} = 1$ if and only if $|x \cap y| = 1$ (edge and non-edge can share at most one vertex). That is, M would be just a matrix of scalar products (over the reals) of the characteristic vectors of edges x and non-edges y, and (even over the reals) the rank of M would not exceed n.

8. Open problems

In the context of this paper the most interesting (and, perhaps, most realistic) problem remains to prove that some explicit graph requires large monotone Σ_3 circuits. Since, by Eq. (5.1), both the graph and its complement must then have large clique covering number, graphs with good Ramsey properties could (apparently) be possible candidates for this purpose. Say that a (bipartite) graph is (a, b)-Ramsey if it contains no copy of K_a (resp., $K_{a,a}$) and its complement contains no copy of K_b (resp., $K_{b,b}$). An indication that (a, b)-Ramsey graphs with *both* parameters a and b small (say, logarithmic in n) may be not good enough is given by the following result about the power of depth-3 circuits. In terms of graphs, this result can be stated as follows.

Let $G(m, r)$ be a random graph on $V = \{1, \ldots, n\}$ represented by a random depth-3 formula of the form

$$\bigoplus_{i=1}^{r} \bigwedge_{j=1}^{m} \bigoplus_{k=1}^{n} (\lambda_{ijk} x_k \oplus \lambda_{ij}) \tag{8.1}$$

where $\{\lambda_{ijk}, \lambda_{ij}\}$ are independent random variables uniformly distributed in $\{0, 1\}$.

Lemma 8.1 (Razborov [39]). *Let H be a graph on k vertices. If $\binom{k}{2} \leq 2^{m-1}$, then $G(m, r)$ contains a copy of H as an induced subgraph with probability at most*

$$\binom{n}{k} \left[2^{-\binom{k}{2}} + e^{-r/2^m} \right].$$

By this lemma, some (a, b)-Ramsey graphs with $a = b = \Theta(\log n)$ *can* be represented by depth-3 circuits (8.1) of poly-logarithmic size. Together with Theorem 3.3, this also shows that using a Parity gate (instead of an OR gate) on the top of $\Sigma_3^{\oplus}$ circuits may exponentially increase the power of such circuits. Note, however, that Lemma 8.1 does *not* work for (a, b)-Ramsey graphs if one of the parameters a or b is,say, constant, just because then the term $2^{-\binom{k}{2}} = \Omega(1)$ is too large. Thus, dense bipartite C_4-free graphs—like the plane graph P_n, which is (a, b)-Ramsey with $a = 2$ and $b = O(n^{3/4})$ [3]—could still be good candidates. Explicit constructions of (non-bipartite) (a, b)-Ramsey graphs with $a = 3$ and $b = O(n^{2/3})$ are also known [4]; moreover, these graphs are d-regular with $d = \Theta(n^{2/3})$. It would be interesting to know whether such graphs are hard for Σ_3 circuits (cf. Problem 1 in [39]).

We known (see Example 2) that the Sylvester graph $H = S(n)$ can be represented by a depth-2 circuit of the form $\bigoplus_{i=1}^{r} \bigvee_{v \in S_i} x_v$ with $r = O(\log n)$ gates. Although small, this circuit uses a Parity gate and it remains not clear what is the Σ_3 circuit complexity of this graph and, in particular, what is $\mathrm{cov}(H)$. The characteristic function of H (the Inner Product function) *can* be computed by a trivial circuit of linear size and logarithmic depth. Together with Valiant's result (mentioned in Introduction) this implies that H *can* be represented by a monotone Σ_3 circuit using $n^{o(1)}$ gates. On the other hand, Lokam [28] has proved that $\Omega\left((\log n)^3/(\log\log n)^5\right)$ gates are necessary.

Problem 8.2. Does $\mathrm{cov}(H) = \exp\left(\omega(\sqrt{\log n})\right)$?

If true, this would imply that the Inner Product function IP_{2m} requires Σ_3 circuits of size $2^{\omega(\sqrt{m})}$, thus improving the highest known lower bound $2^{\Omega(\sqrt{m})}$ for such circuits.

