Electronic Colloquium on Computational Complexity, Report No. 36 (1998)

ECCC

Lower Bounds for (MOD p - MOD m) Circuits

Preliminary version

Vince Grolmusz * Gábor Tardos †

Abstract

Modular gates are known to be immune for the random restriction techniques of Ajtai [Ajt83], Furst, Saxe, Sipser [FSS84], Yao [Yao85] and Håstad [Hås86]. We demonstrate here a random clustering technique which overcomes this difficulty and is capable to prove generalizations of several known modular circuit lower bounds of Barrington, Straubing, Thérien [BST90], Krause and Pudlák [KP94], and others, characterizing symmetric functions computable by small $(\mathrm{MOD}_p, \mathrm{AND}_t, \mathrm{MOD}_m)$ circuits.

Applying a degree-decreasing technique together with random restriction methods for the AND gates at the bottom level, we also prove a hard special case of the Constant Degree Hypothesis of Barrington, Straubing, Thérien [BST90], and other related lower bounds for certain $(\mathrm{MOD}_p, \mathrm{MOD}_m, \mathrm{AND})$ circuits.

Most of the previous lower bounds on circuits with modular gates used special definitions of the modular gates (i.e., the gate outputs one if the sum of its inputs is divisible by m, or is *not* divisible by m), and were not valid for more general MOD_m gates. Our methods are applicable – and our lower bounds are valid – for the most general modular gates as well.

1 Introduction

Boolean circuits are perhaps the most widely examined models of computation. They are used in VLSI design, and in complexity theory as well as in the theory of parallel computation.

A majority of the strongest and deepest lower bound results for computational complexity were proved using the Boolean circuit model of computation (for example [Raz85], [Yao85], [Hås86], [Raz87], [Smo87], or see [vL90] for a survey).

Unfortunately, lots of questions — even for very restricted circuit classes — have been unsolved for a long time.

Bounded depth and polynomial size is a natural restriction. Ajtai [Ajt83], Furst, Saxe, and Sipser [FSS84] proved that no polynomial sized, constant depth circuit can compute the PARITY function. Yao [Yao85] and Håstad [Hås86] generalized this result for sub-logarithmic depths. Their technique involved a sophisticated use of *random restriction techniques*, in which randomly assigned 0-1 values to the input variables fixed the output of large fan-in AND and OR Boolean gates.

*Department of Computer Science, Eötvös University, Budapest, Múzeum krt.6-8, H-1088 Budapest, Hungary; E-mail: grolmusz@cs.elte.hu

ISSN 1433-8092

†Mathematical Institute of the Hungarian Academy of Science, Reáltanoda u. 13-15, H-1055 Budapest, Hungary; E-mail: tardos@cs.elte.hu

Since the modular gates are very simple to define, and they are immune to the random restriction techniques in lower bound proofs for the PARITY function, the following natural question was asked by Barrington, Smolensky and others: How powerful will become the Boolean circuits if — beside the standard AND, OR and NOT gates — MOD_m gates are also allowed in the circuit? Here a MOD_m^A gate outputs 1 iff the sum of its inputs is in a set $A \subset \{0, 1, 2, \ldots, m-1\}$ modulo m.

Razborov [Raz87] showed that for computing MAJORITY with AND, OR, NOT and MOD_2 gates, exponential size is needed with constant depth. This result was generalized by Smolensky [Smo87] for MOD_p gates instead of MOD_2 ones, where p denotes a prime.

We know very little, however, if both MOD_p and MOD_q gates are allowed in the circuit for different primes p, q, or, if the modulus is a non-prime power composite, e.g., 6. For example, it is consistent with our present knowledge that depth-3, linear-size circuits with MOD_6 gates *only*, recognize an NP-complete language (see [BST90]).

It is not difficult to see that constant-depth circuits with MOD_p gates only, (p prime), cannot compute even very simple functions: the n-fan-in OR or AND functions, since they can only compute constant degree polynomials of the input variables over GF_p (see [Smo87]).

But depth-2 circuits with MOD_2 and MOD_3 gates, or MOD_6 gates can compute the n-fan-in OR and AND functions [KM91], [BST90]. Consequently, these circuits are more powerful than circuits with MOD_p gates only. The sketch of the construction: we take a MOD_3 gate at the top of the circuit, and 2^n MOD_2 gates on the next level, and each subset of the n input variables is connected to exactly one MOD_2 gate, then this circuit computes the n-fan-in OR, since if at least one of the inputs is 1, then exactly half of the MOD_2 gates evaluate to 1.

Barrington, Straubing and Thérien in [BST90] conjectured that any $(\mathrm{MOD}_p^B, \mathrm{MOD}_m^A, \mathrm{AND}_d)$ circuit needs exponential size to compute the n fan-in AND function, where the prime p and the positive integers m and d are fixed, and AND_d denotes the fan-in d AND function. They called it the *Constant Degree Hypothesis* (CDH), and proved the $d = 1$ case, with highly non-trivial algebraic techniques. Their proof also works for depth-$(\ell + 1)$

$$(\overbrace{\mathrm{MOD}_{p^k}^B, \mathrm{MOD}_{p^k}^B, \ldots, \mathrm{MOD}_{p^k}^B}^{\ell}, \mathrm{MOD}_m^A) \tag{1}$$

circuits, computing the AND function.

