Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 5 (2004)

On Graph Complexity

S. Jukna [*†‡]

Abstract

A circuit complexity of a graph is the minimum number of union and intersection operations needed to obtain the whole set of its edges starting from stars. Our main motivation to study this measure of graphs is that it is related to the circuit complexity of boolean functions.

We prove some lower bounds to the circuit complexity of explicitly given graphs. In particular, we use the graph theoretic frame to prove that some explicit subsets of $GF(2)^n$ cannot be covered by fewer than $2^{\Omega(n)}$ affine subspaces of $GF(2)^n$.

We conclude with several graph-theoretic problems whose solution would have intriguing consequences in computational complexity.

Keywords: Graph complexity, depth-3 circuits, C_4-free graphs, clique covering number

AMS subject classification: 05C62, 05C35, 05C99, 68Q17, 68Q05

1 Introduction

A major challenge in computational complexity is to exhibit an *explicit* boolean function $f : \{0,1\}^m \to \{0,1\}$ which has high computational complexity, i.e. cannot be computed using a small number of basic boolean operations 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 problem of proving super-linear lower bounds is widely open even if we assume the additional restriction of circuit depth be[1] $O(\log m)$. The main difficulty here is that we want the function f be *explicitely constructed*—easy counting shows that almost all functions require circuits of size $2^{\Omega(m)}$.

*Research supported by a DFG grant SCHN 503/2-2.

†Universität Frankfurt, Institut für Informatik, Robert-Mayer-Str. 11-15, D-60054 Frankfurt, Germany (current address)

‡Institute of Mathematics and Informatics, Akademijos 4, LT-08663 Vilnius, Lithuania

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

ISSN 1433-8092

Pudlák, Rödl and Savický [32] observed that in oder to construct boolean functions requiring large circuits it would be enough to construct graphs that cannot be computed with a small number of union and intersection operations starting from some "simplest" graphs, like stars or cliques. In this paper we follow this approach in the case when input graphs are stars.

1.1 Graphs and circuits

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* a subset of vertices $S \subseteq V$ if $f(a_S) = 1$ where $a_S \in \{0,1\}^V$ is a binary vector with 1's in positions u for all $u \in S$, and 0's elsewhere.

Definition 1.1. *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, the *quadratic function* $f_G(X) = \bigvee_{uv \in E} x_u x_v$ of a graph $G = (V, E)$ represents G and can be computed by a trivial monotone depth-2 formula containing $|E| + 1$ gate: $|E|$ fanin-2 AND gates and one OR gate of fanin $|E|$. Hence, no n-vertex graph G of 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 bipartite covering number of graphs implies (see Proposition 5.2 below) that every n-vertex graph G of degree d can be represented by a monotone CNF (conjunctive normal form) $f(X) = \bigwedge_{i=1}^{r} \bigvee_{v \in S_i} x_v$ with $r = O(d^2 \log n)$. In particular, every graph of constant degree can be represented by a monotone depth-2 formulas of logarithmic size.

In this paper we are interested in graphs which cannot be represented by small circuits. Although proving lower bounds on the circuit complexity of graphs may be of independent interest, we (just like the authors of [32]) 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 $G \subseteq U \times W$ with $n = 2^m$ and $U = W = \{0,1\}^m$ one may associate a boolean function f in $2m$ variables—the *characteristic function* of G—such that $f_m(uv) = 1$ if and only if $uv \in G$.

Magnification Lemma: Given a circuit computing the characteristic function f of a bipartite graph G, it is possible to replace its input literals by ORs of variables

so that the obtained monotone circuit (no negated inputs) represents the graph G. The same holds when Parity gates are used instead of OR gates.

This fact may be 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 (or Parity) gates on the bottom (next to the inputs) level, then the obtained circuit represents G and has 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 $n^\varepsilon = 2^{\varepsilon m}$ gates, which is *exponential* in the number $2m$ of variables of f_m (this is where the term "magnification" comes from). That is, even moderate lower bounds on the computational complexity of graphs would yield high lower bounds on the computational complexity of boolean functions.

Note, however, that proving lower bounds for graphs may be 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 ([17]) whereas the corresponding to this function graph is just a union $(A \times \overline{B}) \cup (\overline{A} \times B)$ of two vertex-disjoint bipartite complete graphs, and can be represented by a circuit $(\bigvee_{u \in A \cup B} x_u) \wedge \left(\bigvee_{v \in \overline{A} \cup \overline{B}} x_v\right)$ using just three 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.

On the other hand, non-trivial lower bounds on graph complexity even in the class of depth-3 circuits would 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*. The total number of gates is the *size* of a circuit. If all gates have fanout 1, then the circuit is a *formula*.

Our motivation to study representation of graphs by depth-3 circuits comes from the following result due to Valiant [38]: 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 formula 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:

Valiant's Lemma: If a bipartite $n \times n$ graph cannot be represented by a monotone Σ_3 formula of size n^ε with $\varepsilon = \omega(1/\log\log\log n)$, 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, 12, 39, 17, 34, 37, 13, 19, 28]. However, for

Σ_3 circuits these lower 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 [29] 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 [29], their argument fails already when bottom fanin is 4.

1.2 Results

In this paper we are trying 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. Among others we prove the following.