Another interesting question is to find graphs whose complements can be represented by depth-3 circuits of much smaller size. Let $\mathcal{E}$ be the set of all bipartite $n \times n$ graphs representable by monotone Σ_3 circuits of size $\exp\left((\log\log n)^{O(1)}\right)$. Let co$-\mathcal{E}$ be the set of complements of graphs from $\mathcal{E}$.

Problem 8.3. Prove that $\mathcal{E} \neq$ co$-\mathcal{E}$.

This would separate the second level of the communication complexity hierarchy introduced by Babai, Frankl and Simon in [6], and hence, resolve a long-standing open question in communication complexity. Note that in the case of Π_3 circuits this problem has a positive solution: an n to n matching M_n can be represented by a monotone Π_3 circuit (in fact, by a CNF) of size $O(\log n)$, but any such circuit for $\overline{M_n}$ requires size $\Omega(\sqrt{n})$ (see Theorem 6.1).

As we have shown in Section 5, the size of Σ_3 circuits can be characterised using the intersection representation of graphs. Good lower bounds on the following related measure would also resolve an old question in boolean circuit complexity.

$d(G) =_{\mathrm{df}}$ minimal number d for which there exists a set $L \subseteq \{0, 1, \ldots\}$ of integers such that each vertex u can be associated with a subset $A_u \subseteq \{1, \ldots, d\}$ so that two vertices u and v are adjacent in G if and only if $|A_u \cap A_v| \in L$.

Hence, $d(G) \leq \mathrm{int}(G)$ because $\mathrm{int}(G)$ is a special case of this measure for $L = \{0\}$. Again, by standard counting, $d(G) = \Omega(n)$ for almost all bipartite $n \times n$ graphs G. Graphs G with $d(G) = \Omega(\log n)$ are also easy to find—such is, for example, an n to n matching M_n. For some subsets L, high lower bounds are also easy to obtain. Say, if $L = \{ip : i = 0, 1, \ldots\}$ for some prime power p, then $d(\overline{M_n}) = \Omega(n)$ can be shown by an easy linear algebra argument. The difficult thing therefore is to do this for *all* possible choices of L.

Problem 8.4. Exhibit a bipartite $n \times n$ graph G with $d(G) = \exp\left((\log\log n)^{\omega(1)}\right)$.

Together with the Magnification Lemma and the results of Yao [43], and Beigel and Tarui [7], this would yield a super-polynomial lower bound for so-called *ACC circuits* computing the characteristic function of G. Such circuits have unbounded fanin AND, OR and MOD_k gates, where $\mathrm{MOD}_k(x_1, \ldots, x_m) = 1$ iff $x_1 + \ldots + x_m = 0$ modulo k. When k is a prime power, exponential lower bounds for such circuits were proved by Razborov [39] and Smolensky [40]. However, the case of composite moduli k—even the case of circuits with AND, OR and MOD_6 gates—remains widely open.

A more ambitious task is to prove non-trivial lower bounds on the *projective* and/or *affine* dimensions of graphs. Given a field F, these measures are defined by:

$\mathrm{pdim}_F(G) =_{\mathrm{df}}$ minimal d such that each vertex u can be associated with a *projective* subspace A_u of F^d such that $uv \in G$ iff $A_u \cap A_v \neq \{\mathbf{0}\}$;

$\mathrm{adim}_F(G) =_{\mathrm{df}}$ minimal d such that each vertex u can be associated with an *affine* subspace A_u of F^d such that $uv \in G$ iff $A_u \cap A_v \neq \emptyset$.

