A Lower Bound Technique for Nondeterministic Graph-Driven Read-Once-Branching Programs and its Applications*

Beate Bollig** and Philipp Woelfel**

FB Informatik, LS2, Univ. Dortmund, 44221 Dortmund, Germany
Fax: +49 (0)231 755-2047
{bollig,woelfel}@Ls2.cs.uni-dortmund.de

Abstract. We present a new lower bound technique for a restricted branching program model, namely for nondeterministic graph-driven read-once branching programs (g.d.-BP1s). The technique is derived by drawing a connection between ω-nondeterministic g.d.-BP1s and ω-nondeterministic communication complexity (for the nondeterministic acceptance modes $\omega \in \{\vee, \wedge, \oplus\}$). We apply the technique in order to prove an exponential lower bound for integer multiplication for ω-nondeterministic well-structured g.d.-BP1s. (For $\omega = \oplus$ an exponential lower bound was already obtained in [5] by using a different technique.) Further, we use the lower bound technique to prove for an explicitly defined function which can be represented by polynomial size ω-nondeterministic BP1s that it has exponential complexity in the ω-nondeterministic well-structured g.d.-BP1 model for $\omega \in \{\vee, \oplus\}$. This answers an open question from Brosenne, Homeister, and Waack [7], whether the nondeterministic BP1 model is in fact more powerful than the well-structured graph-driven variant.

1 Introduction and Results

Branching programs (BPs) or equivalently Binary Decision Diagrams (BDDs) belong to the most important nonuniform models of computation. (For a history of results on branching programs see e.g. the monograph of Wegener [22].)

Definition 1.1. *A* branching program *on the variable set $\mathcal{X}_n = \{x_1, \ldots, x_n\}$ is a directed acyclic graph with one source and two sinks. The internal nodes are marked with variables in $\mathcal{X}_n$ and the sinks are labeled with the boolean constants 0 and 1. Further, each internal node has two outgoing edges, marked with 0 and 1, respectively.*

Let B_n denote the set of boolean functions $f_n : \{0,1\}^n \to \{0,1\}$. A branching program on $\mathcal{X}_n$ represents at each node v a function $f_v \in B_n$ in the following way. If v is a c-sink, $c \in \{0,1\}$, then $f_v = c$ and if v is an internal node with 0-successor v_0 and 1-successor v_1, then $f_v = \overline{x_i} f_{v_0} \vee x_i f_{v_1}$. The function represented by the branching program itself is the function represented at the source. The size *of a*

* An extended abstract of this paper has been presented at MFCS 2002.
** Supported in part by DFG grant We 1066

branching program G is the number of its nodes, denoted by $|G|$, and the branching program complexity *of a boolean function f is the size of the smallest branching program representing f.*

Nondeterminism is one of the most powerful concepts in complexity theory. In analogy to the definition of Turing machines, different modes of acceptance have been studied for branching programs. The following definition is due to Meinel [18].

Definition 1.2. *Let Ω be a set of binary operations. An Ω-*nondeterministic branching program *is a branching program of which some internal nodes are labeled with an operation $\omega \in \Omega$ instead of a variable. Such nodes are called* nondeterministic nodes*, and the function represented at the nondeterministic node v, labeled with ω and with 0-successor v_0 and 1-successor v_1, is $f_v = f_{v_0}\, \omega\, f_{v_1}$. As in the deterministic case, a nondeterministic branching program represents the function which is represented at the source. The size of an Ω-nondeterministic branching program is the number of its deterministic nodes.*

For the ease of notation, we write ω instead of $\{\omega\}$ if the considered set Ω of binary operations is a singleton. In this paper, we investigate the most common acceptance modes OR, AND, and PARITY, denoted by $\vee$, $\wedge$, and $\oplus$, respectively (although our lower bound technique is not limited to these acceptance modes). For certain acceptance modes ω, an alternative way to determine the function value of a function represented by an ω-nondeterministic branching program is to count the number of computation paths of an input a which lead to the 1-sink. (A source-to-sink path is a computation path of the input $a = (a_1 \dots a_n)$ if it leaves any deterministic node labeled by x_i over the edge labeled by a_i and any nondeterministic node over an arbitrary edge.) E.g. a $\oplus$-nondeterministic BP accepts an input a if and only if an odd number of computation paths of a lead to the 1-sink.

Deterministic and nondeterministic BPs can be simulated by the corresponding Turing machines, and the BP complexity of a boolean function is a measure for the space complexity of the corresponding model of sequential computation. Therefore, one is interested in large lower bounds for BPs. Until today, no superpolynomial lower bounds for general BPs representing an explicitly defined function are known. Therefore, various types of restricted BPs have been investigated, and one is interested in refining the proof techniques in order to obtain lower bounds for less restricted BPs. (For the latest breakthrough see e.g. [1], [2], and [3].) There are several reasonable possibilities to restrict BPs, among them restrictions concerning the multiplicity of variable tests or the order in which variables may be tested.

Definition 1.3. (i) *A (nondeterministic)* read-once branching program (*short:* BP1) *is a (nondeterministic) branching program where each variable appears on each computation path at most once.*

(ii) *A (nondeterministic) branching program is called* s-oblivious*, for a sequence of variables* $s = (s_1, \ldots, s_l)$*,* $s_i \in X_n$*, if the set of decision nodes can be partitioned into disjoint sets* V_i*,* $1 \leq i \leq l$*, such that all nodes from* V_i *are labeled with* s_i *and the edges which leave* V_i*-nodes reach a sink or a* V_j*-node where* $j > i$*.*

Besides the theoretical viewpoint people have used BPs in applications. Oblivious BP1s, introduced by Bryant [8] under the term OBDDs, have found a large variety of applications, e.g. in circuit verification. Obliviousness, though, is a very strong restriction. Gergov and Meinel [13] and Sieling and Wegener [20] have independently generalized the concept of obliviousness in the deterministic read-once case in order to show how to use BP1s for verification.

Definition 1.4. *A* graph order *is a branching program with a single sink, where on each path from the source to the sink all variables appear exactly once. A (nondeterministic)* graph-driven BP1 (*short:* g.d.-BP1) *is a (nondeterministic)* BP1 G *for which there exists a graph order* G_0 *with the following property: If for an input* a*, a variable* x_i *appears on the computation path of* a *in* G *before the variable* x_j*, then* x_i *also appears on the unique computation path of* a *in* G_0 *before* x_j*.*

A (nondeterministic) g.d.-BP1 G *with graph order* G_0 *is called* well-structured*, if there exists a mapping* α *from the node set of* G *to the node set of* G_0 *such that for every node* v *in* G *the node* $\alpha(v)$ *is labeled with the same variable as* v*, and such that if a computation path of an input* a *passes through* v*, then the computation path of* a *in* G_0 *passes through* $\alpha(v)$*.*

In graph-driven BP1s according to a fixed graph order, for each input the variables are tested in the same order, whereas (different from OBDDs) for different inputs different orders may be used. For nondeterministic BP1s graph orders do not exist in general. If we generalize OBDDs to nondeterministic OBDDs we gain the possibility of nondeterministic guesses, but the variable order remains the same on all computation paths. For read-once branching programs the situation is different. It is possible to guess nondeterministically and moreover, for each input arbitrary orders of the variables are allowed.