1. If a bipartite $n \times n$ graph G has M edges and contains no copies of $K_{a,b}$, then any Σ_3 circuit with Parity gates on the bottom needs at least $M/(a+b)n$ gates to represent G (Theorem 3.2). This immediately yields strongly exponential lower bounds for many explicit boolean functions. In particular, this implies that any such circuit, detecting whether two given subsets of $\{1,\ldots,m\}$ are disjoint, needs $2^{\Omega(m)}$ gates (Corollary 3.5).

2. Any Σ_3 circuit which has Parity gates on the bottom and an arbitrary threshold gate on the top needs at least $\Omega(\sqrt{n})$ gates to represent an $n \times n$ Hadamard graph (Corollary 3.9). Again, this yields strongly exponential lower bounds for explicit boolean functions—including the Inner Product function—in this more general model.

3. There are explicit $n \times n$ graphs G such that both G and $\overline{G}$ can be represented by monotone Σ_3 circuits of size $O(\log n)$ but every monotone Π_3 circuit for G must have size $\Omega(\sqrt{n})$; a Π_3 circuit is a dual version of Σ_3 circuits. Nothing similar holds in the context of boolean functions: the size of minimal a Π_3 circuit computing a function f does not exceed the size of any Σ_3 circuit computing $\neg f$.

4. If a graph $G = (V, E)$ has no triangles and no 4-cycles, then any monotone formula (of arbitrary depth) computing the quadratic function $f_G(X) = \bigvee_{uv \in E} x_u x_v$ has length $\Theta(|E|)$ (Theorem 7.3).

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 [7, 36, 32, 35, 30, 26], and its potential seems to be far from being exhausted. As we will see, the circuit complexity of a graph is just a generalisation of its (edge) clique covering number, a well-known and widely studied measure of graphs. We show that non-trivial lower bounds on this generalised measure would have interesting consequences in computational complexity.

Notation

We shall use standard graph theory notation. In particular, $K_{a,b}$ is a *biclique* (bipartite clique, complete bipartite graph) $A \times B$, $A \cap B = \emptyset$ with parts of size $a = |A|$ and $b = |B|$. We shall look at bipartite graphs G with a fixed bipartition $V = U \cup W$ as sets $G \subseteq U \times W$ of their edges. By an n to n matching we will mean a bipartite graph M_n consisting of n vertex-disjoint edges. The (bipartite) complement of a bipartite graph $G \subseteq U \times W$ is a bipartite graph $\overline{G} = (U \times W) \setminus G$. A *non-edge* is a pair uv of non-adjacent vertices. A non-edge in a bipartite graph with a bipartition $V = U \cup W$ 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. The *degree* of a graph is the maximum degree of its vertices.

We will also use some well known graph covering measures of graphs. A *fat matching* is a union of vertex-disjoint bipartite cliques. A *fat covering* of a graph G is a family of fat matchings such that each of these fat matchings is a (spanning) subgraph of G and every edge of G is an edge of at least one member of the family. Similarly, a *bipartite clique covering* of G is a family $A_1 \times B_1, \ldots, A_t \times B_t$ of complete bipartite subgraphs of G such that every edge of G is an edge of at least one member of the family. The number t of subgraphs in such a covering is the size and the total number $\sum_{i=1}^{t}(|A_i| + |B_i|)$ of vertices is the weight of the covering.

Let $\mathrm{cc}(G)$ denote the minimum size and $\mathrm{cc}_{\mathrm{w}}(G)$ the minimum weight of a bipartite covering of G. These measures were first studied in [10], and now are the subject of extensive literature. In particular, it is known that $\mathrm{cc}_{\mathrm{w}}(G) = O(n^2/\log n)$ for every n-vertex graph, and there exist graphs matching this upper bound [8]. 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, [9, 11, 2, 31] among others (here fat matchings are called "equivalence graphs"). In particular, it is known that $\mathrm{fat}(G) = O(n/\log n)$ for every n-vertex graph [31]. In this paper we will use these measures to describe the circuit complexity of graphs.

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 2.1. 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 [27] in his famous proof of Kneser's conjecture [23] 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:

$$f(X) = \bigwedge_{i=1}^{r} \bigvee_{v \in S_i} x_v \tag{1}$$

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 $\exists i \in u \cap v$ 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 2.2. 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

$$f(X) = \bigoplus_{i=1}^{r} \bigvee_{v \in S_i} x_v \tag{2}$$

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. Again, the representation is quite compact: the circuit has only $r+1 = \log(2n)$ 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 [32], 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) to 0, we obtain that this Ramsey graph can be represented by a depth-2 circuit of size $O(\log n)$. Razborov in [36] has shown that also some other "combinatorially complicated" graphs can be represented by small circuits of constant depth.

On the other hand, some "combinatorially simple" graphs—like an n to n matching M_n 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 (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) 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) representing M_n.

These examples show that proving high lower bounds for *depth*-2 circuits is a relatively easy task. However, the case of *depth*-3 circuits turns out to be much more difficult. And this is not surprising because, as we already mentioned above, high lower bounds on the size of such circuits (of the form n^{ε} where ε may even tend slowly to 0) would resolve some old problems in computational complexity of boolean functions, including the problem of proving a super-linear lower bound for log-depth circuits.