Yan and Parberry [YP94] – using Fourier-analysis – proved also the $d = 1$ case for $(\mathrm{MOD}_p^{\{1,2,\ldots,p-1\}}, \mathrm{MOD}_2^{\{1\}})$ circuits, but their method also works for the special case of the CDH where the sum of the degrees of the monomials g_i on the input-level satisfies:

$$\sum_{\deg(g_i)\geq 1} (\deg(g_i) - 1) \leq \frac{n}{2(p-1)} - O(1).$$

Krause and Waack [KW95] applied communication-complexity techniques to show that any $(\mathrm{MOD}_m^{\{1,2,\ldots,m-1\}}, \mathrm{SYMMETRIC})$ circuit, computing the ID function:

$$\mathrm{ID}(x, y) = \begin{cases} 1, \text{ if } x = y, \\ 0 \text{ otherwise,} \end{cases}$$

for $x, y \in \{0, 1\}^n$, should have size at least $2^n / \log m$, where SYMMETRIC is a gate, computing an arbitrary symmetric Boolean function. Since (non-weighted) MOD_m gates are

also SYMMETRIC gates, this lower bound is valid for $(\mathrm{MOD}_m^{\{1,2,\ldots,m-1\}}, \mathrm{MOD}_m^A)$ circuits. However, when mod m coefficients (or multiple wires) are allowed on the input-level, then the MOD_m gates are *not* SYMMETRIC gates. Caussinus [Cau96] proved, that the result of [BST90] also implies a similar lower bound for the AND function. Unfortunately, results [KW95], [Cau96] do not generalize for the more general MOD_m^A gates at the top.

Krause and Pudlák [KP94] proved that any $(\mathrm{MOD}_{p^k}^{\{0\}}, \mathrm{MOD}_q^{\{0\}})$ circuit which computes the $\mathrm{MOD}_r^{\{0\}}$ function has size at least 2^{cn}, for some $c > 0$, where p and r are different primes and q is not divisible by either of them.

Our main result is a characterization of those symmetric Boolean functions which are computable by quasipolynomial-size

$$(\overbrace{\mathrm{MOD}_{p^k}^B, \mathrm{MOD}_{p^k}^B, \ldots, \mathrm{MOD}_{p^k}^B}^{\ell}, \mathrm{MOD}_m^A)$$

circuits. We prove (Theorem 4), that the *only* symmetric functions that are computable by such circuits are the MOD_{mp^j} functions with small j. Consequently, the non-trivial threshold functions, (so also AND and OR), and the $\mathrm{MOD}_r^{\{0\}}$ functions if r does not divide $p^j m$ need exponential size on that circuits. Even MOD_4 requires exponential size $(\mathrm{MOD}_{3^r}, \mathrm{AND}_t, \mathrm{MOD}_2)$ circuits for constant t and r. Note the asymmetry: MOD_4 is easy to compute with a polynomial size $(\mathrm{MOD}_2, \mathrm{AND}_3)$ circuit. These results generalize the theorems of Barrington, Straubing, Thérien [BST90] and Krause and Pudlák [KP94], and give a characterization of the computable symmetric functions, instead of singular lower bounds.

Grolmusz [Gro98] generalized the results of [BST90], [YP94], [KW95], [KP94] for $(\mathrm{MOD}_q, \mathrm{MOD}_p, \mathrm{AND}_{cn})$ circuits, where the input-polynomials of each MOD_p gate is constructible from linear terms using at most $cn - 1$ multiplications (or, equivalently, can be computed by an arithmetic circuit of an arbitrary number of mod p additions and at most $cn - 1$ fan-in 2 multiplications). In particular, one can allow the sum of *an arbitrary function* of cn variables and a linear polynomial of the n variables as inputs for each MOD_p gate. We generalize this result, too (Lemma 19). The main tool of the proof of [Gro98] is a Degree Decreasing Lemma, which we also generalize here for non-prime moduli (Lemma 16), and we use it both for lower- and upper bound proofs.

Here we give further generalizations for this circuit lower-bound result. We prove a lower bound on the size of the $(\mathrm{MOD}_p, \mathrm{MOD}_m, \mathrm{AND})$ circuits computing AND_n, if m is a positive integer, p is a prime, and each MOD_m gate has not-too-many AND gates as inputs and those AND gates have low fan-in. For the exact statement see Theorem 6. This is an important special case of the Constant Degree Hypothesis of [BST90]. The lower bound also applies to circuits computing some other functions besides AND.

2 Our Results

2.1 Ideas

MOD_m gates are immune to random restriction techniques, since if at least m input variables of n remain unrestricted, then on the rest it is still a MOD_m gate and has the same complexity.

We overcome this difficulty by a *random clustering* technique, which force some randomly chosen variables to be equal. Each equivalence class (or cluster) will make a new variable of

the MOD_m gate, and each new variable will be invisible (i.e., its coefficient will be a multiple of m) for the gate with a constant probability. (Lemma 11)

We use this for $(\sum_p, \mathrm{AND}_t, \mathrm{MOD}_m)$ circuits, computing symmetric functions. Suppose, that the equivalence classes are of size m, then the resulting function of the new, clustered variables, is a unique symmetric function.

Almost all symmetric functions (except the $\mathrm{MOD}_{p^k m}$ functions) have large restrictions, whose unique factor resulting from the clustering above cannot be expressed as a modulo p sum of functions, none of which depends on all variables. An exponential lower bound follows for the number of AND gates on level 2. (Theorem 3)

If we have $o(n^2/\log n)$ constant-degree monomials as inputs for each MOD_m gates on level 2, then by random restrictions, one can essentially decrease their number, and a small number of low-degree monomials can be converted to linear polynomials with the help of the Degree Decreasing Lemma (Lemma 16), and we can apply Theorem 3 to get lower bounds. (Theorem 6)

2.2 Preliminaries

Definition 1 *A fan-in n gate is an n-variable Boolean function. Let $G_1, G_2, \ldots, G_\ell$ be gates of unbounded fan-in. Then a $(G_1, G_2, \ldots, G_\ell)$- circuit denotes a depth-ℓ circuit with a G_1-gate on the top, G_2 gates on the second level, G_3 gates on the third level from the top,..., and G_ℓ gates on the last level.* AND_t *denotes the t fan-in* AND *gate. The size of a circuit is defined to be the total number of the gates in the circuit.*

All of our modular gates are of unbounded fan-in, and we allow to connect inputs to gates or gates to gates with multiple wires. Note, that by this definition, our modular gates – generally – are not symmetric gates.