Sieling and Wegener [20] have observed that there is a time-space trade-off between graph-driven and well-structured graph-driven BP1s in the deterministic case. The stronger structural property of the latter model leads to the design of simpler and faster algorithms but the storage space of these algorithms is larger than the storage space of the algorithms for graph-driven BP1s. The difference between the two models is the following one. For graph-driven BP1s G according to a graph order G_0 it is possible that a node v with label x_i is reached on the computation paths for two inputs a and b in G whereas the nodes with label x_i on the computation paths for the inputs a and b in G_0 are different. This is not allowed in the well-structured case. Since we are interested in lower bounds, we may assume that each graph order does not contain identical subgraphs. Then well-structured graph-driven BP1s according to a fixed graph order G_0 can be obtained in the following way. We start by a complete decision tree which is ordered

according to G_0, afterwards we merge all identical subgraphs. Finally, all nodes which have the same 0-and 1-successor are deleted. Any OBDD is well-structured since there exists exactly one x_i-node in any variable order for each variable x_i. In [4] it has been shown that even restricted nondeterministic well-structured graph-driven BP1s, called nondeterministic tree-driven BP1s, are a proper generalization of nondeterministic OBDDs.

The concept of graph-driven branching programs has turned out to be also useful in other settings, see e.g. [16] and [21]. Gergov and Meinel [12] were the first ones who suggested parity graph-driven BP1s as a data structure for boolean functions. Another reason for investigating parity graph-driven BP1s is that until now exponential lower bounds on the size of parity read-once branching programs for explicitly defined boolean functions are unknown. One step towards the proof of such bounds might be to investigate BP models "inbetween" deterministic and parity BP1s. Nondeterministic and parity graph-driven BP1s have been investigated more intensely in [4], [7], and [5].

Since for nondeterministic BP1s graph order do not exist in general, it is an intriguing question, whether nondeterministic (well-structured) g.d.-BP1s are in fact significantly more restricted than general nondeterministic BP1s. One of the main contributions of this paper is that we answer this question for the well-structured case in an affirmative way for the most important nondeterministic acceptance modes. This is done by presenting a function called $n/2$-MRC$_n$, which can be represented in polynomial size by ω-nondeterministic BP1s but has exponential complexity in the ω-nondeterministic well-structured g.d.-BP1 model (for $\omega \in \{\vee, \oplus\}$). Note that an analogous separation result for $\omega = \wedge$ follows right away from de Morgan's rules for the complement of $n/2$-MRC$_n$.

In order to prove the separation result, we derive a new lower bound technique. Until now, there was only one general lower bound technique known for nondeterministic well-structured g.d.-BP1s, which in addition worked only for the parity-acceptance mode [7]. We follow a more general approach by drawing connections to communication complexity. Hence, our lower bound technique can be applied to all acceptance modes, where corresponding lower bounds for communication complexity can be proven.

As another application of our lower bound technique, we prove an exponential lower bound for integer multiplication. Lower bounds for integer multiplication are motivated by the general interest in the complexity of important arithmetic functions and the insight into the structure of such functions which is often gained by lower bound proofs. Furthermore, since exponential lower bounds are often proven for functions which are "designed" in such a way that they fit to a given lower bound technique, the lower bound proofs for important functions can lead to refinements of the proof techniques.

Definition 1.5. *The boolean function* $\mathrm{MUL}_{i,n} \in B_{2n}$ *maps two* n*-bit integers* $x = x_{n-1} \dots x_0$ *and* $y = y_{n-1} \dots y_0$ *to the* i*th bit of their product, i.e.,* $\mathrm{MUL}_{i,n}(x, y) = z_i$, *where* $x \cdot y = z_{2n-1} \dots z_0$.

Since the middle bit (the bit z_{n-1}) of integer multiplication is the hardest bit to compute, one is interested mainly in the complexity of $\mathrm{MUL}_n := \mathrm{MUL}_{n-1,n}$. Bryant [9] has proven an exponential lower bound of $2^{n/8}$ for the function MUL_n in the OBDD model, and Gergov has presented an exponential lower bound for nondeterministic linear-length oblivious branching programs [11]. Later Ponzio has shown that the complexity of this function is $2^{\Omega(\sqrt{n})}$ for BP1s [19], and Bollig [4] has proven an exponential lower bound for nondeterministic tree-driven BP1s (i.e. g.d.-BP1s where the graph order is a tree of polynomial size).

Recently, progress in the analysis of MUL_n has been achieved by a new approach using universal hashing. Woelfel [23] has improved Bryant's lower bound to $\Omega\big(2^{n/2}\big)$ and Bollig and Woelfel [6] have presented a lower bound of $\Omega\big(2^{n/4}\big)$ for BP1s. Finally, Bollig, Waack, and Woelfel [5] have proven a lower bound of $2^{(n-46)/12}/n$ for $\oplus$-nondeterministic well-structured g.d.-BP1s. Their proof, though, is limited to this type of acceptance mode.

One step towards proving exponential lower bounds for MUL_n for unrestricted nondeterministic BP1s might be to investigate BP models "inbetween" deterministic and nondeterministic BP1s. This was also the motivation behind a result in [24] where an exponential lower bound has been proven for nondeterministic BP1s which have only a restricted number of nondeterministic nodes at the top of the BP1.

The lower bound for integer multiplication presented here is $2^{n/12-4} \cdot n^{-1/3}$ and is valid for all ω-nondeterministic well-structured g.d.-BP1s where $\omega \in \{\vee, \wedge, \oplus\}$. Comparing with the algebraic approach of [5], one advantage is that using methods from communication complexity, all important types of nondeterminism can be handled simultaneously.

2 A Lower Bound Technique for Nondeterministic Graph-Driven BP1s

Methods from communication complexity have been used to prove lower bounds in several branching program models, e.g. for OBDDs. (See e.g. [14, 17] for the theory of communication complexity.) Consider a boolean function $f \in B_n$ which is defined on the variables in $\mathcal{X}_n = \{x_1, \ldots, x_n\}$, and let $\Pi = (\mathcal{X}_A, \mathcal{X}_B)$ be a partition of $\mathcal{X}_n$. Assume that Alice has access only to the input variables in $\mathcal{X}_A$ and Bob has access only to the input variables in $\mathcal{X}_B$. In a one-way communication protocol, upon a given input x, Alice is allowed to send a single message (depending on the input variables in $\mathcal{X}_A$) to Bob who must then be able to compute the answer $f(x)$. In an ω-nondeterministic communication protocol, $\omega \in \{\vee, \wedge, \oplus\}$, Alice is allowed to "guess" a message. The function value is one if the number of guesses upon which Bob accepts the input matches the corresponding acceptance mode ω (e.g. is at least one in the case of $\omega = \vee$ or odd in case of $\omega = \oplus$). The *ω-nondeterministic one-way communication complexity* of the function f is the

number of bits of communication which need to be transmitted by such a protocol that computes f. It is denoted by $\mathrm{ND}_{\omega}^{A\to B}(f,\Pi)$.