3 Depth-3 circuits with Parity gates on the bottom

By Valiant's Lemma, high lower bounds on the size of monotone Σ_3 circuits representing an explicit bipartite graphs would give us super-linear lower bounds for logarithmic depth circuits. Pudlák, Rödl and Savický asked in [32] whether C_4-free graphs are hard for such circuits.

It turns out that this question can be answered affirmatively if the OR gates on the bottom are replaced by Parity gates; 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 $GF(2)$:

$$g(x_1,\ldots,x_m) = \prod_{i=1}^{r} \bigoplus_{j\in I_i} x_i \oplus \lambda_i$$

where $\lambda_i \in \{0,1\}$ and $I_i \subseteq \{1,\ldots,m\}$. If all scalars λ_i are equal 0, the circuit is *positive*.

Remark 3.1. 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 [34, 37, 13]. However, these techniques seem incapable of proving such high 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.

Theorem 3.2. *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 computing the characteristic function of G has top fanin at least*

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

Theorem 3.2 follows directly from the Magnification Lemma and the following two lemmas.

Lemma 3.3. *For every bipartite graph $G \subseteq U \times W$, $\mathrm{fat}(G)$ is 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 $\bigoplus_{u \in A \cup \overline{B}} x_u$ represents the same graph as $1 \oplus \bigoplus_{u \in A \cup B} x_u$.

Let $g = \bigoplus_{v \in A \cup B} x_v$ with $A \subseteq U$ and $B \subseteq W$ 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 $\mathrm{fat}(G)$, assume that G is a union of s fat matchings. Observe that every fat matching $M = A_1 \times B_1 \cup \cdots \cup A_r \times B_r$ can be obtained as an intersection of r fat matchings $A_i \times B_i \cup (A \setminus A_i) \times (B \setminus B_i)$ in $A \times B$ where A (B) is the union of all A_j's (B_j's). Since each fat matching $A \times B \cup \overline{A} \times \overline{B}$ can be represented by the Parity gate $g(X) = \bigoplus_{v \in A \cup \overline{B}} x_v$, we are done. □

Lemma 3.4. *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

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

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 3.2 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 $m = 2 \log n$ 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 3.5. *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 $K(r) \subseteq U \times V$ where U and W consist of all $n = 2^r$ subsets of $[r] = \{1, \ldots, r\}$, and $uv \in K$ iff $u \cap v = \emptyset$. The graph $K(r)$ can contain a complete bipartite $a \times b$ subgraph $\emptyset \neq A \times B \subseteq K$ only if $a \leq 2^k$ and $b \leq 2^{r-k}$ for some $0 \leq k \leq r$, because $(\bigcup_{u \in A} x_u) \cap (\bigcup_{v \in B} x_v) = \emptyset$. In particular, $K(r)$ does not contain a copy of $K_{a,a}$ with $a > 2^{r/2} = \sqrt{n}$. Since this graph has

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

edges, Theorem 3.2 yields that any $\Sigma_3^{\oplus}$ circuit representing $K(r)$ must have top fanin at least $|K(r)|/(2an) = \Omega(n^{0.08})$. It remains to apply the Magnification Lemma. $\square$

Remark 3.6. 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 [17]) 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 $(\bigvee_{u \in S} x_u) \wedge \left(\bigvee_{u \notin S} x_u\right)$ of two OR gates; the obtained monotone Σ_3 circuit will represent the same graph. Moreover, the graph $K(r)$ shows that $\Sigma_3^{\oplus}$ circuits may be even *exponentially weaker*: this graph can be represented by a monotone CNF of length $O(\log n)$ (see (1)) but requires $\Sigma_3^{\oplus}$ circuits of top fanin at least $\Omega(n^{\varepsilon})$.

A prominent example of a dense bipartite graph without $K_{2,2}$ 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 k-regular with $k = q + 1$ and contains no copies of $K_{2,2}$. Hence, the graph has $\Omega(n^{3/2})$ edges. This is almost optimal because, as shown in [24], no bipartite $n \times n$ graph without a copy of $K_{2,2}$ can have more than $(1 + o(1))n^{3/2}$ edges. By Singer's theorem ([16], p. 128) there exist $0 \leq a_1 < a_2 < \ldots < a_{q+1} < n$ such that P_n is isomorphic to the bipartite graph will parts $U = W = \{0, 1, \ldots, n-1\}$ in which $u \in U$ is joined to $v \in W$ iff $u = (v + a_i) \bmod n$ for some $1 \leq i \leq q + 1$.

For the characteristic function π_{2m} of this graph, Theorem 3.2 yields

Corollary 3.7. *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 [25]; explicit graphs without $K_{s,r}$ with larger values of r and s were earlier constructed by Andreev [4]. For the characteristic functions f_{2m}^a of these graphs, Theorem 3.2 yields

Corollary 3.8. *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 [14] for the Inner Product function

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

Quite recently Pudlák and Rödl [33] have also proved such a lower bound for the characteristic functions of particular pseudorandom sets. Both proofs employ non-trivial facts—the probabilistic communication complexity of IP in [14] and some properties of pseudorandom sets in [33].

Actually, the lower bounds in [14] and [33] 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, let us 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 [5].