These measures were introduced, respectively, by Pudlák and Rödl [33], and by Razborov [38] as a tool to prove lower bounds for branching programs and formulas: if G is a bipartite graph and f is its characteristic function f, then $L(f) \geq \mathrm{adim}_F(G)$ ([38]), and $\mathrm{pdim}_F(G)$ is a lower bound on the size of branching programs computing f ([33]). Karchmer and Wigderson [23] have shown that $\mathrm{adim}_F(G)$ is also a lower bound on the size of span programs computing f. In order to improve the best known lower bounds for these models we need an explicit bipartite graph G with $\mathrm{pdim}_F(G) = \Omega\left(\log^2 n\right)$ and/or with $\mathrm{adim}_F(G) = \Omega\left(\log^3 n\right)$. Even a lower bound $\mathrm{adim}_F(G) = \Omega\left(\alpha \log n\right)$ with α growing faster than the iterated logarithm of n would improve the best lower bound for span programs, given in [23]. Lovasz's subspace version of Bollobá's theorem [30] yields $\mathrm{adim}_F(\overline{M_n}) = \Omega(\log n)$ and, so far, this is the only nontrivial lower bound.

In the proof of Theorem 7.1 we have mentioned yet another possibility to obtain lower bounds on $L(f)$. Given a graph G, let M be a matrix whose rows are labelled by edges and columns by non-edges of G; the (x, y)-th entry is $M_{x,y} = x \setminus y$. That is, each entry of M is either a single vertex or a pair of adjacent vertices. As before, a rectangle in M is a submatrix $A \times B \subseteq M$ with the property that there is a vertex v such that $v \in x$ and $v \notin y$ for all $x \in A$ and $y \in B$.

Problem 8.5. Exhibit an n-vertex graph G with $R(G) = \Omega\left(n \log^a n\right)$.

If proved with $a = 2$ this would give a graph-theoretic proof of Khrapchenko's classical lower bound $\Omega(m^2)$ on the size of non-monotone formula on m variables [24]. If proved with $a \geq 3$ this would improve the strongest currently known lower bound $\Omega(m^{3-o(1)})$ due to Håstad [20].

Appendix: Proof of the Magnification Lemma

In the lemma below, by a *circuit* we will mean an arbitrary computational model whose inputs are literals, i.e. variables $x_i^1 = x_i$ and their negations $x_i^0 = \overline{x}_i$. A boolean function is *isolating* if it rejects the all-0 vector $(0, \ldots, 0)$ and accepts all vectors containing precisely one 1; on other vectors the function can take arbitrary values. Since OR and Parity functions are isolating, the Magnification Lemma is a special case of Lemma 8.6 below.

Let $G \subseteq U \times W$ be a bipartite graph with $U = W = \{0,1\}^m$, and

$$f(y_1, \dots, y_m, z_1, \dots, z_m)$$

be its characteristic function; that is, $f(uv) = 1$ iff $uv \in G$. Suppose we have a circuit F computing f. A *positive extension* of F has 2^{m+1} variables $\{x_u : u \in U\} \cup \{x_v : v \in W\}$, and is obtained from F by replacing input literals y_i^a and z_i^a by functions

$$Y_i^a = g\left(\{x_u : u \in U,\, u(i) = a\}\right) \quad \text{and} \quad Z_i^a = h\left(\{x_v : v \in W,\, v(i) = a\}\right)$$

where g and h are arbitrary isolating functions, and $u(i)$ is the i-th bit of $u \in \{0,1\}^m$.

Lemma 8.6. *Let $G \subseteq U \times W$ be a bipartite $n \times n$ graph. If a circuit F computes the characteristic function of G, then every its positive extension F^+ represents the graph G.*

Proof. For an arc $uv \in U \times W$, let $a_{u,v}$ be the vector in $\{0,1\}^{U \cup W}$ with precisely two 1's in positions u and v. Let F be a circuit computing the characteristic function of G. Then $uv \in G$ iff $F(uv) = 1$. Hence, it is enough to show that $F^+(a_{u,v}) = 1$ iff $F(uv) = 1$.

The only difference of the circuit F^+ from F is that instead of input literals it takes the corresponding isolating functions as inputs. Hence, it is enough to show that on an input vector $a_{u,v}$ these isolating functions output the same values as the corresponding literals do on the input vector uv. We show this only for y-literals (for z-literals the argument is the same).

