Complex Fourier Technique for Lower Bounds on the Mod-m Degree [1]

Frederic Green
Department of Mathematics and Computer Science
Clark University
Worcester, Massachusetts 01610
fgreen@black.clarku.edu

[1] Revised and expanded version of "Lower Bounds for Depth-Three Circuits with Equals and Mod-Gates," in *12th Annual Symposium on Theoretical Aspects of Computer Science*, Springer-Verlag (1995) 71-82.

Abstract

We say an integer polynomial p, on Boolean inputs, weakly m-represents a Boolean function f if p is non-constant and is zero (mod m), whenever f is zero. In this paper we prove that if a polynomial weakly m-represents the Mod_q function on n inputs, where q and m are relatively prime and m is otherwise arbitrary, then the degree of the polynomial is $\Omega(n)$. This generalizes previous results of Barrington, Beigel and Rudich (Computational Complexity **4** (1994), pp. 367-382) and Tsai (Structures 1993, pp. 96-101), which held only for constant or slowly growing m. In addition, the proof technique given here is quite different. We use a method (adapted from Barrington and Straubing, (Computational Complexity **4** (1994), pp. 325-338) in which the inputs are represented as complex q^{th} roots of unity. In this representation it is possible to compute the Fourier transform using some elementary properties of the algebraic integers. As a corollary of the main theorem and the proof of Toda's theorem, if q, p are distinct primes, any depth-three circuit which computes the Mod_q function, and consists of an exact threshold gate at the output, Mod_p-gates at the next level, and AND-gates of polylog fan-in at the inputs, must be of exponential size. We also consider the question of how well circuits consisting of one exact gate over ACC(p)-type circuits (where p is an odd prime) can approximate parity. It is shown that such circuits must have exponential size in order to agree with parity for more than $1/2 + o(1)$ of the inputs.

1 Introduction

One of the challenges facing circuit complexity is to prove lower bounds on bounded-depth circuits with Mod_m gates (that determine if the sum of the inputs is not divisible by the number m). If m is prime, exponential lower bounds are known for general bounded-depth circuits including AND's and OR's in addition to the MOD's [Sm]. Unfortunately the proof techniques for these results make essential use of the fact that $\mathbf{Z}/m\mathbf{Z}$ is a field if m is prime, and the situation for composite m (in which case $\mathbf{Z}/m\mathbf{Z}$ is a ring which is not even an integral domain) has proved to be more difficult. Algebraic or combinatorial techniques other than those used in [Sm] or [Raz] appear to be necessary.

Recently, progress has been made in the simple but important case of a single Mod_m-gate over small AND's [BBR], [Tsai]. This paper, for the most part, continues this line of investigation. Since this computational model is really based on polynomials, we formulate the problem accordingly. Let $p : \{0,1\}^n \to \mathbf{Z}$ be an integer polynomial, and $q, m \geq 2$ be relatively prime natural numbers. What is the minimum degree of p such that for all $(x_1, ..., x_n) \in \{0,1\}^n$, $p(x_1, ..., x_n) \equiv 0 \ \ (m)$ iff $\sum_{i=1}^n x_i$ is divisible by q? This question was studied by Barrington, Beigel and Rudich [BBR], and, in subsequent work, by [Tsai]. The strongest lower bound, obtained by Tsai, states that the polynomial must have degree linear in n. Analysis of the proofs of these previous results shows that they are valid only if m is a constant or a slowly growing function of n. But what if m is some arbitrary function of n, where the only constraint on m is that it is prime to q? This question has a number of motivations, which are discussed below. It is also of some intrinsic interest, since allowing m to vary in this way gives a considerably stronger computational model. Intuitively it would make sense that the degree lower bound has nothing to do with m's dependence on n, and is only due to the relative primality of q and m. However, it is not clear how the techniques of [BBR] and [Tsai] can be adapted to handle this more general case. In this paper we introduce an alternative technique which can.

Here we therefore consider a generalization of the problems studied in [BBR], [Tsai]. The most significant lower bounds will be for the *modulo* functions,

$$\mathrm{Mod}_{m,r}(x_1, ..., x_n) = \begin{cases} 0 & \text{if } \sum_{i=1}^n x_i \equiv r \ (m) \\ 1 & \text{otherwise,} \end{cases}$$