In order to state the lower bound technique for nondeterministic g.d.-BP1s, we have to introduce some further notation, first. A *filter* of a set X is a closed upward subset of 2^X (i.e. if $S \in \mathcal{F}$, then all supersets of S are in $\mathcal{F}$). Let $\mathcal{F}$ be a filter of $\mathcal{X}_n = \{x_1, \ldots, x_n\}$. A subset $B \subseteq \mathcal{X}_n$ is said to be in the *boundary* of $\mathcal{F}$ if $B \notin \mathcal{F}$ but $B \cup \{x_i\} \in \mathcal{F}$ for some $x_i \in \mathcal{X}_n$.

Let f be a function in B_n defined on the variables in $\mathcal{X}_n$ and $\mathcal{F}$ be a filter of $\mathcal{X}_n$. For a subset $Z \subseteq \mathcal{X}_n$, we denote by $\mathcal{A}(Z)$ the set of all possible assignments to the variables in Z. Let $\Pi = (\mathcal{X}_A, \mathcal{X}_B)$ be a partition of $\mathcal{X}_n$. If $\mathcal{X}_B$ is in the boundary of $\mathcal{F}$, then Π is called *$\mathcal{F}$-partition* of $\mathcal{X}_n$. Finally, a function $f' \in B_n$ is called *(ϵ,Π)-close* to f, if there exists a set $R \subseteq \mathcal{A}(\mathcal{X}_A)$ with $|R| \geq \epsilon \cdot 2^{|\mathcal{X}_A|}$, such that f and f' coincide on all inputs in $R \times \mathcal{A}(\mathcal{X}_B)$.

Theorem 2.1. *Let $\mathcal{F}$ be a filter of $\mathcal{X}_n$, $f \in B_n$, $0 < \epsilon \leq 1$, and $\ell \in \mathbb{N}$. If for every $\mathcal{F}$-partition Π of $\mathcal{X}_n$ and for every function f' which is (ϵ,Π)-close to f it is $\mathrm{ND}_{\omega}^{A\to B}(f',\Pi) > \ell$, then any ω-nondeterministic graph-driven* BP1 *representing f either has a size of at least $2^{\ell}+1$ or its graph order has a size of more than $1/\epsilon$ (for $\omega \in \{\vee, \wedge, \oplus\}$).*

Proof. Let G be an ω-nondeterministic g.d.-BP1 representing f and assume that $|G| \leq 2^{\ell}$ and that the graph order G_0 of G has a size of at most $1/\epsilon$. It suffices to show that there exist an $\mathcal{F}$-partition Π and a function f' which is (ϵ,Π)-close to f such that $\mathrm{ND}_{\omega}^{A\to B}(f',\Pi) \leq \ell$.

Let v^+ for each node v of G_0 be the set of variables which are assigned to a node reachable from v (including the variable which is assigned to v). The filter $\mathcal{F}$ defines a cut, called *frontier*, through the edges of the graph order in the following way. An edge $e = (v,w)$ is in the frontier and thus called *frontier-edge* if the set $v^+ \in \mathcal{F}$ but $w^+ \notin \mathcal{F}$. We know that $s^+ = \mathcal{X}_n$ for the source s, $t = \emptyset$ for the sink t and finally - since G_0 is a complete BP1 - $w^+ = v^+ \setminus \{x_i\}$ for each edge (v,w) where v is marked with a variable x_i. Hence, each source-to-sink path passes through exactly one frontier-edge, and if (v,w) is a frontier-edge then w^+ is in the boundary of $\mathcal{F}$.

Using simple graph-theoretical arguments it is easy to see that G_0 contains at most $|G_0|$ frontier-edges. Hence, by the pigeonhole principle there exists a frontier-edge $e = (v,w)$ such that the computation paths of at least $2^n/|G_0|$ inputs pass through e. Let $\mathcal{X}_B = w^+$, $\mathcal{X}_A = \mathcal{X}_n \setminus w^+$ and $\Pi = (\mathcal{X}_A, \mathcal{X}_B)$. Clearly, each input reaching e is uniquely determined by its assignment to the variables in $\mathcal{X}_A$, and therefore there exists a set $R \subseteq \mathcal{A}(\mathcal{X}_A)$ with $|R| \geq 2^{|\mathcal{X}_A|}/|G_0|$ such that all inputs in $R \times \mathcal{A}(\mathcal{X}_B)$ reach the edge e. Furthermore, since w^+ is in the boundary of $\mathcal{F}$, Π is an $\mathcal{F}$-partition of $\mathcal{X}_n$.

Consider now the following ω-nondeterministic one-way communication protocol with respect to Π. For the input $x = (x_A, x_B) \in \mathcal{A}(\mathcal{X}_A) \times \mathcal{A}(\mathcal{X}_B)$, Alice tests whether $x_A \in R$. If $x_A \notin R$, we do not care about the result of the computation;

thus Alice may send an arbitrary ℓ-bit string to Bob. If on the other hand $x_A \in R$, then the input x reaches the edge e in G_0 after testing exactly the variables in $\mathcal{X}_A$. This means that in G on all computation paths of x the $\mathcal{X}_A$-variable tests appear before the $\mathcal{X}_B$-variable tests, and the nodes which are reached on these computation paths after all $\mathcal{X}_A$-variable tests have been performed are uniquely defined by x_A. Hence, Alice may guess nondeterministically a path which corresponds to the values of the $\mathcal{X}_A$-variables, and determine the unique node u on this path where the first variable in $\mathcal{X}_B$ is being tested. She finally sends an ℓ-bit string describing the node to Bob (recall that G consists of at most 2^ℓ nodes). Then Bob knows the node u and can compute (nondeterministically) $f(x) = f_u(x_B)$.

Obviously, by the above described protocol, Alice and Bob compute nondeterministically the function value $f(x)$ if $(x_A, x_B) \in R \times \mathcal{A}(\mathcal{X}_B)$. Since Alice sends at most ℓ bits to Bob, there exists a function f' which coincides with f on all inputs in $R \times \mathcal{A}(\mathcal{X}_B)$ such that $\mathrm{ND}_\omega^{A \to B}(f', \Pi) \leq \ell$. Finally, f' is (ϵ, Π)-close to f because by construction $|R| \geq 2^{|\mathcal{X}_A|}/|G_0| \geq \epsilon\, 2^{|\mathcal{X}_A|}$. □

The above technique does not yield lower bounds for nondeterministic g.d.-BP1s directly, because the size of the graph order of such a branching program is not part of the nondeterministic g.d.-BP1 size. Until now it is unknown whether there exists a class of functions f_n which has polynomial complexity in the nondeterministic g.d.-BP1 model whereas the size of every graph order of a polynomial size nondeterministic g.d.-BP1 for f_n is exponential. The situation is different in the well-structured case as Bollig, Waack, and Woelfel [5] have shown by the following proposition.