Let y_i^a be some input literal of F, and $u, v \in \{0,1\}^m$. By the definition, the function $Y_i^a = g\left(\{x_u : u \in U,\, u(i) = a\}\right)$ depends only on the variables x_u corresponding to the left part U of the bipartition such that $u(i) = a$. Each input of the form $a_{u,v}$ assigns precisely one 1 to these variables, and this 1 is in the position x_u. Hence, $Y_i^a(a_{u,v}) = 1$ iff Y_i^a depends on x_u, which can happen if and only if $u(i) = a$. On the other hand, we also have that $y_i^a(uv) = 1$ if and only if $u(i) = a$. Thus, $Y_i^a(a_{u,v}) = y_i^a(uv)$, and we are done. □

Acknowledgements

I am grateful Noga Alon and Alexander Razborov for interesting discussions, and to an anonymous referee for useful comments.

References

[1] Ajtai, M. (1983) Σ_1^1-formulae on finite structures, *Ann. Pure and Appl. Logic* **24** 1–48.

[2] Alon, N. (1986) Covering graphs by the minimum number of equivalence relations, *Combinatorica* **6** 201–206.

[3] Alon, N. (1986) Eigenvalues, geometric expanders, sorting in rounds, and Ramsey Theory, *Combinatorica* **6** 207–219.

[4] Alon, N. (1994) Explicit Ramsey graphs and orthonormal labelings, *Electronic J. Combinatorics* **1**:R12, 8pp.

[5] Andreev, A. E. (1986) On a family of boolean matrices, *Moscow Univ. Math. Bull.* **41** 79–82; translation from *Vestnik Mosk. Univ.* **41** (1986) 97–100

[6] Babai, L., Frankl, P. and Simon, J. (1986) Complexity classes in communication complexity. In *Proc. of 26th Ann. IEEE Symp. on Foundations of Comput. Sci.* 337–347.

[7] Beigel, R. and J. Tarui (1994) On ACC, *Computational Complexity* **4** 350–366.

[8] Bublitz, S. (1986) Decomposition of graphs and monotone size of homogeneous functions, *Acta Informatica* **23** 689–696.

[9] Chung, F. R. K., Erdős, P. and Spencer, J.(1983) On the decomposition of graphs into complete bipartite subgraphs. In *Studies in pure mathematics*, Mem. of P. Turán, 95–101.

[10] Duchet, P. (1979) *Représentations, noyaux en théorie des graphes et hypergraphes.* Thése de doctoral d'Etat, Université Paris VI.

[11] Erdős, P. and Rényi, A. (1962) On a problem in the theory of graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* **7**, 215–235.

[12] Erdős, P., Goodman, A.W. and Pósa, L. (1966) The representation of a graph by set intersections, *Can. J. Math.* **18** 106–112.

[13] Frankl, P. (1982) Covering graphs by equivalence relations, *Annals of Discrete Math.* **12** 125–127.

[14] Furst, M., Saxe, J. and Sipser, M. (1984) Parity, circuits and the polynomial time hierarchy, *Math. Systems Theory* **17** 13–27.

[15] Grigoriev, D. and Razborov, A. (2000) Exponential lower bounds for depth 3 arithmetic circuits in algebras of functions over finite fields, *Applicable Algebra in Engineering, Communication and Computing* **10**:6 465–487.

[16] Grolmusz, V. (1998) A lower bound for depth-3 circuits with MOD m gates, *Information Processing Letters*, **67** 87–90.

[17] Hajnal, A., Maass, W., Pudlaák, P., Szegedy, M. and Turán, G. (1993) Threshold circuits of bounded depth, *J. of Comp. Syst. Sci.* **46** 129–154.

[18] Hall, M. (1967) *Combinatorial Theory.* Wiley and Sons, New York and London.

[19] Håstad, J. (1989) Almost Optimal Lower Bounds for Small Depth Circuits. In S. Micali (ed.), *Advances in Computing Research* **5** 143–170.

[20] Håstad, J. (1998) The shrinkage exponent of de Morgan formulas is 2, *SIAM J. Comput.* **27**:1 48–64.