In the literature MOD_m gates are sometimes defined to be 1, iff the sum of their inputs is divisible by m, and sometimes they are defined to be 1, iff the sum of their inputs is not divisible by m. The following, more general definition covers both cases.

Definition 2 *We say that gate G is a* MOD_m-gate, *if there exists $A \subset \{0, 1, \ldots, m-1\}$ such that*

$$G(x_1, x_2, \ldots, x_n) = \begin{cases} 1, & \text{if } \sum_{i=1}^n x_i \bmod m \in A \\ 0 & \text{otherwise.} \end{cases}$$

A is called the 1-set of G. MOD_m *gates with 1-set A are denoted by* MOD_m^A.

2.3 Theorems

Here we list the three main results of this paper.
To be concise we use $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A)$ to denote circuits of type (1).

Theorem 3 *Suppose that a circuit of type $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A)$ with p prime computes a symmetric Boolean function f on n variables, such that $f \neq \mathrm{MOD}_{p^j m}^A$ for any A. Then its size S is exponential in p^j, i.e., there exists a number $c > 1$ depending on p, m, k, and ℓ such that $S > c^{p^j}$.*

As a special case we get that the size S of an n-variable circuit of type $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A)$ with p prime computing any of the nontrivial threshold functions (including AND *and* OR*)*

or the $\mathrm{MOD}_r^{\{0\}}$ *function (where* r *does not divide* mp^j *for any* j*) is exponential in* n*. We have* $S > c^n$ *for a number* $c > 1$ *depending only on* p*,* m*,* k*, and* ℓ*.*

Theorem 4 *Let the prime* p*, and the positive integers* m*,* k*, and* ℓ *be fixed with* m *not a power of* p*. The symmetric functions computed by a type* $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A)$ *circuit of quasipolynomial size are exactly the functions* $\mathrm{MOD}_{mp^j}^C$ *with* $j = O(\log\log n)$ *and* $C \subset \{0, 1, \ldots, mp^j - 1\}$*.*

On the other hand, all the functions $\mathrm{MOD}_{mp^j}^C$ *with* $j = O(\log\log n)$ *can be computed by quasipolynomial size* $(\sum_p, \mathrm{AND}_2, \mathrm{MOD}_m)$ *circuits.*

To state our result that proves a special case of the Constant Degree Hypothesis we need the following definition:

Definition 5 *Let* G *be a gate of a circuit on the second level from the inputs computing some function of* AND*'s of variables. We say that* G *relates two input variables if they appear as inputs in a common* AND *gate below* G*.*

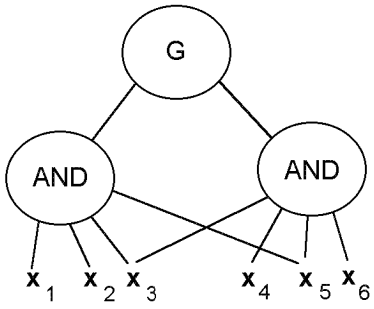

Figure 1: Gate G relates e.g., x_1 and x_2, or x_3 and x_6, but does not relate x_1 and x_4.

Theorem 6 *Let* p *be prime and* m*,* k*, and* ℓ *fixed positive integers. Suppose that a* $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A, \mathrm{AND})$ *circuit computes* AND_n*. If each* MOD_m *gate in the circuit relates* $o(n^2/\log n)$ *pairs of input variables then the size of the circuit is super-polynomial.*

As a special case we get that the size of the circuit is super-polynomial if each MOD_m *gate has fan-in* $o(n^2/\log n)$ *and each* AND *gate has constant fan-in.*

We remark that there is a trade-off between the number of pairs the MOD_m gates relate and the lower bound on the circuit size. See Theorem 20 If no MOD_m gate relates more than $n^{2-\varepsilon}$ pairs then the circuit has to be exponential size. Note also, that similar bounds can be proved for circuits computing many other natural functions, like threshold or MOD_r functions.

3 The Proof

3.1 Eliminating the top gate

The top-gate elimination is widely used in the literature (c.f., [KP94], Lemma 5.2, or [BT91]). It replaces the top MOD_{p^r} gate with constant fan-in AND gates and a simple summation

modulo p with a polynomial increase in the size.

Lemma 7 *Let p be a prime, k a positive integer, and $A \subset \{0,1,\ldots,p^k-1\}$. There is a modulo p polynomial of degree p^k-1 computing the $\mathrm{MOD}^A_{p^k}$ function.*

□

One can repeatedly use this lemma to eliminate a constant-depth sub-circuit of MOD_{p^r} gates from the top of any circuit, as stated by the next lemma. For stating it, we need

Notation 8 *Let $\Sigma_p(x_1,x_2,\ldots,x_s)=\sum_{i=1}^{s} x_i \bmod p$.*

In general, Σ_p is not a Boolean gate, since its value is from $\{0,1,\ldots,p-1\}$. But, in the following lemma, its value will always be 0 or 1.

Lemma 9 *Suppose that $f:\{0,1\}^n \to \{0,1\}$ is computed by a depth-$(\ell+1)$*

$$(\overbrace{\mathrm{MOD}^A_{p^k},\ldots,\mathrm{MOD}^A_{p^k}}^{\ell},G)$$

circuit, where p is a prime and on the input level we have arbitrary gates (or sub-circuits) G. Suppose the number of these gates G is S. Then f can also be computed from the same gates G by a $(\Sigma_p,\mathrm{AND}_t,G)$ circuit, with $t<p^{k\ell}$ and at most $S^{p^{k\ell}}$ AND_t gates on the middle level.

Proof: By Lemma 7 all $\mathrm{MOD}^A_{p^k}$ can be replaced by a modulo p polynomial of degree less than p^k thus f is degree $< p^{k\ell}$ polynomial of the output of the G gates. The bound on the size comes from counting all the possible monomials in such a polynomial. □

Note, that the size of the new circuit is still polynomial in S and the fan-in of the AND gates is constant if depth ℓ and modulus p^k are constants. Note that AND_t gates with $t<p^k$ can be considered as special MOD_{p^k} gates and thus constant fan-in AND gates can be eliminated the same way.