Proposition 2.2 ([5]). *For any nondeterministic well-structured graph driven* BP1 *G on n variables, there exists a graph order G_0 such that G is G_0-driven and $|G_0| \leq 2n|G|$.*

Corollary 2.3. *Let $f \in B_n$ be a function satisfying the conditions of Theorem 2.1 for some filter $\mathcal{F}$ of $\mathcal{X}_n$ and the parameters ϵ and ℓ. Then any ω-nondeterministic well-structured graph driven* BP1 *for f has a size of more than* $\min\left\{2^\ell, (\epsilon \cdot 2n)^{-1}\right\}$.

3 An Exponential Lower Bound for Integer Multiplication

As a first application of the lower bound technique, we prove a lower bound for integer multiplication. We consider here the boolean function $\mathrm{MUL}_n^* \in B_{2n-2}$; this is the subfunction of MUL_n, which takes as inputs only odd integers (i.e. the least significant bits of the two n-bit factors are fixed to 1). Obviously, a lower bound on the (nondeterministic) communication complexity of MUL_n^* implies the same lower bound for MUL_n.

The following lemma describes the connection between integer multiplication and nondeterministic communication complexity, which we need to apply Corollary 2.3. It is well known that a large nondeterministic communication complexity

can be shown by proving that the communication matrix according to a given partition Π contains a large triangular submatrix (this follows e.g. from the methods in [10]). Note that we use the term *submatrix* here in the common combinatorial sense, which means that each submatrix is obtained from a matrix M by selecting an arbitrary set of rows and columns of M and ordering them arbitrarily.

Lemma 3.1. *Let $A, B \subseteq \mathbb{Z}_{2^n}$ and $Y \subseteq \mathbb{Z}^*_{2^n} := \{1, 3, \ldots, 2^n - 1\}$ and assume that $|B| = 2^\beta$ and $|Y| = 2^\mu$. Consider the $|A| \times |B \times Y|$-matrix M, where each row is identified with an integer $a \in A$ and each column is identified with a pair $(b, y) \in B \times Y$, and the entry of the matrix in row a and column (b, y) equals $\mathrm{MUL}^*_n(a + b, y)$. Then M contains a triangular $s \times s$-submatrix where $s = \min\left\{|A|/2 - 1,\ 2^{(3\mu+\beta-3n-10)/4} - 1\right\}$.*

In order to prove Lemma 3.1, we need to recall some properties about integer multiplication which have been derived by Bollig and Woelfel [6] and Bollig, Woelfel, and Waack [5] using universal hashing. Let $\mathbb{Z}^*_{2^n}$ be the set of odd n-bit integers.

Lemma 3.2 ([5, 6]). *Let $X \subseteq \mathbb{Z}_{2^n}$ and $Y \subseteq \mathbb{Z}^*_{2^n}$. If $|X| \cdot |Y| \geq 2^{n+2r+1}$, $r \geq 0$, then there exists an element $y \in Y$ such that*

$$\forall q \in \{0, \ldots, 2^r - 1\}\, \exists x \in X :\ q \cdot 2^{n-r} \leq (xy) \bmod 2^n < (q+1) \cdot 2^{n-r}.$$

Lemma 3.3 ([5]). *Let $Y \subseteq \mathbb{Z}^*_{2^n}$, $1 \leq r \leq n-1$, and $(z_i, z_i') \in \mathbb{Z}_{2^n} \times \mathbb{Z}_{2^n}$, where $z_i \neq z_i'$ for $1 \leq i \leq t$. Then there exists a subset $Y' \subseteq Y$, $|Y'| \geq |Y| - t \cdot 2^{n-r+1}$, such that for all pairs (z_i, z_i'), $1 \leq i \leq t$,*

$$\forall y \in Y' :\ 2 \cdot 2^{n-r} \leq \big((z_i - z_i')y\big) \bmod 2^n \leq 2^n - 2 \cdot 2^{n-r}.$$

Proof (of Lemma 3.1). We show below that there exist an element $y \in Y$, a subset $\{a_1, \ldots, a_{s+1}\} \subseteq A$, and a subset $\{b_1, \ldots, b_s\} \subseteq B$ such that for all $1 \leq j \leq s+1$ and $1 \leq i \leq s$

$$\mathrm{MUL}^*_n(a_j + b_i, y) \ = \ \begin{cases} 0 & \text{if } i \geq j \\ 1 & \text{if } i < j. \end{cases} \tag{1}$$

This means that the $s \times s$-submatrix of M consisting of the rows $a_2, \ldots, a_{s+1}$ and of the columns $(b_1, y), \ldots, (b_s, y)$ is triangular.

Let $r = (\mu+\beta-n)/2-1$. If $|A| \leq 2^{(3\mu+\beta-3n-6)/4}$, then we let $A' = A$. Otherwise, we let A' be an arbitrary subset of A containing exactly $2^{(3\mu+\beta-3n-6)/4}$ elements.

Consider now the $t = |A'|(|A'|-1)$ pairs (z_i, z_i'), $1 \leq i \leq t$, with $z_i, z_i' \in A'$ and $z_i \neq z_i'$. Applying Lemma 3.3, we obtain a subset $Y' \subseteq Y$, $|Y'| \geq |Y| - |A'|^2 \cdot 2^{n-r+1}$, such that for all different $a, a' \in A'$ it holds

$$\forall y \in Y' :\ 2 \cdot 2^{n-r} \leq \big((a - a')y\big) \bmod 2^n \leq 2^n - 2 \cdot 2^{n-r}. \tag{2}$$

Then

$$\begin{aligned} |B| \cdot |Y'| \ &\geq\ |B| \cdot |Y| - |B| \cdot |A'|^2 \cdot 2^{n-r+1} \ \geq\ 2^{\beta+\mu} - 2^{\beta+(3\mu+\beta-3n-6)/2+n-r+1} \\ &=\ 2^{\beta+\mu} - 2^{\beta+\mu+(\mu+\beta-n)/2-1-r-1} \ =\ 2^{\beta+\mu} - 2^{\beta+\mu-1} \ =\ 2^{\beta+\mu-1} \\ &=\ 2^{n+2r+1}. \end{aligned}$$

Therefore, we may apply Lemma 3.2 (with $X = B$) in order to see that there exists an element $y \in Y'$ such that

$$\forall q \in \{0, \ldots, 2^r - 1\}\, \exists b \in B :\ q \cdot 2^{n-r} \leq (by) \bmod 2^n < (q+1) \cdot 2^{n-r}. \quad (3)$$

We let this element $y \in Y'$ be fixed from now on. Further, let

$$A'_< = \left\{a \in A' \;\middle|\; (ay) \bmod 2^n < 2^{n-1}\right\}$$

and

$$A'_\geq = \left\{a \in A \;\middle|\; (ay) \bmod 2^n \geq 2^{n-1}\right\}.$$

We choose A^* to be the set which has at least as many elements as the other one. Hence,

$$|A^*| \ \geq\ |A'|/2 \ \geq\ \min\left\{|A|,\, 2^{(3\mu+\beta-3n-6)/4}\right\}/2 \ =\ s+1.$$