Corollary 3.9. *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 3.9 and the Magnification Lemma immediately yield a lower bound $\Omega(2^{m/2})$ for IP_{2m} in this class of circuits.

For the proof of Corollary 3.9 we need the so-called "discriminator lemma" for threshold gates. Let $\mathcal{F}$ be a family of subsets of a finite set X. For a subset $A\subseteq X$, let $\mathrm{thr}_{\mathcal{F}}(A)$ denote the minimum number t for which there exist t members $B_1,\ldots,B_t$ of $\mathcal{F}$ and a number $0\le k\le t$ 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

$$\left|\frac{|A\cap B|}{|A|}-\frac{|\overline{A}\cap B|}{|\overline{A}|}\right|\ge\varepsilon.$$

Lemma 3.10. ([15]) *If $\mathrm{thr}_{\mathcal{F}}(A)\le t$ then A is a $1/t$-discriminator for some $B\in\mathcal{F}$.*

Proof. Let $B_1,\ldots,B_t\in\mathcal{F}$ 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,

$$1\ \le\ \sum_{i=1}^{t}\frac{|A\cap B_i|}{|A|}-\sum_{i=1}^{t}\frac{|\overline{A}\cap B_i|}{|\overline{A}|}$$

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

□

Proof of Corollary 3.9. Let A be an $n \times n$ Hadamard graph. Lindsey's lemma (see, e.g. [3] or [5]) 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 $\sqrt{abn}$. This, in particular, implies that both A and $\overline{A}$ have $\Theta(n^2)$ edges. Hence, by Lemmas 3.3 and 3.10, it is enough to show that $||A \cap B| - |\overline{A} \cap B|| = O(n^{3/2})$ 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 $\sum_{i=1}^{t} s_i \leq n$, we obtain

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

□

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 $2^{\Omega(m)}$.*

A trade-off $sr = \Omega\left(m^3/(\log m)^5\right)$ between these parameters for IP_{2m} was recently proved by Lokam [26] (also using the 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 $2^{\Omega(m^{1-\varepsilon})}$.

Theorem 4.1 follows directly from Lemma 3.4 and the following

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 $\mathrm{cc}(G) \leq s2^r$ and $\mathrm{cc}(\overline{G}) \leq r^s$.*

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 $\mathrm{cc}(G) \leq s2^r$.

To prove $\mathrm{cc}(\overline{G}) \leq r^s$, 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 that $\mathrm{cc}(\overline{G}) \leq \prod_{i=1}^{s} \mathrm{cc}(H_i) \leq r^s$. □

5 Combinatorics of Σ_3 circuits

In this section we give a combinatorial characterisation of graphs represented by monotone Σ_3 circuits. Recall that each such circuit is just an OR of CNFs (conjunctive normal forms), where a CNF of length t is an AND

$$g(X) = \Big(\bigvee_{u \in S_1} x_u \Big) \wedge \cdots \wedge \Big(\bigvee_{u \in S_t} x_u \Big) \tag{3}$$

of t clauses, each of which is an OR of variables. Let $\mathrm{cnf}(G)$ denote the minimum length of a CNF representing G, and let $\Sigma_3(G)$ be the minimum number t such that G can be represented as a union of at most t graphs H such that $\mathrm{cnf}(H) \leq t$.

The length of CNFs can be described combinatorially in terms of the clique covering number as well as in terms of set-intersections. Say that a bipartite graph $G \subseteq U \times W$ admits an *intersection representation of size* t if it is possible to associate with every vertex $u \in U \cup W$ a subset A_u of $\{1, \ldots, t\}$ so that for every arc $uv \in U \times W$, $uv \in G$ iff $A_u \cap A_v = \emptyset$. Let $\mathrm{int}(G)$ denote the smallest t for which G admits such a representation.

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

Proof. An OR of variables represents a complement of a biclique (and each complement of a biclique can be represented by an OR gate). Hence, a bipartite graph G can be represented by a CNF of the form (3) iff G is an intersection of complements of t bicliques, or equivalently, iff the complement $\overline{G}$ can be represented as a union of t bicliques, implying that $\mathrm{cnf}(G) = \mathrm{cc}(\overline{G})$.

The equality $\mathrm{cc}(\overline{G}) = \mathrm{int}(G)$ is also easy to show. Given an intersection representation of $G \subseteq U \times W$ by subsets A_u of $\{1, \ldots, t\}$ for $u \in U \cup W$, 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 $A_u = \{i : u \in I_i\}$. □

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 a general upper bound

$$\Sigma_3(G) \leq \min\left\{\mathrm{cc}(G), \mathrm{cc}(\overline{G})\right\}. \tag{4}$$

By (4), $\Sigma_3(G) \leq \mathrm{cc}(G) \leq n$ is a trivial upper bound for every n-vertex graph G. For graphs of small degree we have a better upper bound.