3.2 Random Clustering

Definition 10 *Let $\sim$ be an equivalence relation on the variables of a function f. By the factor $f/\sim$ of f we mean the function obtained from f by identifying variables according to $\sim$. The variables of $f/\sim$ correspond to the equivalence classes of $\sim$. For an integer m we call the $f/\sim$ an m-factor of f if each equivalence class in $\sim$ consists of m variables.*

We say that the Boolean function f is p-simple (p is a positive integer) if it can be expressed as a modulo p sum of functions none of which depend on all of the variables.

Example. *Suppose that f has 6 variables, and $x_1{\sim}x_2, x_3{\sim}x_4, x_5{\sim}x_6$. Then $f/\sim$ is a 2-factor of f, has three variables, and is defined as*

$$f/{\sim}(y_1,y_2,y_3)=f(y_1,y_1,y_2,y_2,y_3,y_3).$$

Notice that any factor of the AND function is again an AND function. The m-factor of a symmetric function is unique and it is also a symmetric function. Note that for prime

numbers p a function f is p-simple if and only if it can be expressed as a modulo p polynomial of degree less than the number of its variables.

The following lemma is about a special type of three level circuits. It is stated in a more general way but the reader may think of polynomial size $(\sum_p, \mathrm{AND}_t, \mathrm{MOD}_m^A)$ circuits with constant t.

Lemma 11 *Let p, m, and t be positive integers, $1 \geq \varepsilon > 0$ and suppose the Boolean function f on n variables satisfies $f \equiv \sum_{i=1}^{S} f_i \pmod p$, where each f_i is computed in an arbitrary way from t of the functions f_{ij} and from $(1-\varepsilon)n$ of the input variables. Each of the functions f_{ij} is in turn a modulo m linear combination of the input variables. Here the functions f_i output modulo p values while f_{ij} output modulo m values. If n is divisible by m and $S < c^n$ then there exists a p-simple m-factor of f, where the constant $c > 1$ depends only on m, t, and ε.*

Proof: The idea is to realize that $f_{ij}/\sim$ is a modulo m linear combination of its variables and the coefficient of a variable — corresponding to an equivalence class in a random $\sim$ — is equal to zero with a positive constant probability. Thus $f_i/\sim$ depends on all of its variables with exponentially small probability. With high probability all the functions $f_i/\sim$ has an invisible variable and thus $f/\sim$ is p-simple.

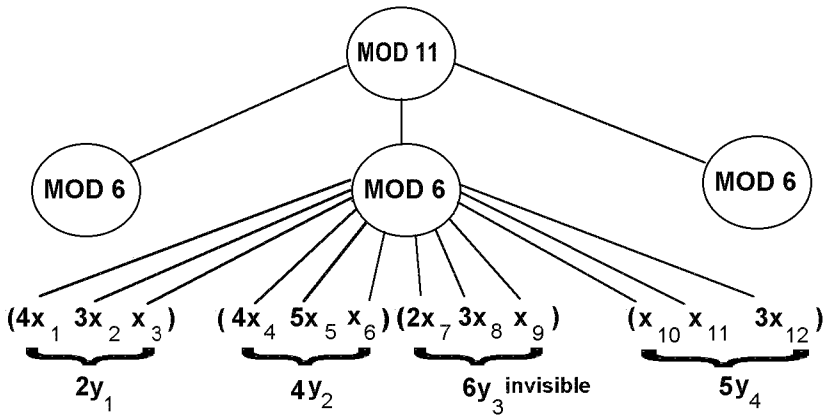

Figure 2: Random clustering in the simplest case: $t = 1, \varepsilon = 1$, every f_{i1} is a MOD_6 gate.

Let us choose $\sim$ uniformly at random from all the partitions of the variables to classes of size m. Consider choosing the equivalence classes one by one. For a fixed $1 \leq i \leq S$ and for the first $\varepsilon n/(2m)$ classes, with probability at least $(\varepsilon m^{-t}/4)^m$, all the m variables in the class has the same coefficient in all the t combinations $f_{ij}, j = 1, 2, \ldots, t$ and all of them are outside the $(1-\varepsilon)n$ variables "seen" by f_i. If this is the case, $f_i/\sim$ does not depend on the variable corresponding to this class. Thus with probability at most $e^{-(\varepsilon n/(2m))\cdot(\varepsilon m^{-t}/4)^m)}$ does $f_i/\sim$ depend on each of its variables. We choose $\ln c = (\varepsilon/4)^{m+1} m^{tm+1}$. If $S < c^n$ then with positive probability none of the functions $f_i/\sim$ depend on all of the variables, thus $f/\sim \equiv \sum_{i=1}^{S} f_i/\sim \pmod p$ is p-simple. □

We remark here that the same proof gives that if S in the lemma is bounded by another exponential function of n then a random m-factor of f can almost always be expressed as a modulo p sum of functions each of which do not depend on an m^{-mt} fraction of their variables.

Notation 12 *Let us denote by $w(x)$ the weight of a zero-one vector x, i.e., the number of ones in x. For a symmetric Boolean function f on n variables and integer $1 \leq i \leq n$ let us denote by $f(i)$ the value of f on inputs of weight i.*

Lemma 13 *Let p be a prime. If f is a symmetric Boolean function on p^k variables with $f(0) \neq f(p^k)$ then f is not p-simple.*

Proof: Notice that

$$\sum_{x \in \{0,1\}^n} (-1)^{w(x)} f(x) \equiv 0 \pmod{p}$$

for p-simple functions f. The left hand side is zero for functions not depending on one of the input variables, thus it is divisible by p for a modulo p sum of such functions.