or their negations (for example, a Mod_m gate computes the function $\mathrm{Mod}_{m,0}$). Let f be any Boolean function on n inputs. We say an integer polynomial p on n Boolean inputs *m-represents* f if for all input settings of $x_1, ..., x_n$, $f(x_1, ..., x_n) = 0$ iff $p(x_1, ..., x_n) \equiv 0 \ (m)$. The *$\mathrm{Mod}_m$-degree* of f is the smallest degree of a polynomial which m-represents f. The *Mod-degree* of f is the smallest Mod_m-degree of f for any m. Let f_q be any one of the functions $\mathrm{Mod}_{q,r}$ or $\neg\mathrm{Mod}_{q,r}$. Define the *relatively prime Mod-degree* (or *rp-Mod-degree*, for short) of f_q as the smallest Mod_m-degree of f_q for any m which is relatively prime to q. Some functions have large Mod-degrees (for example, the AND function; see Proposition 9 below). Trivially, the Mod-degree of any f_q is 1. Our main theorem is,

Theorem 1 *For any $n \in \mathbf{N}$, the rp-Mod-degree of any of the functions* $\mathrm{Mod}_{q,r}(x_1, \ldots, x_n)$ *or* $\neg\mathrm{Mod}_{q,r}(x_1, \ldots, x_n)$ *is at least* $\lfloor \frac{n}{2(q-1)} \rfloor$.

Actually a stronger result is proved. Introducing terminology analogous to that of [ABFR], say a polynomial *weakly* m-represents a Boolean function f if p is not a constant function and $f(x_1, ..., x_n) = 0 \Longrightarrow p(x_1, ..., x_n) \equiv 0\ (m)$. The *weak Mod$_m$-degree* and the *weak rp-Mod-degree* are defined in the obvious ways. Theorem 5 in section 3 of this paper says that the weak rp-Mod-degree of the modulo functions is $\Omega(n)$. As a by-product (see Corollary 6), we present a new proof that, if m is independent of n and is not a prime power, the minimal degree of a polynomial which m-represents the $\neg\mathrm{Mod}_m$ function is $\Omega(n)$, a result originally obtained in [Tsai]. Another application of this stronger version (in the special case of $q = 2$) is described below.

The proof strategy is substantially different than those used in [BBR] and [Tsai]. We make use of the representation of Boolean inputs by complex roots of unity, introduced by Barrington and Straubing [BS]. There the goal was to show that the *sign* of p could not agree with $\mathrm{Mod}_{q,r}$. Related results, via similar techniques, for lower bounds on circuits consisting of one threshold over Mod_q gates, have been obtained recently by Krause and Pudlák [KP]. As in this previous work, the formulation in terms of roots of unity permits the application of the powerful technique of spectral analysis. The main difference between the setting here and that of [BS] and [KP] is that here we work in a finite ring (i.e., $\mathbf{Z}/m\mathbf{Z}$) rather than an infinite field (i.e., the complex numbers). As we will see, this requires the explicit use of some elementary properties of rings of algebraic integers.

In [BS], it is shown that the sign of a low-degree polynomial can't agree with $\mathrm{Mod}_{q,r}$ even for a constant fraction of the input settings. We prove a similar result that, for odd prime p, a low-degree polynomial cannot p^k-represent the parity function for more than a fraction $1/2 + o(1)$ of all input settings (Theorem 12, in section 5). Strong lower bounds on approximately m-representing parity for any odd m, or on approximately m-representing Mod_q for any relatively prime q and m, remain open.

The topic of this paper was originally motivated by some questions about depth-three circuits involving both threshold and mod gates. As is well known, threshold and parity gates are a potent combination. Inspired by Toda's theorem [Tod] that $\mathrm{PP}^{\mathrm{PH}} \subseteq \mathrm{PP}^{\oplus\mathrm{P}}$, Allender [Al] showed that any quasi-polynomial size "generalized perceptron," consisting of a threshold gate over $\mathrm{AC}^{(0)}$-type circuits, can be simulated by a quasi-polynomial-size depth-three circuit consisting of a threshold gate at the root, parity gates at the next level and AND gates of polylog fan-in connected to the inputs (which we refer to here as "threshold-of-$\oplus^+$" circuits; see section 2 for notation). On the negative side, drawing on work of Hajnal et al. [HMPST] and Boppana and Håstad [Has], it was shown in [Gre 91] that small generalized perceptrons cannot compute the parity function (in terms of complexity classes, that there is an oracle A such that $\oplus \mathrm{P}^A \not\subseteq \mathrm{PP}^{\mathrm{PH}^A}$). Improvements and generalizations have appeared in [Bei 92], [BRS], [BS]. However, as of this writing, no non-polynomial lower bounds are known for the more powerful threshold-of-$\oplus^+$ circuits (and hence no oracle is known separating $\mathrm{PP}^{\oplus \mathrm{P}}$ from PSPACE). In fact the best one can do, by Smolensky's theorem [Sm], is $\Omega(n^{1/2-o(1)})$ for threshold-of-Mod_q^+ circuits computing parity (q an odd prime). Although this problem has not received as much press as the ACC problem [Bar], it shares at least some of its difficulty.