Proposition 5.2. *For every n-vertex graph G of degree d, $\mathrm{cnf}(G) = O(d^2 \log n)$ and $\Sigma_3(G) = O(d^{2/3} \log n)$*

Proof. Alon [2] has proved (using a probabilistic argument) that $\mathrm{cc}(\overline{G}) = O(d^2 \log n)$. Together with Proposition 5.1 this yields the first claim. To get the second claim, simply break G into $d^{2/3}$ subgraphs of maximal degree $d^{1/3}$ each. □

Proposition 5.3. *If H is a fat matching then both $\mathrm{cnf}(H)$ and $\Sigma_3(\overline{H})$ are at most $2\log n$. Moreover, if M_n is an n to n matching then $\Sigma_3(M_n) = \Omega(\log n)$.*

Proof. Let us first look how an n to n matching M_n can be represented by a monotone CNF of length $O(\log n)$. 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, $\mathrm{cnf}(M_n) = \mathrm{int}(M_n) \leq t = 2\log n$. The same argument clearly works for any fat matching H. The upper bound $\Sigma_3(\overline{H}) \leq 2\log n$ follows from the upper bound (4).

To prove the lower bound $\Sigma_3(M_n) = \Omega(\log n)$, let $t = \Sigma_3(M_n)$. Then there exists a matching $H \subseteq M_n$ containing $|H| \geq n/t$ edges and admitting an intersection representation of size t. Since the sets of neighbours of any two vertices in H are distinct, all the sets A_u associated with vertices u on the left (resp. on the right) part of the bipartition must be distinct. This implies $2^t \geq |H| \geq n/t$, and hence, $t = \Omega(\log n)$. □

So far we do not know of any explicit n-vertex graphs G with $\Sigma_3(G)$ substantially larger than $\log n$. The best what we know is the lower bound of the form $\Sigma_3(H_n) \geq (\log n)^{3/2-o(1)}$ proved by Lokam in [26] for an Hadamard graph H_n.

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:

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

by the size of such a circuit we again mean $\max\{s, r\}$. Since ANDs of more than 3 variables do not contribute to the representation of graphs, such circuits are just a slight generalisation of CNFs. And indeed, the representation power of such circuits is much weaker than that of monotone Σ_3 circuits. To see this, take an n to n matching. We already know that both M_n and $\overline{M}_n$ can be represented by monotone Σ_3 circuits of size $O(\log n)$. Moreover, M_n has also a monotone Π_3 circuit of logarithmic size, because $\mathrm{cnf}(M_n) = O(\log n)$. On the other hand, we have

Theorem 6.1. *Every monotone Π_3 circuit representing $\overline{M}_n$ must have size at least $\Omega(\sqrt{n})$.*

Note that nothing similar holds in the case of boolean functions: every Σ_3 circuit for a boolean function f can be tranformed to a Π_3 circuit of the same size for its complement $\neg f$.

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 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 $G \setminus H$ of remaining edges is non-empty, since otherwise we would have $H = G$, meaning that $\overline{G} \cup E$ is just a complete graph. By what was said, the CNF $(\bigvee_{u \in S_1} x_u) \wedge \cdots \wedge (\bigvee_{u \in S_t} x_u)$ 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. Let now 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 cliques. At least one of these cliques, say $A \times B$, must contain at least $|M_n|/t = n/t$ edges of the matching M_n. But this means that

$$|H \cap (A \times B)| \geq (n/t)^2 - (n/t).$$

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 cliques $R_1, \ldots, R_t$, 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 clique $R \in \{R_1, \ldots, R_t\}$ 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

$$N_1 + N_0 = |R| \geq \frac{N}{t} \geq \frac{n^2}{4t} \geq 4n,$$

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 Quadratic functions of graphs

In oder to obtain high lower bounds on the *non-monotone* circuit complexity of boolean functions it would be enough, by the Magnification Lemma, to show that any monotone boolean function $f(X)$ representing a given graph G requires large *monotone* circuits. That is, it is enough to deal with monotone circuits, but the lower bound must hold for *all* monotone boolean functions $f(X)$ representing G.

A natural monotone boolean function representing a given graph $G = (V, E)$ is the *quadratic function* f_G defined by

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

It is therefore useful (as the first step) to understand for what graphs these functions require large monotone circuits.

7.1 Quadratic functions and Σ_3 circuits

A *complete star* in a graph with n vertices is a set of $n-1$ edges sharing one endpoint in common. If the graph is bipartite, then a complete star is a set of edges joining all vertices of one part with a fixed vertex of the other part. A graph is *star-free* if it contains no complete stars.

Theorem 7.1 ([20]). *Let $G = (V, E)$ be a star-free graph of degree d, and let F be a monotone Σ_3 circuit computing f_G. Then F has at least $\sqrt{|E|}/d$ gates. If F is a formula, then F has at least $|E|/d^2$ gates.*

Since every Σ_3 circuit is an OR of CNFs, the theorem is an easy consequence of the following lemma—we include its proof to demonstrate the kind of difficulties one faces when trying to obtain similar lower bounds for circuits *recognising* the graph G. Let $\mathrm{cnf}(f_G)$ denote the minimum length of (i.e., the number of clauses in) a monotone CNF computing f_G. A *complete star* in a graph with n vertices is a set of $n-1$ edges sharing one endpoint in common.