For a symmetric function on $n = p^k$ variables the left hand side of the last equation is

$$\sum_{i=0}^{n} (-1)^i \binom{n}{i} f(i) \equiv f(0) - f(n) \pmod{p},$$

since p divides $\binom{p^k}{i}$ unless $i = 0$ or $i = p^k$. Thus $f(0) \neq f(n)$ implies that f has full p-degree as claimed. □

Theorem 14 *Let p be a prime, m, t, k, and S positive integers and $1 \geq \varepsilon > 0$. Suppose the symmetric Boolean function f on n variables is the modulo p sum of S of the functions f_i, where each of the f_i computed in an arbitrary way from t of the functions f_{ij} and from $(1-\varepsilon)n$ of the input variables. Each of the functions f_{ij} is in turn a modulo m linear combination of the input variables. Here the functions f_i output modulo p values while f_{ij} output modulo m values. Suppose f is not equal to any MOD_{mp^k} gate. Then $S > c^{p^k}$ for a constant $c > 1$ depending only on m, t, and ε.*

Proof: Since f is not a MOD_{mp^k} gate, there exist numbers $0 \leq i < i + mp^k = j \leq n$ such that $f(i) \neq f(j)$. Restrict the function f by assigning 0 to $n - j$ of its variables and assigning 1 to i of them. The resulting function f' is a symmetric function of its mp^k variables satisfying $f'(0) \neq f'(mp^k)$. Notice that the restriction does not increase the size of the circuit computing the function. The unique m-factor of f' is a symmetric function f'' on p^k variables satisfying $f''(0) \neq f''(p^k)$. By Lemma 13 f'' is not p-simple. Thus Lemma 11 gives the claimed bound on S. □

We are ready now to prove Theorem 3.

Proof: (Theorem 3) We apply Lemma 9 to get rid of the MOD_{p^k} gates and get a $(\sum_p, \mathrm{AND}_t, \mathrm{MOD}_m)$ circuit for our symmetric function. The size of the circuit blows up polynomially, i.e., it is bounded by S^b, where b and t depend on p, m, k, and ℓ. Then Theorem 14 bounds S. Notice that we did not use the feature of Theorem 14 that the middle gates can directly depend on many input variables.

The statement on the specific functions follows from the observation that every function mentioned there satisfies that it is not of the form $\mathrm{MOD}^A_{mp^j}$ unless $mp^j > n$. □

The following lemma nicely complements Theorem 14.

Lemma 15 *Consider the Boolean function $f(x_1, x_2, \ldots, x_n) = \mathrm{MOD}^A_{mp^k}(x_1, x_2, \ldots, x_n)$. If m is not a power of the prime p then f can be computed by a $(\sum_p, \mathrm{AND}_2, \mathrm{MOD}_m)$ circuit of size at most $(mn)^{2p^{k'}}$, where $p^{k'}$ is the largest power of p dividing mp^k.*

Notice that the assumption that m is not a power of p is necessary. Otherwise, if $m = p^\ell$, arbitrary size constant depth circuits of constant fan-in AND and arbitrary MOD_p and MOD_m gates could only compute Boolean functions expressible as constant degree modulo p polynomials, and that constant degree does not depend on k. Consequently, it cannot compute f, which is a degree-(p^k-1) polynomial.

Proof: Suppose first that all elements of the 1-set A are congruent to a single number a modulo m. There is a degree $p^{k'}-1$ polynomial on the input computing $\mathrm{MOD}^A_{p^{k'}}$ modulo p (Lemma 7). This polynomial can be implemented by a modulo p sum of AND gates of at most $p^{k'}-1$ variables. The number of AND gates is bounded by $n^{p^{k'}-1}$. Let q be prime factor of m different from p and stick a redundant $\mathrm{MOD}^{\{1\}}_q$ gate above each AND gate. Apply the Degree Decreasing Lemma (Lemma 16) to replace each AND gate by a collection of at most $(2q)^{p^{k'}-2}$ MOD_q gates summing to the same value modulo p. First replace each MOD_q gate by a MOD_m gate computing the same function then replace each MOD_m gate G by the AND of G and the $\mathrm{MOD}^{\{a\}}_m$ gate on all the inputs. The resulting circuit computes the AND of the $\mathrm{MOD}^{\{a\}}_m$ and the $\mathrm{MOD}^A_{p^{k'}}$ functions, thus it computes the $\mathrm{MOD}^A_{mp^{k'}}$ function as desired.

To remove our assumption on A notice that every set A can be decomposed into m sets A_i satisfying this assumption. The equation $\mathrm{MOD}^A_{mp^k} = \sum_i \mathrm{MOD}^{A_i}_{mp^k}$ proves the lemma. □

Consider the smallest $(\sum_p, \mathrm{AND}_t, \mathrm{MOD}_m)$ circuit computing the function $\mathrm{MOD}^{\{0\}}_{mp^j}$ and notice that the lower bound on the circuit size for this function in Theorem 14 is c^{p^j} while the upper bound in Lemma 15 is n^{cp^j}. The gap is too wide to characterize polynomial size circuits, but we can characterize quasipolynomial size circuits as in Theorem 4.

Proof: (Theorem 4) Apply Lemma 9 as in Theorem 3 to eliminate the MOD_{p^k} gates. Use Theorem 14 and Lemma 15 to get the two sides of the characterization. □

3.3 The Degree Decreasing Lemma

Lemma 16 exploits a surprising property of (MOD_s, MOD_m)-circuits, which lacks in (MOD_p, MOD_p) circuits, since constant-depth circuits with MOD_p gates are capable only to compute a constant degree polynomial of the inputs, and this constant depends on the depth, and not on the size. Here we generalize the original version of [Gro98] for non-prime moduli as well.