Some progress has been made in a restricted version of the problem in [CGT], in which the functions computed by the parity sub-circuits are symmetric. Here we consider a different kind of restriction: the parity sub-circuits are general, but the top gate is an "exact threshold" or "exact" gate (the resulting circuits are called "exact-of-$\oplus^+$" circuits). In terms of Turing machine complexity classes, we thus consider the relativized power of classes such as $\mathrm{C}_=\mathrm{P}^{\oplus \mathrm{P}}$. While the resulting problems appear to be considerably simpler than those associated with $\mathrm{PP}^{\oplus \mathrm{P}}$, it is hoped that some insight will be gained towards attacking the more general problem.

The natural conjecture is that if q and p are distinct primes, an exact-of-Mod_p^+ circuit which computes the Mod_q function would have to be of exponential size. We prove this in section 4. By the proof of the second part of Toda's theorem (see Theorem 7), this problem easily reduces to Theorem 1.

How well can an exact-of-Mod_p^+ circuit *approximate* Mod_q? We can answer this question for $q = 2$, and indeed, building on the main result and extending some known techniques, obtain an improvement of part of Smolensky's theorem about ACC(p)-type circuits for odd prime p. We find that exact-of-ACC(p) circuits can equal the parity function for at most $1/2 + o(1)$ of the inputs, unless they are of exponential size. This is a consequence of Theorem 12 (mentioned above) and one of Smolensky's lemmas, and its proof is also given in section 5.

In section 6, we show how these results translate to oracle separations of polynomial-time Turing machine counting classes.

2 Preliminaries and Notation

It is assumed the reader has some familiarity with circuit complexity and polynomial-time complexity classes as in, e.g., [BDG]. More detailed background can be found in [Str] and [Bei 93].

Following the convention of [GKRST], if G is a Boolean gate, G^+ denotes a family of circuits with G at the root and AND gates of polylog fan-in at the input level. If G is a Boolean gate and $\mathcal{C}$ is a circuit class, G-of-$\mathcal{C}$ denotes the class of circuits with $\mathcal{C}$-type circuits serving as inputs to G-type gates. Examples of these notations were given in the introduction. It should be noted that the size of the $\mathcal{C}$ subcircuits is to be regarded as a function of the number of inputs to the global G-of-$\mathcal{C}$ circuit.

An *exact* gate over n Boolean inputs $x_1, ..., x_n$ with weights $w_i \in \mathbf{Z}$ $(1 \leq i \leq n)$ and threshold $t \in \mathbf{Z}$, returns 1 iff $\sum_{i=1}^n w_i x_i = t$. We interpret the term $w_i x_i$ as a sum of $|w_i|$ identical inputs x_i, so that the quantity $\sum_{i=1}^{n} |w_i|$ is called the *size* of the exact gate. Thus in an exact-of-$\mathcal{C}$ circuit, the size of the exact gate is the number of subcircuits of type $\mathcal{C}$, including repetitions. (This convention simplifies the statements of the results. If we called n the size and $\sum_{i=1}^{n} |w_i|$ the weight, the results would have to be expressed as tradeoffs between size and weight.) For the purposes of this paper, one may assume that the inputs x_i to an exact gate are never negated. Since $\overline{x}_i = 1 - x_i$, negations can be implemented by choosing w_i to be negative and adjusting the threshold t appropriately; the proof methods used here are independent of t.

A circuit family is of type ACC(p) if it is of bounded depth and each circuit in the family has AND, OR and Mod_p-gates.

If R is a ring and $m \in \mathbf{N}$, $a \equiv b\ (mR)$ means that for some $x \in R$, $a = b + mx$. When it is clear that $R = \mathbf{Z}$, we use the notation $a \equiv b\ (m)$. As usual, R/mR denotes the residue class ring of R modulo R-multiples of m.