Lemma 7.2. *If H is a star-free graph with M edges and degree d, then $\mathrm{cnf}(f_H) \geq M/d^2$.*

Proof. Let F be monotone CNF of length $t = \mathrm{cnf}(f_H)$ computing f_H. Since H has no complete stars, this CNF must contain at least two clauses. Take any of these clauses $C = \bigvee_{u \in S} x_u$ and consider the shrinked CNF $F' = F \setminus \{C\}$. Since C must accept all edges of H, each of these edges must have at least one endpoint in S. But any one vertex in S can be an endpoint of at most d edges, implying that $|S| \geq M/d$.

Since F is a shortest CNF computing f_H, the shrinked CNF F' must make an error, i.e. it must (wrongly) accept some independent set of H. That is, there must be an independent set I such that every clause of F' contains a variable x_v with $v \in I$. Since F' has only $t-1$ clauses, we may assume that $|I| \leq t-1$. This error must be corrected by the clause C, implying that every vertex $u \in S$ must be adjacent (in H) with at least one vertex in I, for otherwise F would wrongly accept the independent set $I \cup \{u\}$ of H. Hence, at least one vertex $v \in I$ must have at least $|S|/|I| \geq M/td$ neighbours in H. Since the degree of v cannot exceed d, the desired lower bound $t \geq M/d^2$ follows. □

Note that the main reason why such a simple argument does not work for $\mathrm{cnf}(H)$ is that in this last case the vertices of S need not have neighbours in I unless $I = \{v\}$ (F' accepts a single vertex), in which case v must have large degree (at least $|S|$). The only interesting errors made by F' are non-edges of H. Let T be the set of endpoints of these non-edges. Observe that T must be an independent set in H: if

uv and $u'v'$ are two non-edges accepted by F', then uv' cannot be an edge, because then we would be forced to include at least one of its endpoints u and v' in S, and the original CNF F would wrongly accept one of these two non-edges. Hence, T is an independent set, and and it is enough that $S \cap T = \emptyset$ to correct all these errors. In particular, if $S = \overline{T}$ then the clause $C = \bigvee_{u \in S} x_u$ rejects all non-edges wrongly accepted by F', and accepts all edges.

7.2 Quadratic functions and boolean formulas

In this section we consider circuits of *arbitrary* depth with unbounded fanin AND and OR gates; inputs are literals (variables and their negations). The *depth* of a circuit is the length of a longest path from an input to the output gate. 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 $2n$ (computing the corresponding OR of variables) so that the resulting monotone formula recognises G. Thus,

$$L(f) \geq L_+(G)/(2n).$$

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 [32] 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 $L_+(G) = \Omega\left(n \log^3 n\right)$ 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 [18].

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 $|S| \leq 2$

$$F(S) = 0 \text{ if and only if } S \text{ is an independent set in } G. \tag{5}$$

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 (5) must hold for subsets $S \subseteq V$ of size $|S| \leq k$.

Note that the quadratic function $f_G(X) = \bigvee_{uv \in E} x_u x_v$ recognises *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 $k < n$? Using a rank-argument it can be shown that, for some graphs, this is *not* possible unless k is smaller than two times the degree of G.

Theorem 7.3. *Let $G = (V, E)$ be a triangle-free graph without 4-cycles and of 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 any monotone circuit computing f has depth at least $\log |E| - 1$, and any monotone formula computing f has length at least $|E|/2$. 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 edges and vertices $y \in V \cup E$, 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

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

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 circuit computing f has depth at least $\log |\mathcal{R}|$ (see [21]), and every monotone formula has length at least $|\mathcal{R}|$ (see [35]). Hence, it remains to prove that $|\mathcal{R}| \geq |G|/2$.

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

$$M_{x,y} = 1 \text{ if and only if } x \cap y \neq \emptyset \tag{6}$$

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. 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 7.3 yields

Corollary 7.4. $L_+(f_{P_n}) = \Theta(n^{3/2})$.

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 (6) 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.

7.3 Saturated graphs

If a circuit computes f_G then it also represents the graph G. This holds for all graphs and all circuits. In general, however, the converse may not hold because the circuit representing a graph needs not to reject independent sets with more than two vertices. Hence, in general, lower bounds for circuits computing f_G do not imply lower bounds for circuits representing G. Still, there are graphs for which this holds.

A natural way to force a circuit representing a graph G to compute f_G is to "kill of" all large independent subsets by including additional edges. This is a standard trick in boolean complexity to obtain so-called "slice functions."

For a bipartite graph $H \subseteq U \times W$, we define its *saturated extension* as a (non-bipartite) graph $G = (V, E)$ such that $V = U \cup W$, and $uv \in E$ iff either $uv \in H$ or both vertices u, v lie in U or in W. That is, the induced subgraphs of G on U as well as on W are complete graphs.

Lemma 7.5. *If G is the saturated extension of a star-free bipartite graph, then every monotone boolean function representing G coincides with f_G.*

Proof. Let $G = (V, E)$ be a saturated extension of a bipartite graph $H \subseteq U \times W$. Suppose that H has no complete stars, and let f be a monotone boolean function representing G. Take an arbitrary subset $S \subseteq V$ of vertices. If $f_G(S) = 1$ then S

contains both endpoints of some edge $uv \in E$. This edge must be accepted by f and, since f is monotone, $f(S) = 1$. If $f_G(S) = 0$ then S is an independent set of G. But the only independent sets in G are single vertices and non-edges of H. Hence, $f(S) = 0$ because f represents H and H contains no complete stars. □

Note, however, that proving lower bounds on the monotone complexity of quadratic functions f_G of saturated graphs is a difficult task. In particular, Theorem 7.1 and Theorem 7.3 both fail for such functions.