We consider only the case where A^* equals $|A'_<|$; the other case is symmetric and can be proven analogously. We label the elements in A^* by $a_1, \ldots, a_{s+1}$ in such a way that

$$0 \ \leq\ (a_1 y) \bmod 2^n \ \leq\ \ldots \ \leq\ (a_{s+1} y) \bmod 2^n \ < 2^{n-1}. \quad (4)$$

Then we obtain by (2) that

$$\forall 1 \leq i \leq s :\ (a_i y) \bmod 2^n + 2 \cdot 2^{n-r} \ \leq\ (a_{i+1} y) \bmod 2^n. \quad (5)$$

For $1 \leq i \leq s$ we let now

$$q_i \ :=\ \left\lfloor \frac{2^{n-1} - (a_i y) \bmod 2^n}{2^{n-r}} \right\rfloor - 1 \quad (6)$$

and choose $b_i \in B$ such that

$$q_i \cdot 2^{n-r} \ \leq\ (b_i y) \bmod 2^n \ <\ (q_i + 1) \cdot 2^{n-r}. \quad (7)$$

(Such a b_i exists because of (3)). Hence, we get for $1 \leq j \leq i$

$$(a_j y) \bmod 2^n + (b_i y) \bmod 2^n \ \overset{(4),(7)}{<}\ (a_i y) \bmod 2^n + (q_i + 1) \cdot 2^{n-r} \ \overset{(6)}{\leq}\ 2^{n-1}. \quad (8)$$

Thus, $\big((a_j + b_i)y\big) \bmod 2^n < 2^{n-1}$, which implies $\mathrm{MUL}^*_n(a_j + b_i, y) = 0$. This already proves the claim (1) for the case $i \geq j$.

We consider now the case $i < j$. First of all, we have

$$\begin{aligned}(a_{i+1} y) \bmod 2^n + (b_i y) \bmod 2^n \ &\overset{(5),(7)}{\geq}\ (a_i y) \bmod 2^n + 2 \cdot 2^{n-r} + q_i \cdot 2^{n-r} \ \overset{(6)}{\geq} \\ &(a_i y) \bmod 2^n + 2 \cdot 2^{n-r} + 2^{n-1} - (a_i y) \bmod 2^n - 2 \cdot 2^{n-r} \ =\ 2^{n-1}. \quad (9)\end{aligned}$$

Hence, by (4) we also obtain $(a_j y) \bmod 2^n + (b_i y) \bmod 2^n \geq 2^{n-1}$. Thus,

$$\begin{aligned} 2^{n-1} &\leq (a_j y) \bmod 2^n + (b_i y) \bmod 2^n \\ &= (a_j y) \bmod 2^n - (a_i y) \bmod 2^n + (a_i y) \bmod 2^n + (b_i y) \bmod 2^n \\ &\overset{(4),(8)}{<} 2^{n-1} + 2^{n-1} \;=\; 2^n. \end{aligned}$$

These inequalities tell us that $\big((a_j + b_i)y\big) \bmod 2^n \geq 2^{n-1}$. Hence, $\mathrm{MUL}_n^*(a_j + b_i) = 1$. Altogether, we have shown (1). □

In order to derive a lower bound for integer multiplication by the use of Theorem 2.1, we need to define an appropriate filter of the input variables. We use the filters $\mathcal{F}_k(Z)$ which are defined on an m-element variable set Z for $1 \leq k < m$ as $\mathcal{F}_k(Z) = \{M \subseteq Z \mid |M| \geq m - k + 1\}$. This definition ensures that (Z_A, Z_B) is an $\mathcal{F}_k$-partition if and only if $|Z_A| = k$.

In the following let $\mathcal{X}_{n-1} = \{x_1, \ldots, x_{n-1}\}$ and $\mathcal{Y}_{n-1} = \{y_1, \ldots, y_{n-1}\}$ be the input variables for the odd x- and the y-integers, which are multiplied by MUL_n^*.

Lemma 3.4. *Let $k = \lceil n/3 + 2/3 \log(n-1) - 9/2 \rceil$ and $\epsilon = 2^{n/4-k-5/2}$. Further, let $\mathcal{X}_A, \mathcal{X}_B \subseteq \mathcal{X}_{n-1}$ and $\mathcal{Y}_A, \mathcal{Y}_B \subseteq \mathcal{Y}_{n-1}$. If $\Pi = (\mathcal{X}_A \cup \mathcal{Y}_A, \mathcal{X}_B \cup \mathcal{Y}_B)$ is an $\mathcal{F}_k(\mathcal{X}_{n-1} \cup \mathcal{Y}_{n-1})$-partition of $\mathcal{X}_{n-1} \cup \mathcal{Y}_{n-1}$ and f' is (ϵ, Π)-close to MUL_n^*, then $\mathrm{ND}_\omega^{A \to B}(f', \Pi) \geq n/12 - \log(n-1)/3 - 3$ for any $\omega \in \{\vee, \wedge, \oplus\}$.*

Proof. Let f' be (ϵ, Π)-close to MUL_n^*. By definition of $\mathcal{F}_k$ we know that $|\mathcal{X}_A \cup \mathcal{Y}_A| = k$. Hence, there exists a subset $R \subseteq \mathcal{A}(\mathcal{X}_A \cup \mathcal{Y}_A)$ such that $|R| \geq \epsilon\, 2^k$ and $f'(z) = \mathrm{MUL}_n^*(z)$ for all $z \in R \times \mathcal{A}(\mathcal{X}_B \cup \mathcal{Y}_B)$. We may assume w.l.o.g. that $|\mathcal{X}_A| \geq |\mathcal{Y}_A|$, and get the following inequalities

$$\begin{aligned} |\mathcal{X}_A| + |\mathcal{Y}_A| = k, \qquad & |\mathcal{X}_B| + |\mathcal{Y}_B| = 2n - 2 - k, \\ |\mathcal{X}_A| \geq k/2, \qquad & |\mathcal{X}_B| \leq n - 1 - k/2. \end{aligned} \tag{10}$$

Now consider all pairs $(x_A, y_A) \in R$, where $x_A \in \mathcal{A}(\mathcal{X}_A)$ and $y_A \in \mathcal{A}(\mathcal{Y}_A)$. Let for each $y_A \in \mathcal{A}(\mathcal{Y}_A)$ the set $R(y_A) := R \cap \{(x_A, y_A) \mid x_A \in \mathcal{A}(\mathcal{X}_A)\}$. Clearly, $R = \bigcup_{y_A \in \mathcal{A}(\mathcal{Y}_A)} R(y_A)$. Hence, by the pigeonhole principle, there exists a partial assignment y_A for which

$$|R(y_A)| \;\geq\; \frac{|R|}{|\mathcal{A}(\mathcal{Y}_A)|} \;\geq\; \frac{\epsilon\, 2^k}{2^{|\mathcal{Y}_A|}} \overset{(10)}{\geq} \frac{\epsilon\, 2^k}{2^{k/2}} \;=\; \epsilon\, 2^{k/2}. \tag{11}$$

From now on, we let the element $y_A \in \mathcal{A}(\mathcal{Y}_A)$ satisfying inequality (11) be fixed.