The proof of the main theorem makes use of the ring of *algebraic integers*, which we denote[1] by $\overline{\mathbf{Z}}$. Most algebra or number theory texts discuss the subject (see, for example, [IR] (chapter 6) for a compact treatment, including proofs, of all the relevant notions). For the sake of completeness, the definition and the necessary facts that are used in this paper are included here. An algebraic integer is a complex root of an integer polynomial (i.e., one with integer coefficients) in which the leading coefficient is one. For example, the roots of the polynomial $z^n - 1$ (i.e., the complex n^{th} roots of unity) are algebraic integers. It is straight-forward (although not trivial) to prove that $\overline{\mathbf{Z}}$ is a ring under complex addition and multiplication. Crucial for this paper is the elementary property that an algebraic integer is a rational number if and only if it is a rational integer. Because of this fact, the notation $a \equiv b \pmod{m}$ means $a \equiv b\ \ (m\overline{\mathbf{Z}})$ as well as $a \equiv b\ \ (m\mathbf{Z})$ if a and b are integers. Nevertheless, we usually resort to the more precise notations $a \equiv b\ \ (m\overline{\mathbf{Z}})$ or $a \equiv b\ \ (m\mathbf{Z})$.

3 Lower Bounds for the rp-Mod-degree of Modulo Functions

We start by reviewing the formulation of [BS] to express circuit inputs as complex roots of unity. Let q be a natural number ≥ 2 and ζ denote the primitive complex root of unity $e^{\frac{2\pi i}{q}}$ (q will be fixed in this section). Let $D = \{1, \zeta, \zeta^2, ..., \zeta^{q-1}\}$. Note that

$$\sum_{y \in D} y^k = \begin{cases} q & \text{if } k \equiv 0\ \ (q) \\ 0 & \text{otherwise.} \end{cases}$$

[1]There is apparently no standard notation for all the algebraic integers. They constitute the *integral closure* of $\mathbf{Z}$ in the field of algebraic numbers (which in turn is the *algebraic* closure of the rationals). The notation here is chosen to suggest the integral closure of $\mathbf{Z}$.

We will consider polynomials over n variables $y_1, ..., y_n \in D$. Let V denote $\{0, ..., q-1\}$. Since $y^q = 1$ for any $y \in D$, a monomial will have the general form $\prod_{i \in S_{\mathbf{e}}} y_i^{e_i}$, where $\mathbf{e} \in V^n$, and by definition $S_{\mathbf{e}}$ denotes the set $\{i | e_i \neq 0\}$. Define $\prod_{i \in \emptyset} y_i^{e_i} = 1$. The integer $|S_{\mathbf{e}}|$ is called the *weight* of the monomial.

The first crucial observation to make here is that the sum of a monomial of non-zero weight less than n over all values of the y_i's subject to the constraint $\prod_{i=1}^n y_i = \zeta^r$ (for any integer r) is zero. Let "$\mathbf{y} : \prod y = \zeta^r$" denote that $\mathbf{y} \in D^n$ varies subject to the constraint $\prod_{i=1}^n y_i = \zeta^r$.

Lemma 2 *Let r, n be integers. Let e_i ($i \in \{1, ..., n\}$) be integers such that for each i, $0 \leq e_i \leq q-1$. Define $S_{\mathbf{e}}$ as above and suppose $|S_{\mathbf{e}}| < n$. Then*

$$\sum_{\mathbf{y}:\prod y = \zeta^r} \prod_{i \in S_{\mathbf{e}}} y_i^{e_i} = \begin{cases} q^{n-1} & \textit{if } S_{\mathbf{e}} = \emptyset \\ 0 & \textit{otherwise.} \end{cases}$$

Proof: If $S_{\mathbf{e}} = \emptyset$, the result is obvious. Suppose $S_{\mathbf{e}} \neq \emptyset$. Since $|S_{\mathbf{e}}| < n$, there is an $i' \in \{1, ..., n\} - S_{\mathbf{e}}$. The constraint $\prod_{i=1}^n y_i = \zeta^r$ is equivalent to,

$$y_{i'} = \zeta^r \cdot \prod_{i \neq i'} y^{-i}.$$

This determines $y_{i'}$ uniquely in terms of the other y_i's. We can then perform the sum over the y_i's ($i \neq i'$) independently. Hence for any $i \in S_{\mathbf{e}}$, the sum has a factor,

$$\sum_{y_i \in D} y_i^{e_i},$$

which is zero, since $e_i \not\equiv 0$ (q). ∎

For any integer r, it will be convenient to define the *r-inner product* $\langle f, g \rangle_r$ for functions f, g over D^n as,

$$\langle f, g \rangle_r = \frac{1}{q^{n-1}} \sum_{\mathbf{y}:\prod y = \zeta^r} \overline{f}(y_1, ..., y_n) g(y_1, ..., y_n).$$

Lemma 2 says that $\langle 1, \prod_{i \in S_{\mathbf{e}}} y_i^{e_i} \rangle_r$ is zero unless $S_{\mathbf{e}} = \emptyset$, in which case it is 1. More generally, the set of monomials of weight $< n/2$, under the r-inner product, form an orthonormal basis for polynomials of weight $< n/2$ over D^n.