8 Open problems

In the context of this paper the most interesting problem remains to prove that some explicit bipartite $n \times n$ graph requires large monotone Σ_3 circuits. The difficult thing here is to prove this for an *explicit* graph—easy counting shows that $\Sigma_3(G) = \Omega(\sqrt{n})$ for almost all bipartite $n \times n$ graphs G. Explicit graphs with $\Sigma_3(G) = \Omega(\log n)$ are easy to find: we have already shown that such is, for example, the an n to n matching M_n. However, to obtain some important consequences in computational complexity, we need an explicit graph G with $\Sigma_3(G) \geq n^{\varepsilon}$. If proved with $\varepsilon = \omega(1/\sqrt{\log n})$, this would give the first explicit boolean function in m variables requiring (non-monotone) Σ_3 circuit of size $2^{\omega(\sqrt{m})}$. If proved with $\varepsilon = (\log\log n)^{\omega(1)}/\log n$, this would give an explicit boolean function outside the second level of the communication complexity hierarchy introduced in [5]. If proved with $\varepsilon = \omega(1/\log\log\log n)$, this would give the first super-linear lower bound for log-depth circuits, thus resolving a long-standing open question in computational complexity.

If a graph G has a small monotone Σ_3 circuit, then some its dense subgraph H must have a small CNF. We already know (see Proposition 5.1) that $\mathrm{cnf}(H)$ is the minimal number of independent sets covering all non-edges of H. Hence, in oder to show that H cannot be represented by a small CNF one could try to show that no independent set of H can cover too many non-edges. If the original graph is a good expander, then one could hope that also H will have good enough expanding properties. As the initial graph with good expanding properties one could take, for example, the incidence graph P_n of the projective plane $PG(2, q)$. As shown in [3], every set X of vertices on one side of P_n has at least $n - n^{3/2}/|X|$ neighbours on the other side. Hence, every independent set of P_n can cover at most $O(n^{3/2})$ non-edges, implying that $\mathrm{cnf}(P_n) = \Omega(\sqrt{n})$. However, removing edges may destroy the expansion property so that the remaining graph H *may* contain large independent sets. For example, the adversary could remove all edges of P_n lying in an $(n/c)\times(n/c)$ clique for a large enough constant $c > 0$. Since no $m \times m$ C_4-free graph can have more than $(1 + o(1))m^{3/2}$ edges [24], the resulting graph H will still have a constant fraction of edges of P_n but the corresponding independent set will already cover a constant fraction of all its non-edges. Hence, we need an argument allowing us to show that $\mathrm{cc}(\overline{H})$ is large even when some independent sets of H are large.

Problem 8.1. Is there a constant $\varepsilon > 0$ such that $\mathrm{cc}(\overline{H}) \geq D^{\varepsilon}$ for every subgraph H of P_n of average degree at least D?

We already know (Theorem 6.1) that the monotone Π_3 circuit size for an n to n matching M_n is exponentially smaller than that for $\overline{M_n}$.

Problem 8.2. Exhibit an explicit bipartite graph G for which $\Sigma_3(\overline{G})$ is exponentially larger than $\Sigma_3(G)$.

The following problem does not require a graph be explicitely given, and hence, may be apparently easier. Let $\mathcal{E}$ be the set of all bipartite $n \times n$ graphs G with $\Sigma_3(G) \leq \exp((\log\log n)^c)$ for some constant $c > 0$. Let co$-\mathcal{E}$ be the set of complements of graphs from $\mathcal{E}$.

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

This would separate the second level of the communication complexity hierarchy ([5]) and hence resolve a long-standing open question in communication complexity.

A somewhat more ambitious task is to prove non-trivial lower bounds on the minimum number of bits sent by the players on the worst case input in the following "edge/non-edge" game between two players, Alice and Bob: Alice gets an edge x of G, Bob gets a non-edge y of G, and their goal is to determine a vertex $v \in x \setminus y$; at the end of the game on input (x, y) this vertex v must be known to both players. The graph G itself is known to players long before the game starts; hence, they may agree upon what characteristics of the graph they will use to encode the information about their inputs. It is clear that at least $\log_2 n$ bits of communication are necessary (Bob must know the answer) and $2\log_2 n$ bits are enough (Alice can just send her entire edge).

As mentioned in Sect. 7.2, the number of communicated bits on the worst case input is at least the logarithm of the minimum number $R(G)$ of disjoint rectangles needed to cover the communication matrix M of this game (see [21]). Rows of this matrix 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 8.4. Exhibit an n-vertex graph G on which the edge/non-edge game needs at least $\log_2 n + a\log\log n$ bits of communication for some $a \geq 2$, or equivalently, for which $R(G) = \Omega(n\log^a n)$.

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 [22]. If proved with $a \geq 3$ this would improve the strongest currently known lower bound $\Omega(m^{3-o(1)})$ due to Håstad [18].