The goal of the rest of this proof is to show that an ω-nondeterministic one-way communication protocol, where Alice gets an input from $R(y_A)$ and Bob gets an input from $\mathcal{A}(\mathcal{X}_B \cup \mathcal{Y}_B)$, requires the communication of a large number of bits. Following standard lower bound techniques of nondeterministic communication complexity, this can be done for $\omega \in \{\vee, \wedge, \oplus\}$ by showing that the communication matrix has a large triangular submatrix. More precisely, if the communication matrix has a triangular $(s \times s)$-submatrix, then $\mathrm{ND}_\omega^{A \to B}(f', \Pi) \geq \log s$.

For the inputs in $R(y_A) \times \mathcal{A}(\mathcal{X}_B \cup \mathcal{Y}_B)$, the communication matrix is a matrix where each row is identified with an element in $R(y_A)$ and each column is identified with an element in $\mathcal{A}(\mathcal{X}_B \cup \mathcal{Y}_B)$. The entry of the matrix in a row identified with (x_A, y_A) and a column identified with (x_B, y_B) is the function value $f'(x_A x_B y_A y_B)$ (by $x_A x_B y_A y_B$ we mean the complete input which is consistent with the partial inputs x_A, x_B, y_A, and y_B).

Let now $|x_A|$ be the n-bit integer obtained from the partial assignment x_A by setting all bits which are not fixed by x_A to 0. Analogously define $|x_B|$, $|y_A|$, and $|y_B|$. Then we have $f'(x_A x_B y_A y_B) = \mathrm{MUL}_n^*(|x_A| + |x_B|, |y_A| + |y_B|)$ (recall that f' coincides with MUL_n^* on the inputs $R(y_A) \times \mathcal{A}(\mathcal{X}_B \cup \mathcal{Y}_B)$). Hence, if we let $A = \{|x_A| \mid (x_A, y_A) \in R(y_A)\}$, $B = \{|x_B| \mid x_B \in \mathcal{A}(x_B)\}$, and $Y = \{|y_A| + |y_B| \mid y_B \in \mathcal{A}(\mathcal{Y}_B)\}$, then the communication matrix as described above corresponds to a communication matrix M, where each row is identified with an integer $a \in A$, each column is identified with a pair $(b, y) \in B \times Y$ and where the entry in row a and column (b, y) is the function value $\mathrm{MUL}_n^*(a + b, y)$.

Using inequality (11), we have

$$|A| \;=\; |R(y_A)| \;\geq\; \epsilon \cdot 2^{k/2},$$

and using the inequalities (10), we get for $|B| = 2^\beta$ and $|Y| = 2^\mu$

$$\beta + \mu \;=\; |\mathcal{X}_B \cup \mathcal{Y}_B| \;=\; 2n - 2 - k \quad \text{and} \quad \beta = |\mathcal{X}_B| \;\leq\; n - 1 - k/2.$$

Therefore, $\beta + 3\mu$ is minimal if $\beta = \mu = n - 1 - k/2$. Thus, $\beta + 3\mu \geq 4n - 4 - 2k$. Using these parameters in Lemma 3.1 together with the precondition $\epsilon = 2^{n/4-k-5/2}$ and $k = \lceil n/3 + (2/3)\log(n-1) - 9/2 \rceil$ yields that M has a triangular $(s \times s)$-submatrix where

$$\begin{aligned} s \;&=\; \min\left\{\epsilon 2^{k/2-1} - 1,\, 2^{n/4-k/2-14/4} - 1\right\} \;=\; 2^{n/4-k/2-7/2} - 1 \\ &\geq\; 2^{n/4-(n/3+(2/3)\log(n-1)-7/2)/2-7/2} - 1 \;=\; 2^{n/12-\log(n-1)/3-7/4} - 1. \end{aligned}$$

By the discussion above, we can finally conclude that

$$\mathrm{ND}_\omega^{A\to B}(f', \Pi) \geq \log s \geq n/12 - \log(n-1)/3 - 3.$$

□

A simple calculation using the parameters from the lemma above shows that $\left(\epsilon \cdot 4(n-1)\right)^{-1} \geq 2^{n/12-\log(n-1)/3-4}$. Using Corollary 2.3, this yields the following exponential lower bound for well-structured g.d.-BP1s representing MUL_n.

Corollary 3.5. *Let $\omega \in \{\vee, \wedge, \oplus\}$. The size of any ω-nondeterministic well-structured graph-driven* BP1 *for* MUL_n *is larger than* $2^{n/12-4} \cdot (n-1)^{-1/3}$.

4 Separating Nondeterministic Well-Structured Graph-Driven BP1s from Nondeterministic BP1s

Here we answer an open question from Brosenne, Homeister, and Waack [7], whether the class of all boolean functions representable in polynomial size by ω-nondeterministic well-structured graph-driven BP1s is a proper subclass of all boolean functions representable in polynomial size by ω-nondeterministic BP1s, in an affirmative way.

The function $n/2$-MRC$_n$ is defined on an $n \times n$ boolean matrix X on the variables $\mathcal{X}_{n\times n} = \{x_{1,1}, \ldots, x_{n,n}\}$. Its function value is 1 if and only if the following two conditions are fulfilled (for the sake of readability we assume that n is an even number.)

1. The number of ones in the matrix is at least $n^2/4+n$ and at most $(3/4)n^2-n$.
2. The matrix either contains exactly $n/2$ monochromatic rows and each non-monochromatic row contains exactly $n/2$ ones, or it contains exactly $n/2$ monochromatic columns and each non-monochromatic column contains exactly $n/2$ ones.

Note that because of condition 1, there cannot be $n/2$ monochromatic rows and $n/2$ monochromatic columns for a satisfying input. Furthermore, if condition 2 is satisfied, then condition 1 is fulfilled if and only if at least one of the monochromatic rows (columns) satisfying condition 2 consists only of ones, and at least one of the monochromatic rows (columns) consists only of zeros.

The branching program model for which we show the upper bound is even more restricted than the general ω-nondeterministic BP1 model.

Definition 4.1. *An (ω, k)-*PBDD *G consists of k* OBDDs *$G_1, \ldots, G_k$ whose variable orders may be different. If $f_1, \ldots, f_k$ are the functions represented by $G_1, \ldots, G_k$, then G represents the function $f_1 \,\omega\, f_2 \,\omega \cdots \omega\, f_k$. The size of G is $|G| = |G_1| + \cdots + |G_k|$.*

Note that we can regard an (ω, k)-PBDD as an ω-nondeterministic BP1 which has $k-1$ nondeterministic nodes at the top, which generate k paths leading to the disjoint OBDDs $G_1, \ldots, G_k$. Motivated by applications, the model of $(\vee, k)$-PBDDs has been introduced in [15].