Lemma 3 *Let $\mathbf{e}', \mathbf{e} \in V^n$ and $|S_{\mathbf{e}} \cup S_{\mathbf{e}'}| < n$. Then,*

$$\left\langle \prod_{i \in S_{\mathbf{e}'}} y_i^{e'_i}, \prod_{i \in S_{\mathbf{e}}} y_i^{e_i} \right\rangle_r = \begin{cases} 1 & \textit{if } \mathbf{e} = \mathbf{e}' \\ 0 & \textit{otherwise.} \end{cases}$$

Proof: Define,

$$S(\mathbf{e},\mathbf{e}') = \left\langle \prod_{i\in S_{\mathbf{e}'}} y_i^{e'_i}, \prod_{i\in S_{\mathbf{e}}} y_i^{e_i} \right\rangle_r .$$

The inner product can be written as,

$$S(\mathbf{e},\mathbf{e}') = \frac{1}{q^{n-1}} \sum_{\mathbf{y}:\prod y=\zeta^r} \prod_{i\in S_{\mathbf{e}'}-S_{\mathbf{e}}} y_i^{q-e'_i} \prod_{i\in S_{\mathbf{e}}-S_{\mathbf{e}'}} y_i^{e_i} \prod_{i\in S_{\mathbf{e}'}\cap S_{\mathbf{e}}} y_i^{e_i-e'_i}.$$

Because $|S_{\mathbf{e}} \cup S_{\mathbf{e}'}| < n$, we can apply Lemma 2. By that lemma, if $S_{\mathbf{e}}\Delta S_{\mathbf{e}'} \neq \emptyset$, $S(\mathbf{e},\mathbf{e}') = 0$. If $S_{\mathbf{e}}\Delta S_{\mathbf{e}'} = \emptyset$, $S(\mathbf{e},\mathbf{e}') = 0$ unless $\mathbf{e} = \mathbf{e}'$. If $\mathbf{e} = \mathbf{e}'$ then the sum is q^{n-1}, so $S(\mathbf{e},\mathbf{e}') = 1$. ∎

The main result concerns integer polynomials which take on values in the ring $\mathbf{Z}/m\mathbf{Z}$ for some m. Because we work with inputs that are roots of unity, we will find it necessary, in proving the main theorem, to consider *algebraic* integer polynomials with values in $\mathbf{Z}/m\mathbf{Z}$. For the moment, we consider the even more general case of polynomials with algebraic integer coefficients and values in $\overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$[2]. In the following lemma, we show, for such a polynomial $t : D^n \to \overline{\mathbf{Z}}$ of weight $< n/2$, that we can solve for the coefficients in terms of special values of t, namely, exactly those values of $t(y_1, ..., y_n)$ such that $\prod_{i=1}^n y_i = \zeta^r$, for any r. This set of coefficients (the $c_{\mathbf{e}}$'s in equation 1 below) is called the Fourier transform of t. Thus the lemma gives a particular method for computing the Fourier transform when the polynomial has weight $< n/2$. The idea of the proof of the main theorem is to use this lemma to show that the Fourier transform is zero if $t(y_1, ..., y_n) \equiv 0 \pmod{m}$ for the indicated values of $\mathbf{y}$.

Lemma 4 *Let m be a natural number relatively prime to q. Suppose $t : D^n \to \overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$ is a polynomial with algebraic integer coefficients and weight $< n/2$, i.e., of the form,*

$$t(y_1, ..., y_n) \equiv \sum_{\mathbf{e}\in V^n:|S_{\mathbf{e}}|<n/2} c_{\mathbf{e}} \prod_{i\in S_{\mathbf{e}}} y_i^{e_i} \quad (m\overline{\mathbf{Z}}) \tag{1}$$

where $c_{\mathbf{e}} \in \overline{\mathbf{Z}}$. Then for any natural number r and $\mathbf{e} \in V^n$ with $|S_{\mathbf{e}}| < n/2$,

$$c_{\mathbf{e}} \equiv \langle \prod_{i\in S_{\mathbf{e}}} y_i^{e_i}, t(y_1, ..., y_n)\rangle_r \quad (m\overline{\mathbf{Z}}). \tag{2}$$