Lemma 16 *(Degree Decreasing Lemma) Let p be a prime, and $s, m > 1$ be integers, satisfying* $\gcd(s,p) = \gcd(s,m) = 1$. *Let x_1, x_2, x_3 be variables of values from $\{0,1,\ldots,p-1\}$, $x'_1 \in \{0,1\}$. Then*

$$\mathrm{MOD}^A_p(x_1x_2 + x_3) \equiv H_0 + H_1 + \cdots + H_{p-1} + \beta \pmod{s}, \tag{2}$$

$$\mathrm{MOD}^A_m(x'_1x_2 + x_3) \equiv H'_0 + H'_1 + \beta' \pmod{s}, \tag{3}$$

where H_i abbreviates

$$H_i = \alpha \sum_{j=0}^{p-1} \mathrm{MOD}^A_p(ix_2 + x_3 + j(x_1 + (p-i)))$$

for $i = 0,1,\ldots,p-1$, and where α is the multiplicative inverse of p modulo s: $\alpha p \equiv 1 \pmod{s}$, and β is a positive integer satisfying $\beta = -|A|(p-1)\alpha \bmod s$, and

where H'_i abbreviates

$$H'_i = \alpha' \sum_{j=0}^{m-1} \mathrm{MOD}_m^A(ix_2 + x_3 + j(x'_1 + (m-i)))$$

for $i = 0, 1$, and where α' is the multiplicative inverse of m modulo s: $\alpha' m \equiv 1 \pmod{s}$, and β' is a positive integer satisfying $\beta' = -|A|\alpha \bmod s$,

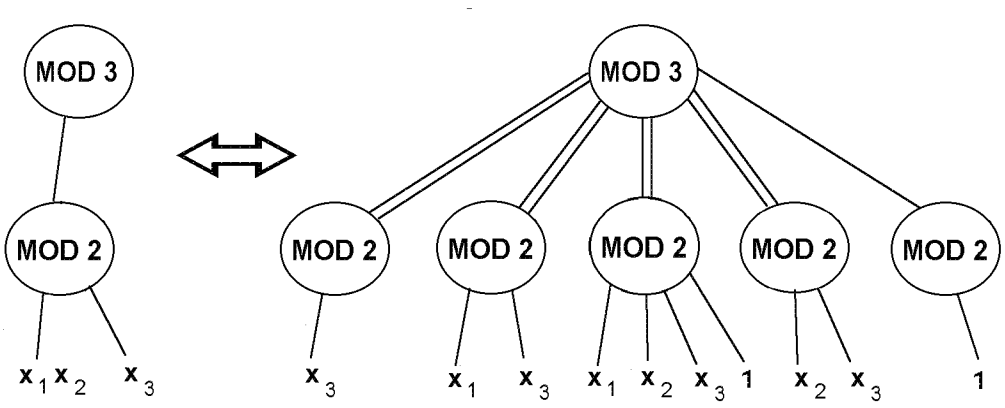

Figure 3: Degree decreasing in $(\mathrm{MOD}_3, \mathrm{MOD}_2^{\{1\}})$ case. On the left the input is a degree-2 polynomial, on the right linear polynomials.

Proof: Let $x_1 = k$ and let $0 \le i \le p-1$, $k \ne i$. Then

$$H_k = \alpha \sum_{j=0}^{p-1} \mathrm{MOD}_p^A(kx_2 + x_3) = \alpha p \mathrm{MOD}_p^A(kx_2 + x_3) \equiv \mathrm{MOD}_p^A(x_1x_2 + x_3) \pmod{s},$$

and

$$H_i = \alpha \sum_{j=0}^{p-1} \mathrm{MOD}_p^A(ix_2 + x_3 + j(k-i)) = \alpha|A|,$$

since for any fixed x_2, x_3, i, k expression $kx_2 + x_3 + j(k-i)$ takes on every value exactly once modulo p while $j = 0, 1, \ldots, p-1$; so $\mathrm{MOD}_p^A(ix_2 + x_3 + j(k-i))$ equals to 1 exactly $|A|$ times. Consequently,

$$H_0 + H_1 + \cdots + H_{p-1} + \beta \equiv \mathrm{MOD}_p^A(x_1x_2 + x_3) + (p-1)\alpha|A| + \beta \equiv \mathrm{MOD}_p^A(x_1x_2 + x_3) \pmod{s}.$$

Similarly, let $x'_1 = k \in \{0, 1\}$ and let $i \in \{0, 1\}$, $k \ne i$. Then

$$H'_k = \alpha' \sum_{j=0}^{m-1} \mathrm{MOD}_m^A(kx_2 + x_3) = \alpha' m \mathrm{MOD}_m^A(kx_2 + x_3) \equiv \mathrm{MOD}_p^A(x'_1x_2 + x_3) \pmod{s},$$

and

$$H'_i = \alpha' \sum_{j=0}^{m-1} \mathrm{MOD}_m^A(ix_2 + x_3 + j(k-i)) = \alpha'|A|,$$

since for any fixed x_2, x_3, i, k, for $i \neq k$ $|i-k| = 1$, so expression $kx_2 + x_3 + j(k-i)$ takes on every value exactly once modulo m while $j = 0, 1, \ldots, m-1$; so $\mathrm{MOD}_m^A(ix_2 + x_3 + j(k-i))$ equals to 1 exactly $|A|$ times. Consequently,

$$H_0' + H_1' + \beta' \equiv \mathrm{MOD}_m^A(x_1'x_2 + x_3) + \alpha'|A| + \beta' \equiv \mathrm{MOD}_p^A(x_1'x_2 + x_3) \pmod{s}.$$

□

3.4 Random Restriction

The Constant Degree Hypothesis of [BST90] states that any $(\sum_p, \mathrm{MOD}_m, \mathrm{AND}_d)$ circuit computing AND has super-polynomial size if p is a prime and m and d are constants. We make progress toward this goal by proving that AND requires super-polynomial size circuits of this type if each MOD_m gate has fan-in $o(n^2/\log n)$. For the stronger version of this statement see Theorem 6.