Theorem 4.2. *For $\omega \in \{\vee, \oplus\}$, the function $n/2$-MRC$_n$ can be represented by $(\omega, 2)$-*PBDD*s with size $O(n^4)$, but its complexity is $\Omega(2^{n/4}/n)$ for ω-nondeterministic well-structured graph-driven* BP1*s.*

Proof. A rowwise (columnwise) variable order is an order, where all variables of one row (column) are tested one after another. We show how to construct two OBDDs G_1, G_2 such that G_1 and G_2 accept exactly the satisfying inputs containing $n/2$ monochromatic rows and $n/2$ monochromatic columns, respectively. Clearly, the set $\{G_1, G_2\}$ then is a $(\vee, 2)$-PBDD for $n/2$-MRC$_n$ and - because any satisfying

input contains either only $n/2$ monochromatic rows or only $n/2$ monochromatic columns, but not both - also a valid $(\oplus, 2)$-PBDD for $n/2$-MRC$_n$.

The OBDD G_1 checks whether there exist $n/2$ monochromatic rows, where at least one row consists only of ones and at least one row consists only of zeros. In addition, it tests whether each non-monochromatic row contains exactly $n/2$ ones. If this is not the case, then it reaches the 0-sink. All variables are tested in a rowwise variable order. During the tests of the variables $x_{i,1}, \ldots, x_{i,n}$ the OBDD stores the information of the number of already tested monochromatic rows, the information whether at least one of them is 1-monochromatic and at least one of them is 0-monochromatic, and the number of already tested ones in the ith row. Hence, it suffices to distinguish at most $4n^2$ situations, and G_1 contains $O(n^4)$ nodes. The OBDD G_2 checks the same as G_1, but for the columns. This can be done in an analogous way but using a columnwise variable order. Altogether the $(2, \omega)$-PBDD size of $n/2$-MRC$_n$ is bounded above by $O(n^4)$ for $\omega \in \{\vee, \oplus\}$.

Now we prove the lower bound. Again we apply the technique from Corollary 2.3. In order to do so, we have to define an appropriate filter $\mathcal{F}_{\mathrm{M}}$ of the variable set $\mathcal{X}_{n\times n}$. A set $T \subseteq \mathcal{X}_{n\times n}$ is in the filter $\mathcal{F}_{\mathrm{M}}$, if T contains all variables from $n/2+1$ arbitrary rows and $n/2+1$ arbitrary columns. If $\Pi = (\mathcal{X}_A, \mathcal{X}_B)$ is an $\mathcal{F}_{\mathrm{M}}$-partition, then by definition $\mathcal{X}_B \notin \mathcal{F}_{\mathrm{M}}$ and there exists a variable $x_{i,j}$ such that $\mathcal{X}_B \cup \{x_{i,j}\} \in \mathcal{F}_{\mathrm{M}}$. Hence, $\mathcal{X}_A$ contains variables from exactly $n/2$ different rows and from at most $n/2$ different columns or vice versa.

The lower bound of Theorem 4.2 follows right away from the following lemma and Corollary 2.3 by choosing $\epsilon = 1/\big(n \cdot 2^{n/4}\big)$.

Lemma 4.3. *Let $0 < \epsilon \leq 1$ and Π be an arbitrary $\mathcal{F}_{\mathrm{M}}$-partition of $\mathcal{X}_{n\times n}$. Then for every function f' which is (ϵ, Π)-close to $n/2$-MRC$_n$, it is $\mathrm{ND}_{\omega}^{A\to B}(f', \Pi) \geq n/2 + \log \epsilon$.*

Proof. Let $\Pi = (\mathcal{X}_A, \mathcal{X}_B)$ be an $\mathcal{F}_{\mathrm{M}}$-partition and f' be (ϵ, Π)-close to $n/2$-MRC$_n$. We may assume w.l.o.g. that $\mathcal{X}_A$ contains variables from exactly the rows $1, \ldots, n/2$, whereas there are at most $n/2$ columns from which variables are contained in $\mathcal{X}_A$. Since f' is (ϵ, Π)-close to $n/2$-MRC$_n$, there exists a subset $R \subseteq \mathcal{A}(\mathcal{X}_A)$, $|R| \geq \epsilon \cdot 2^{|\mathcal{X}_A|}$, such that f' coincides with $n/2$-MRC$_n$ on all inputs in $R \times \mathcal{A}(\mathcal{X}_B)$. For $1 \leq i \leq n/2$ let k_i be the number of variables in row i which are contained in $\mathcal{X}_A$. We consider the mapping

$$\mu : \mathcal{A}(\mathcal{X}_A) \to \{0, \ldots, k_1\} \times \cdots \times \big\{0, \ldots, k_{n/2}\big\},$$

which maps a partial assignment α to the tuple $\mu(\alpha) = (z_1, \ldots, z_{n/2})$, where z_i is the number of bits in row i being fixed to 1 by α.

Let $\mu(R) = \{\mu(\alpha) \mid \alpha \in R\}$. Below, we show the following two inequalities from which the lemma follows right away.

(I1) $\mathrm{ND}_{\omega}^{A\to B}(f', \Pi) \geq \log |\mu(R)|$.
(I2) $|\mu(R)| \geq \epsilon \cdot 2^{n/2}$.

Proof of (I1): We show that the communication matrix contains a diagonal $s \times s$-submatrix, where $s = |\mu(R)|$. For an arbitrary partial assignment $\alpha \in R$ let $\mu(\alpha) = \big(\mu_1(\alpha), \ldots, \mu_{n/2}(\alpha)\big)$. We fix for each such α a corresponding partial assignment $\beta \in \mathcal{A}(\mathcal{X}_B)$ as follows. In row i, $1 \leq i \leq n/2$, β sets exactly $n/2 - \mu_i(\alpha)$ variables to 1 and the other variables to zero. (Recall that $\mathcal{X}_A$ contains variables from at most $n/2$ columns, and hence at least $n/2$ variables from each row are in $\mathcal{X}_B$.) All the variables in the rows $n/2+1, \ldots, n-1$ are fixed to 0 and the variables in row n are all set to 1. Then $(\alpha\beta)$ contains exactly $n/2$ rows with exactly $n/2$ ones each (the rows $1, \ldots, n/2$), and it contains $n/2-1$ 0-monochromatic rows and one 1-monochromatic row. Hence, $n/2\text{-MRC}_n(\alpha\beta) = 1$.