Proof: Note that since q and m are relatively prime, q has an inverse in $\mathbf{Z}/m\mathbf{Z}$, and hence also in $\overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$. Hence the r-inner product is well-defined in $\overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$. Let $\mathbf{e}' \in V^n$ be such that $|S_{\mathbf{e}'}| < n/2$. Take the r-inner product of both sides of equation (1) with $\prod_{i\in S_{\mathbf{e}'}} y_i^{e'_i}$. We obtain,

$$\langle \prod_{i\in S_{\mathbf{e}'}} y_i^{e'_i}, t(y_1, ..., y_n)\rangle_r \equiv \sum_{\mathbf{e}\in V^n:|S_{\mathbf{e}}|<n/2} c_{\mathbf{e}} \langle \prod_{i\in S_{\mathbf{e}'}} y_i^{e'_i}, \prod_{i\in S_{\mathbf{e}}} y_i^{e_i}\rangle_r \quad (m\overline{\mathbf{Z}}).$$

[2]This is the simplest approach, although it should be pointed out that it is possible to work in the finite ring $\mathbf{Z}[\zeta]/m\mathbf{Z}[\zeta]$. See section 7 for a brief discussion.

Note $|S_{\mathbf{e}'} \cup S_{\mathbf{e}}| < n$. Apply Lemma 3 to the above equation to obtain equation (2). ∎

We now state and prove the main theorem. Theorem 1 is then immediate.

Theorem 5 *Let q and N be natural numbers. Let $p : \{0,1\}^N \to \mathbf{Z}$ be a polynomial over the Boolean variables $x_1, ..., x_N$ with integer coefficients, and let $m \in \mathbf{N}$ be such that q and m are relatively prime. Suppose that for some integer $0 \leq r \leq q-1$, for any $x_1, ..., x_N$ it holds that,*

$$\sum_{i=1}^{N} x_i \equiv r \quad (q) \Longrightarrow p(x_1, ..., x_N) \equiv 0 \quad (m),$$

and that p is not a constant function mod m. Then the degree of p is at least $\lfloor \frac{N}{2(q-1)} \rfloor$.

Proof: The result is trivial if $N < 2(q-1)$, so suppose wlog that $N \geq 2(q-1)$. Let d be the degree of the polynomial p. Suppose $d < \lfloor \frac{N}{2(q-1)} \rfloor$. Then we will show that, under the given hypotheses, p is always zero mod m.

Define ζ, D, V as in the discussion above. Exactly as in [BS], we restrict the input settings so that p can be regarded as a polynomial over inputs in D. For this purpose, assume wlog that $(q-1)|N$; other values of N can be reduced to this by restricting $< q-1$ of the inputs to 1 (this is the origin of the floor in the degree lower bound). Let $n = \frac{N}{q-1}$. Group the N inputs into n sets of size $q-1$, each of the form $\{x_{(q-1)i+1}, ..., x_{(q-1)i+q-1}\}$ with $0 \leq i \leq n-1$. Let $y_1, ..., y_n$ be n inputs each taking on values in D. Define the polynomials mentioned in [BS] (though not written down explicitly),

$$u_j(y) = \frac{1}{q} \sum_{i=0}^{q-1} \sum_{l=j+1}^{q} \zeta^{-il} y^i,$$

for $y \in D$ and $1 \leq j \leq q-1$. Writing $y = \zeta^s$, with $1 \leq s \leq q$, it is easy to verify that $u_j(y) = 0$ if $j \geq s$ and $u_j(y) = 1$ if $j < s$. Thus $u_1(y)u_2(y)...u_{q-1}(y)$, regarded as a string of bits, is $1^{s-1}0^{q-s}$. Restricting the input settings appropriately, we can then encode y_i as the string $x_{(q-1)i+1}x_{(q-1)i+2}...x_{(q-1)i+q-1}$ by setting

$$x_{(q-1)i+j} = u_j(y_{i+1})$$

for $0 \leq i \leq n-1$ and $1 \leq j \leq q-1$. Writing $y_i = \zeta^{s_i}$ for $1 \leq i \leq n$, note that,

$$\sum_{j=1}^{q-1} x_{(q-1)i+j} = s_i - 1,$$

and therefore,

$$\sum_{i=1}^{N} x_i \equiv r \ (q) \ \text{ iff } \ \sum_{i=1}^{n} (s_i - 1) \equiv r \ (q) \ \text{ iff } \ \prod_{i=1}^{n} y_i = \zeta^{n+r}. \tag{3}$$