Definition 17 *We say that a* MOD_m^A *gate is* ε-linear, *if there is a subset* H *of the input-variables, containing at most the* ε*-fraction of all variables, such that the gate does not relate two input variables outside* H*, i.e., the input of the* MOD_m^A *gate is linear in the variables outside* H *with coefficients that are arbitrary functions of the variables in* H*.*

We start with a simple application of the Degree Decreasing Lemma (Lemma 16).

Lemma 18 *Let* p *and* m *be relative prime integers and consider an* n *variable Boolean function* f *be computed by a* $(\mathrm{MOD}_m^B, \mathrm{AND})$ *circuit, where the top* MOD_m^B *gate is an* ε*-linear gate. Then* f *can be computed by a* $(\sum_p, \mathrm{MOD}_m)$ *circuit with* $(2m)^{|H|}$ MOD_m *gates.*

Proof: We use induction on $|H|$. In the $|H| = 0$ base case the AND gates have fan-in one, thus they can be removed.

As the AND gates implement multiplication of the 0-1 variables we can consider the input of the MOD_m gate a polynomial P of the input variables with all of its monomials having at most a single variable outside H. We may suppose that P is multi-linear. If $x_i \in H$ for some $1 \leq i \leq n$ we can write this input in the form $P = Qx_i + R$, where the polynomials Q and R do not depend on x_i and all their monomials contain at most a single variable outside H. We apply Lemma 16 to replace our MOD_m gate with the modulo p sum of $2m$ MOD_m gates. The inputs of these MOD_m gates are linear combinations of x_i, Q, and R. To finish the proof we apply the inductive hypothesis with $H \setminus \{x_i\}$ to replace each of these new MOD_m gates with the modulo p sum of $(2m)^{|H|-1}$ MOD_m gates on the input variables. □

Lemma 19 *Let prime* p *and positive integers* k*,* ℓ*, and* m *be fixed. Then there exist constants* $c > 1$ *and* $\varepsilon > 0$ *such that if a circuit* $((\mathrm{MOD}_{p^k}^A)^\ell, \mathrm{MOD}_m^B, \mathrm{AND})$ *computes* AND_n*, and every* MOD_m *gate is an* ε*-linear gate, then the size of the circuit is* $S > c^n$*.*

The condition on H_G in the lemma is equivalent to saying that the input of gate G is a linear combination of the input variables outside H_G with coefficients arbitrarily depending on variables in H_G.

The proof of this lemma is simpler for the case p not dividing m. We need Theorem 14 in its full generality for the the remaining case.

Proof: Suppose first that p does not divide m.

We apply Lemma 18 for the MOD_m gates. The resulting circuit computes a modulo p polynomial of degree less than $p^{k\ell}$ of the at most $S(2m)^{\varepsilon n}$ MOD_m gates (Lemma 9). The size is therefore at most $(S(2m)^{\varepsilon n})^{p^{k\ell}}$. But Theorem 3 claims an exponential lower bound on this size, thus for a small enough ε, size S must be exponential in n.

In the general case where p may divide m we decompose $m = p^a m_0$ where p does not divide m_0. We decompose each MOD_m^B gate into the sum of at most m gates $\mathrm{MOD}_m^{\{0\}}$ gates that realize that the latter are computed as the AND of the corresponding MOD_{m_0} and MOD_p^k gates. (We used similar decomposotion in the proof of Lemma 15.) We have increased the size of the circuit by a factor of at most $2m$ so far. We apply Lemma 18 to the MOD_{m_0} gates. This increases the size by a factor of at most $(2m)^{\varepsilon n}$. The resulting circuit has MOD_{m_0} and AND gates at the bottom level and MOD_{p^a}, MOD_{p^k}, $\sum_p$, and AND_2 gates everywhere else. As the last two types can be replaced with MOD_p gates we can apply Lemma 9. We get a three level circuit computing AND_n with a $\sum_p$ gate on top, AND_t gates in the middle (with a constant t depending on m, p, k, and ℓ). The bottom gates are MOD_m and AND gates. Notice that the number S_2 of the gates in the middle level is at most S_1^t, where S_1 is the number of gates on the bottom level, and $S_1 \leq (2m)^{\varepsilon n+1} S$.

The fanin of these bottom AND gates are bounded by $\varepsilon n+1$. We choose $\varepsilon < 1/4t$. Merging the bottom AND gates with the middle AND gates one gets that AND_n is the modulo p sum of AND functions on at most $n/2$ inputs and at most t MOD_m gates. Applying Theorem 14 one gets that $S_2 > c^n$ with some $c > 1$ depending on p, m, k, and ℓ. Thus $S^t > c^n/(2m)^{t(\varepsilon n+1)}$ proving an exponential lower bound on S if $c > (2m)^{\varepsilon t}$. □

Now we turn to prove Theorem 6. It is a special case of the following result proving an optimal tradeoff between size and the new measure of the maximal number of related pairs.

Theorem 20 *Let p be a prime and m, k, and ℓ positive integers. Suppose that a $((\mathrm{MOD}_{p^k}^B)^\ell, \mathrm{MOD}_m^A, \mathrm{AND})$ circuit computes AND_n. If each MOD_m gate in the circuit relates at most $X \geq n$ pairs of input variables then the size of the circuit is at least $c_0^{n^2/X}$, with a constant $c_0 > 1$ depending on p, m, k, and ℓ.*

Proof: We fix the values c and ε claimed in Lemma 19. We take a restriction on the circuit by leaving a variable unrestricted with probability $P = \varepsilon n/(22X)$ independently for each of the variables. We assign 1 to the rest of the variables. Clearly, the restricted circuit computes the AND of the remaining variables.

With probability of at least $1/2$, the number of the remaining variables is at least $n_0 = Pn/2 = \varepsilon n^2/(44X)$.

The pairs related by a MOD_m gate in the restricted circuit are simply those pairs related by this gate in the unrestricted circuit that both remain unrestricted.