We consider now s arbitrary partial assignments $\alpha_1, \ldots, \alpha_s \in R$ such that $\mu(\alpha_i) \neq \mu(\alpha_j)$ for $i \neq j$. Let $\beta_1, \ldots, \beta_s$ be the corresponding partial assignments in $\mathcal{A}(\mathcal{X}_B)$. (It is obvious that also $\beta_i \neq \beta_j$ for $i \neq j$.) Clearly, the $s \times s$-matrix which has in row i and column j the entry $n/2\text{-MRC}_n(\alpha_i\beta_j)$ is a submatrix of the communication matrix of $n/2\text{-MRC}_n$. Hence, for the claim (I1), it suffices to show that this matrix is a diagonal matrix. For the diagonal elements, we have already proven above that $n/2\text{-MRC}_n(\alpha_i\beta_i) = 1$. Consider now an element in row i and column j, $i \neq j$. Since $\alpha_i \neq \alpha_j$, there exists an index $1 \leq t \leq n/2$ for which $\mu_t(\alpha_i) \neq \mu_t(\alpha_j)$. Hence, by construction the matrix X defined by the input $\alpha_j\beta_i$ contains in row t not exactly $n/2$ ones. But the construction also ensures that none of the rows $n/2+1, \ldots, n$ of X contains exactly $n/2$ ones, thus there exist less than $n/2$ rows with exactly $n/2$ ones. Finally, the property that row n is 1-monochromatic and the row $n-1$ is 0-monochromatic ensures that there exists no monochromatic column. Altogether, this yields that $n/2\text{-MRC}_n(\alpha_i\beta_j) = 0$.

Proof of (I2): Recall that $\mathcal{X}_A$ contains k_i variables in row i of the matrix X ($1 \leq i \leq n/2$). Hence, there are exactly 2^{k_i} possible settings of those variables in row i and among these, there are $\binom{k_i}{z_i}$ settings for which row i contains exactly z_i ones. Hence, for every tuple $z = (z_1, \ldots, z_{n/2}) \in \{0, \ldots, k_1\} \times \cdots \times \{0, \ldots, k_{n/2}\}$ we obtain that

$$\frac{|\mu^{-1}(z)|}{|\mathcal{A}(\mathcal{X}_A)|} \;=\; \frac{\binom{k_1}{z_1} \cdots \binom{k_{n/2}}{z_{n/2}}}{2^{k_1} \cdots 2^{k_{n/2}}} \;\leq\; \frac{2^{k_1-1} \cdots 2^{k_{n/2}-1}}{2^{k_1} \cdots 2^{k_{n/2}}} \;=\; 2^{-n/2}. \tag{12}$$

Since R is the union of all $\mu^{-1}(z)$ for $z \in \mu(R)$, there exists by the pigeonhole principle an element $z \in \mu(R)$ for which $|\mu^{-1}(z)| \geq |R|/|\mu(R)|$. Using the precondition that $|R| \geq \epsilon \cdot 2^{|\mathcal{X}_A|}$ together with inequality (12) yields

$$|\mu(R)| \;\geq\; \frac{|R|}{|\mu^{-1}(z)|} \;\geq\; \frac{\epsilon \cdot 2^{|\mathcal{X}_A|}}{2^{-n/2} \cdot |\mathcal{A}(\mathcal{X}_A)|} \;=\; \epsilon \cdot 2^{n/2}.$$

This finally proves (I2). □

Acknowledgment

We would like to thank Martin Sauerhoff and Ingo Wegener for fruitful discussions about the subject of this paper and helpful comments.

References

1. M. Ajtai. A non-linear time lower bound for boolean branching programs. In *Proc. of 40th FOCS*, pp. 60–70. 1999.
2. P. Beame, M. Saks, X. Sun, and E. Vee. Super-linear time-space tradeoff lower bounds for randomized computation. In *Proc. of 41st FOCS*, pp. 169–179. 2000.
3. P. Beame and E. Vee. Time-space tradeoffs, multiparty communication complexity, and nearest neighbor problems. In *Proc. of 34th ACM STOC*, pp. 688–697. 2002.
4. B. Bollig. Restricted nondeterministic read-once branching programs and an exponential lower bound for integer multiplication. *RAIRO*, 35:149–162, 2001.
5. B. Bollig, S. Waack, and P. Woelfel. Parity graph-driven read-once branching programs and an exponential lower bound for integer multiplication. In *Proc. of 2nd TCS*, pp. 83–94. 2002.
6. B. Bollig and P. Woelfel. A read-once branching program lower bound of $\Omega(2^{n/4})$ for integer multiplication using universal hashing. In *Proc. of 33rd ACM STOC*, pp. 419–424. 2001.
7. H. Brosenne, M. Homeister, and S. Waack. Graph-driven free parity BDDs: Algorithms and lower bounds. In *Proc. of 26th MFCS*, pp. 212–223. 2001.
8. R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. on Comp.*, C-35:677–691, 1986.
9. R. E. Bryant. On the complexity of VLSI implementations and graph representations of boolean functions with applications to integer multiplication. *IEEE Trans. on Comp.*, 40:205–213, 1991.
10. C. Damm, M. Krause, C. Meinel, and S. Waack. Separating counting communication complexity classes. In *Proc. of 9th STACS*, pp. 281–292. 1992.
11. J. Gergov. Time-space tradeoffs for integer multiplication on various types of input oblivious sequential machines. *Information Processing Letters*, 51:265–269, 1994.
12. J. Gergov and C. Meinel. Frontiers of feasible and probabilistic feasible boolean manipulation with branching programs. In *Proc. of 11th STACS*, pp. 576 –585. 1993.
13. J. Gergov and C. Meinel. Efficient analysis and manipulation of OBDDs can be extended to FBDDs. *IEEE Trans. on Comp.*, 43:1197–1209, 1994.
14. J. Hromkovič. *Communication Complexity and Parallel Computing*. Springer, 1997.
15. J. Jain, J. Bitner, D. S. Fussell, and J. A. Abraham. Functional partitioning for verification and related problems. In *Brown MIT VLSI Conf.*, pp. 210–226. 1992.
16. M. Krause. BDD-based cryptanalysis of keystream generators. In *Proc. of EUROCRYPT*, pp. 222–237. 2002.
17. E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
18. C. Meinel. The power of polynomial size Ω-branching programs. In *Proc. of 5th STACS*, pp. 81–90. 1988.
19. S. Ponzio. A lower bound for integer multiplication with read-once branching programs. *SIAM Journal on Computing*, 28:798–815, 1998.
20. D. Sieling and I. Wegener. Graph driven BDDs – a new data structure for Boolean functions. *Theor. Comp. Sci.*, 141:283–310, 1995.
21. D. Sieling and I. Wegener. A comparison of free BDDs and transformed BDDs. *Formal Methods in System Design* 19, pp. 223–236, 2001.
22. I. Wegener. *Branching Programs and Binary Decision Diagrams - Theory and Applications*. SIAM, 2000.
23. P. Woelfel. New bounds on the OBDD-size of integer multiplication via universal hashing. In *Proc. of 18th STACS*, pp. 563–574. 2001.
24. P. Woelfel. On the complexity of integer multiplication in branching programs with multiple tests and in read-once branching programs with limited nondeterminism. In *Proc. of 17th Conf. on Comp. Compl.*, pp. 80–89. 2002.