Up to this point, the argument has followed the one given in [BS]. However we must now remember that we are working in $\mathbf{Z}/m\mathbf{Z}$, while the polynomials u_j have coefficients which are not necessarily integers. Therefore instead of computing the polynomial $p(x_1, ..., x_N)$ mod $m\mathbf{Z}$, we compute it mod $m\overline{\mathbf{Z}}$. Since $p(x_1, ..., x_N) \equiv 0 \ \ (m\mathbf{Z})$ iff $p(x_1, ..., x_N) \equiv 0 \ \ (m\overline{\mathbf{Z}})$, we lose no generality in doing this. Recall from the proof of Lemma 4 that q has an inverse in $\overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$. Thus, working in $\overline{\mathbf{Z}}/m\overline{\mathbf{Z}}$, we regard u_j as a polynomial of degree at most $q-1$ (and weight at most 1) with algebraic integer coefficients. Using the u_j's we can define a polynomial $t : D^n \to \mathbf{Z}$ with coefficients in $\overline{\mathbf{Z}}$ which agrees with p when the x_i's encode y_i's according to the above scheme:

$$t(y_1, ..., y_n) \equiv p(u_1(y_1), ..., u_{q-1}(y_1), ..., u_1(y_n), ..., u_{q-1}(y_n)) \ \ (m\overline{\mathbf{Z}}).$$

Note that by the construction of t and relation (3),

$$\prod_{i=1}^{n} y_i = \zeta^{n+r} \Longrightarrow t(y_1, ..., y_n) \equiv 0 \ \ (m\overline{\mathbf{Z}}). \tag{4}$$

It is easy to see that the *weight* of any monomial in t is at most d, the degree of p. Note that $d < \lfloor \frac{N}{2(q-1)} \rfloor = \lfloor \frac{n}{2} \rfloor$ and hence $d < \frac{n}{2}$. Thus t fulfills the requirements of Lemma 4. We solve for the coefficients $c_{\mathbf{e}}$ of t by that lemma:

$$c_{\mathbf{e}} \equiv \langle \prod_{i \in S_{\mathbf{e}}} y_i^{e_i}, t(y_1, ..., y_n) \rangle_{n+r} \ \ (m\overline{\mathbf{Z}}).$$

By implication (4), every term on the right hand side is zero. Hence for any $\mathbf{e}$, $c_{\mathbf{e}} \equiv 0 \ \ (m\overline{\mathbf{Z}})$, from which we conclude that for any $y_1, ..., y_n \in D$, $t(y_1, ..., y_n) \equiv 0 \ \ (m\overline{\mathbf{Z}})$ and therefore, since $t(y_1, ..., y_n)$ is an integer, $t(y_1, ..., y_n) \equiv 0 \ \ (m\mathbf{Z})$. This immediately implies that $p(x_1, ..., x_N) \equiv 0 \ \ (m)$ for those q^n input settings of the x_i's that encode y_i's. By re-labeling the x_i's and repeating the argument for each re-labeling we conclude that $p(x_1, ..., x_N) \equiv 0 \ \ (m)$ for all 2^N settings of the x_i's. ∎

Corollary 6 [Tsai] *Suppose m is fixed and is not a prime power, and let q^e_{max} be the largest prime power dividing m. Then any polynomial which m-represents $\neg\mathrm{Mod}_m$ must have degree at least $\lfloor \frac{n}{2q^e_{max}} \rfloor$.*

Proof: Let p be a polynomial of degree less than $\lfloor \frac{n}{2q^e_{max}} \rfloor$ which m-represents $\neg\mathrm{Mod}_m$. For any prime factor s of m, let $e(s) = ord_s(m)$, i.e., $e(s)$ is the greatest integer such that $s^{e(s)}$ divides m. Let r, s be two distinct prime factors of m. Since $\sum_{i=1}^{n} x_i \not\equiv 0 \ (m) \Longrightarrow p(x_1, ..., x_n) \equiv 0 \ (m)$, it follows that,

$$\sum_{i=1}^{n} x_i \not\equiv 0 \ (s^{e(s)}) \Longrightarrow p(x_1, ..., x_n) \equiv 0 \ (r^{e(r)}).$$

Since the degree of p is less than $\lfloor \frac{n}{2s^{e(s)}} \rfloor$, it follows from Theorem 5 that p is the zero function mod $r^{e(r)}$. As this holds for any prime divisor r of m, we conclude that p is identically zero mod m, which is a contradiction. ∎

4 Simulations and Lower Bounds for Exact-of-Mod_p^+ Circuits

Let p be a prime. The main result of this section is that the negation of any exact-of-Mod_p^+ circuit of size 2^{n^ϵ} (with $\epsilon < 1$) can be p^s-represented, for some $s \in \mathbf{N}$, by a polynomial of sub-linear degree. By Theorem 1, such a polynomial cannot p^s-represent $\mathrm{Mod}_{q,r}$ or $\neg\mathrm{Mod}_{q,r}$ if q is a prime $\neq p$. Hence exact-of-Mod_p^+ circuits of size 2^{n^ϵ} cannot compute these functions.