Finally, let us recall yet another problem for graphs which would solve an old problem in circuit complexity. An ACC circuit is a constant-depth circuit with unbounded fanin AND, OR and MOD_k gates for arbitrary positive integers k, where $\mathrm{MOD}_k(x_1,\ldots,x_m)=1$ iff $x_1+\ldots+x_m=0$ modulo k. Exponential lower bounds for such circuits are known only if one is allowed to use MOD_k gates for prime numbers k [36, 37]. However, the case of composite moduli k— even the case of circuits with gates AND, OR and MOD_6—remains widely open.

Let $\alpha(G)$ denote the minimum number t for which there exist t bipartite complete graphs $H_1,\ldots,H_t$ and a subset $L\subseteq\{0,1,\ldots,t\}$ such that for every arc uv of G, $uv\in E$ if and only if the number of the graphs H_i that contain uv is a member of L. Equivalently, $\alpha(G)$ is the minimal number t for which there exists a subset $L\subseteq\{0,1,\ldots,t\}$ and an assignment $u\mapsto A_u\subseteq\{0,1,\ldots,t\}$ such that for every arc uv of G, $uv\in E$ if and only if $|A_u\cap A_v|\in L$.

Problem 8.5. Exhibit a bipartite $n\times n$ graph G with $\alpha(G)=\exp\left((\log\log n)^{\omega(1)}\right)$.

Together with the Magnification Lemma and the reduction [40, 6] of ACC circuits to depth-2 circuits with a symmetric gate on the top, this would yield an exponential lower bound for ACC circuits computing the characteristic function of G.

More combinatorial problems related to circuit complexity of boolean functions can be found in a survey paper of Pudlák and Rödl [31].

Acknowledgements

I would like to thank Noga Alon and Alexander Razborov for interesting discussions.

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] 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

[5] Babai, L., Frankl, P. and Simon, J. (1986): Complexity classes in communication complexity, In: *Proc. of* 26-*th IEEE FOCS*, pp. 337–347.

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

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

[8] 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, pp. 95–101 (1983).

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

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

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

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

[13] 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.

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

[15] Hajnal, A., Maass, W., Pudlaák, P., Szegedy, M. and G. Turán (1993): Threshold circuits of bounded depth, *J. of Computer and System Science* **46**, 129–154.

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

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

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

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

[20] Jukna, S. (2005): Disproving the single level conjecture, *Electronic Colloquium on Computational Complexity*, Report Nr. 21, 1–17.

[21] Karchmer, M. and Wigderson, A. (1988): Monotone circuits for connectivity require superlogarithmic depth. In: *Proc. 20th ACM STOC,* 539–550.

[22] 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.

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

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

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

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

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

[28] Paturi, R., Pudlák, P., Zane, F. (1997): Satisfiability coding lemma. In: *Proc. of* 39*-th IEEE FOCS*, pp. 566–574.

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

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

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

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

[33] Pudlák, P. and Rödl, V. (2004): Pseudorandom sets and explicit constructions of Ramsey graphs. Manuscript.

[34] 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.

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

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

[37] Smolensky, R. (1987): Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In: *Proc. of* 19*-th Ann. ACM Symp. Theor. Comput.*, 77–82.

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

[39] Yao, A. C. (1985): Separating the polynomial time hierarchy by oracles. In: *Proc.* 26*-th Ann. IEEE Symp. Found. Comput. Sci.*, pp. 1–10.

[40] Yao, A. C. (1990): On ACC and threshold circuits. In In: *Proc.* 31*-st Ann. IEEE Symp. Found. Comput. Sci.*, pp. 619–627.

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 g is *isolating* if g 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

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

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 $\{x_u : u \in U\} \cup \{x_v : v \in V\}$, and is obtained from F by replacing input literals y_i^a and z_i^a by functions

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

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 $X_{u,v}$ be the vector in $\{0,1\}^{U \cup V}$ 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 $F^+(X_{u,v}) = 1$ 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 $X_{u,v}$ 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 $Y_i^a = g(\{x_u : u \in U, u(i) = a\})$ 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 $X_{u,v}$ assigns precisely one 1 to these variables, and this 1 is in the position x_u. Hence, $Y_i^a(X_{u,v}) = 1$ 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, $Y_i^a(X_{u,v}) = y_i^a(uv)$, and we are done. □

Note that Lemma 8.6 holds not only for bipartite graphs but also for arbitrary k-partite hypergraphs $G \subseteq V_1 \times V_2 \times \cdots \times V_k$ with $|V_1| = |V_2| = \ldots = |V_k| = n = 2^m$. The only difference is that then the characteristic function f of such a hypergraph has km instead of $2m$ variables.

If we use OR gates as isolating functions in Lemma 8.6 then the only we need to obtain a positive extension F^+ of F is to simltineausly compute $4m = 4 \log n$ boolean sums. Pudlák, Rödl and Savický [32] have shown that, for any $k \geq 1$, any r boolean sums built out of n variables can be computed using at most $kn + k2^{\lceil r/k \rceil + 1}$ fanin-2 AND and OR gates. Together with Lemma 8.6 this yields the following lower bound for fanin-2 circuits.

Corollary 8.7. *Let f be the characteristic function of a bipartite $n \times n$ graph G. If G cannot be represented by a monotone circuit of size L, then f cannot be computed by a non-monotone circuit using fewer than $L - 24n$ gates.*

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'