The expected number of pairs related by a single gate in the restricted circuit is at most $XP^2 = \varepsilon n_0/11$. Unfortunately, the deviation can be large, it is easy to construct n gates relating $n-1$ pairs each, such that *any* restriction to n' variables has a gate relating $n'-1$ pairs. Thus it is important, that for Lemma 19 we need not bound the number of related pairs, only the size of a set covering each.

Lemma 19 proves the Theorem if there is a restriction leaving n_0 variables unrestricted and making every MOD_m gate G admit a set of H_G of size at most εn_0 containing at least one of every pair related by G in the restricted circuit.

Let us bound the probability that this is not the case for a fixed MOD_m gate G. Take a maximal matching on pairs of unrestricted variables that are related by G. Clearly the set H of the endpoints of these edges satisfies that no pair is related outside H. Thus it suffices to bound the probability that $|H| \geq \varepsilon n_0$. This is exactly the event that all the variables involved in the $j = \lceil \varepsilon n_0/2 \rceil$ pairs of the matching remain unrestricted. Consequently, the probability that there exists an H, $|H| \geq \varepsilon n_0$ is at most the number of choices for the matching multiplied by P^{2j} which is the probability that the variables remain unrestricted. Hence if $S\binom{X}{j}P^{2j} < 1/2$, where S is the size of our circuit then Lemma 19 proves our Theorem. The alternative is $S \geq (2\binom{X}{j}P^{2j})^{-1} \geq \left(\frac{j}{eXP^2}\right)^j$ proving the same bound. □

Next we show that the bound in Theorem 20 is tight.

Theorem 21 *If m is not a power of the prime p, and $X > 0$ is arbitrary, then the n variable* AND *function is computable by a $(\sum_p, \mathrm{MOD}_m, \mathrm{AND})$ circuit of size $(2m)^{n^2/(2X)}$ such that the total number of pairs of variables related by any MOD_m gate in the circuit is at most X.*

Proof: Compute AND of the variables in two levels with AND gates, first computing the AND of $\lceil n^2/(2X) \rceil$ classes of at most $\lceil 2X/n \rceil$ variables each. Then place a $\mathrm{MOD}_m^{\{1\}}$ gate of fan-in 1 onto the top. Apply Lemma 18 to replace the top two levels by the modulo p sum of $(2m)^{n^2/(2X)}$ MOD_m gates. The inputs of these new gates are linear combinations of the outputs of the gates computing AND for a single class.

Note that Lemma 18 works only if m is not a multiple of p. Otherwise use that MOD_m gates can simulate MOD_q gates if q divides m. □

We remark that the proofs of Lemma 19 and Theorem 20 use Theorem 14 for the lower bound, so they apply to circuits computing OR or MOD_r with r not dividing mp^s, not just for AND. The upper bound in Theorem 21 can also be applied to these functions.

Acknowledgment. We are grateful to Zoltán Király for helpful discussions. This work was supported in part by grants OTKA F014919, FKFP 0835, AKP 97-56 2,1.

References

[Ajt83] Miklós Ajtai. $\sum_1^1$ formulae on finite structures. *Annals of Pure and Applied Logic*, 24:1–48, 1983.

[BST90] David A. Mix Barrington, Howard Straubing, and Denis Thérien. Non-uniform automata over groups. *Information and Computation*, 89:109–132, 1990.

[BT91] Richard Beigel and Jun Tarui. On ACC. In *Proc. 32nd Ann. IEEE Symp. Found. Comput. Sci.*, pages 783–792, 1991.

[Cau96] Herve Caussinus. A note on a theorem of Barrington, Straubing and Thérien. *Information Processing Letters*, 58:31–33, 1996.

[CG85] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. In *Proc. 26th Ann. IEEE Symp. Found. Comput. Sci.*, pages 429–442, 1985. Appeared also in SIAM J. Comput. Vol. 17, (1988).

[FSS84] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits and the polynomial time hierarchy. *Math. Systems Theory*, 17:13–27, 1984.

[Gro98] Vince Grolmusz. A degree-decreasing lemma for (MOD q, MOD p) circuits. In *Proc. ICALP'98, Aalborg, Denmark, LNCS ????*, page ???, 1998.

[Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *Proc. 18th Ann. ACM Symp. Theor. Comput.*, pages 6–20, 1986.

[KM91] Jeff Kahn and Roy Meshulam. On mod p transversals. *Combinatorica*, 10(1):17–22, 1991.

[KP94] Matthias Krause and Pavel Pudlák. On the computational power of depth 2 circuits with threshold and modulo gates. In *Proc. 26th Ann. ACM Symp. Theor. Comput.*, 1994.

[KW95] Matthias Krause and Stephan Waack. Variation ranks of communication matrices and lower bounds for depth-two circuits having nearly symmetric gates with unbounded fan-in. *Mathematical Systems Theory*, 28(6):553–564, November/December 1995.

[Raz85] Alexander Razborov. Lower bounds for the monotone complexity of some Boolean functions. *Sov. Math. Dokl.*, 31:354–357, 1985.

[Raz87] Alexander Razborov. Lower bounds on the size of bounded depth networks over a complete basis with logical addition, (in Russian). *Mat. Zametki*, 41:598–607, 1987.

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proc. 19th Ann. ACM Symp. Theor. Comput.*, pages 77–82, 1987.

[vL90] J. van Leeuwen, editor. *Handbook of Theoretical Computer Science*, volume A, chapter 14. The complexity of finite functions, by R.B. Boppana and M. Sipser. Elsevier-MIT Press, 1990.

[Yao85] Andrew C. Yao. Separating the polynomial-time hierarchy by oracles. In *Proc. 26th Ann. IEEE Symp. Found. Comput. Sci.*, pages 1–10, 1985.

[Yao90] Andrew C. Yao. On ACC and threshold circuits. In *Proc. 31st Ann. IEEE Symp. Found. Comput. Sci.*, pages 619–627, 1990.

[YP94] P.Y. Yan and Ian Parberry. Exponential size lower bounds for some depth three circuits. *Information and Computation*, 112:117–130, 1994.

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'