[21] Håstad, J., Jukna, S. and Pudlák, P. (1995) Top-down lower bounds for depth 3 circuits, *Computational Complexity* **5** 99–112.

[22] Karchmer, M. and Wigderson, A. (1988) Monotone circuits for connectivity require super-logarithmic depth. In *Proc. of 20th Ann. ACM Symp. on the Theory of Computing* 539–550.

[23] Karchmer, M. and Wigderson, A. (1993) On span programs. In *Proc. of the 8th Ann. Structures in Complexity Conference* 102–111.

[24] Khrapchenko, V. M. (1971) A method of determining lower bounds for the complexity of Π–schemes, *Math. Notes of the Acad. of Sci. of the USSR* **10**:1 474–479.

[25] Kneser, M. (1955) Aufgabe 300, *Jahresber. Deutsch. Math.-Verein* **58**.

[26] Kövári, P., Sós, V. T. and Túran, P. (1954) On a problem of Zarankiewicz, *Colloq. Math* **3** 50–57.

[27] Kollár, J., Rónyai, L. and Szabó, T. (1996) Norm-graphs and bipartite Turán numbers, *Combinatorica* **16**:3 399–406.

[28] Lokam, S. V. (2003) Graph complexity and slice functions, *Theory of Computing Systems* **36**:1 71–88.

[29] Lovász, L. (1978) Kneser's conjecture, chromatic numbers and homotopy, *J. Comb. Th. (A)* **25** 319–324.

[30] Lovász, L. (1977) Flats in matroids and geometric graphs. In P. J. Cameron (ed.) *Combinatorial surveys*, *Proc. of 6th British Combin. Conf.* 45–86. Academic Press, London.

[31] Paturi, R., Pudlák, P. and Zane, F. (1997) Satisfiability coding lemma. In *Proc. of 39th Ann. IEEE Symp. on Foundations of Comput. Sci.* 566–574.

[32] Paturi, R., Saks, M. and Zane, F. (2001) Exponential lower bounds for depth three boolean circuits, *Computational Complexity* **9**:1 1–15.

[33] Pudlák, P. and Rödl, V. (1992) A combinatorial approach to complexity, *Combinatorica* **12**:2 221–226.

[34] Pudlák, P. and Rödl, V. (1994) Some combinatorial-algebraic problems from complexity theory, *Discrete Mathematics* **136** 253–279.

[35] Pudlák, P. and Rödl, V. (2004) Pseudorandom sets and explicit constructions of Ramsey graphs. In J. Krajiček (ed.) *Quaderni di Matematica*, Vol. 13 327–346.

[36] Pudlák, P., Rödl, V. and Savický, P. (1988) Graph complexity, *Acta Informatica* **25** 515–535.

[37] Razborov, A. A. (1987) Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$, *Math. Notes of the Academy of Sciences of the USSR* **41**:4 333–338.

[38] Razborov, A. A. (1990) Applications of matrix methods to the theory of lower bounds in computational complexity, *Combinatorica,* **10**:1 81–93.

[39] Razborov, A. A. (1988) Bounded-depth formulae over the basis $\{\&, \oplus\}$ and some combinatorial problem. In S.I. Adian (ed.), *Problems of Cybernetics, Complexity Theory and Applied Mathematical Logic* (VINITI, Moscow) 149–166. (Russian)

[40] Smolensky, R. (1987) Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proc. of 19th Ann. ACM Symp. on the Theory of Computing* 77–82.

[41] Valiant, L. (1977) Graph-theoretic methods in low-level complexity. In *Proc. of 6th Conf. on Mathematical Foundations of Computer Science*, *Springer Lect. Notes in Comput. Sci.* **53** 162–176. Springer-Verlag.

[42] Yao, A. C. (1985) Separating the polynomial time hierarchy by oracles. In *Proc. of 26th Ann. IEEE Symp. on Foundations of Comput. Sci.* 1–10.

[43] Yao, A. C. (1990) On ACC and threshold circuits. In *Proc. of 31th Ann. IEEE Symp. on Foundations of Comput. Sci.* 619–627.