Theorem 7 *Fix any real number $\epsilon < 1$. Let p be a prime, and $\{C_n\}$ a family of exact-of-Mod_p-of-AND circuits of size less than 2^{n^ϵ} and with bottom fan-in $o(n^{1-\epsilon})$. Then for any n, for some $s \in \mathbf{N}$, $\neg C_n$ can be p^s-represented by a polynomial of degree $o(n)$.*

Proof: We follow the proof of the second part of Toda's theorem, as applied to circuits as in [BT]. We make use of the *modulus-amplifying* polynomials Q_d, which have the property that for every $m \geq 1$ and $X \geq 0$,

$$\begin{aligned} X \equiv 0 \ (m) &\Rightarrow Q_d(X) \equiv 0 \ (m^d), \\ X \equiv 1 \ (m) &\Rightarrow Q_d(X) \equiv 1 \ (m^d). \end{aligned}$$

(The modulus-amplifying polynomials constructed by Beigel and Tarui [BT] have degree $2d-1$.) Denote the i-th Mod_p-of-AND subcircuit of C_n by s_i. By hypothesis, for some k where $k \leq 2^{n^\epsilon}$, C_n outputs 1 iff $\sum_{i=1}^k s_i = \ell$, for some integer ℓ. In turn, s_i outputs 1 iff a certain sum of AND gates is nonzero mod p. Denote this sum by $\sigma_i(x_1, ..., x_n)$, so that s_i outputs 1 iff $\sigma_i(x_1, ..., x_n) \not\equiv 0 \ (p)$. Since the AND gates have fan-in $o(n^{1-\epsilon})$, $\sigma_i(x_1, ..., x_n)$ is a polynomial in the inputs of degree $o(n^{1-\epsilon})$. Without loss of generality, we may assume that $\sigma_i(x_1, ..., x_n)$ mod p is always either 0 or 1 (this can be assured by raising σ_i to the power $p-1$ and invoking Fermat's little theorem; the powering only multiplies the degree of σ_i by the constant $p-1$). Now by the properties of Q_d given above,

$$\begin{aligned} \sigma_i(x_1, ..., x_n) \equiv 0 \ (p) &\Rightarrow Q_{n^\epsilon}(\sigma_i(x_1, ..., x_n)) \equiv 0 \ (p^{n^\epsilon}), \\ \sigma_i(x_1, ..., x_n) \equiv 1 \ (p) &\Rightarrow Q_{n^\epsilon}(\sigma_i(x_1, ..., x_n)) \equiv 1 \ (p^{n^\epsilon}). \end{aligned}$$

Since $p^{n^\epsilon} \geq 2^{n^\epsilon}$, we conclude that C_n outputs 1 iff,

$$\sum_{i=1}^k Q_{n^\epsilon}(\sigma_i(x_1, ..., x_n)) \equiv \ell \ (p^{n^\epsilon}).$$

Now $Q_{n^\epsilon} \circ \sigma - \ell$ is a polynomial t of degree $O(n^\epsilon) \cdot o(n^{1-\epsilon}) = o(n)$. Furthermore, C_n outputs 1 if and only if $t(x_1, ..., x_n) \equiv 0 \ (p^{n^\epsilon})$. Thus t is a polynomial of degree $o(n)$ that p^{n^ϵ}-represents $\neg C_n$.

∎

Corollary 8 *Let p, q be distinct primes and $\epsilon < 1$. No family of exact-of-*Mod_p^+ *circuits of size 2^{n^ϵ} can compute the functions* $\mathrm{Mod}_{q,r}$ *or* $\neg\mathrm{Mod}_{q,r}$.

We also make an easy observation here that was mentioned in the introduction.

Proposition 9 *The Mod-degree of the AND function is n.*

Proof: This is immediate from a result of Tarui [Ta] (viz. proposition 1 in that reference). ∎

Applying Theorem 7, it follows that,

Corollary 10 *For any prime p and $\epsilon < 1$, no family of exact-of-*Mod_p^+ *circuits of size 2^{n^ϵ} can compute the OR function.*

5 Approximating Parity with Exact-of-ACC(p) Circuits

In this section another application of the main theorem is presented. We show that a circuit consisting of an exact gate over bounded-depth circuits with AND, OR and Mod_p gates (for odd prime p) can agree with parity for at most a fraction $1/2 + o(1)$ of all input settings, unless it is of exponential size. The machinery of Smolensky's is not sufficient to prove this because, once again, we need to deal with variable powers of the prime p. Unfortunately the current technique does not seem to be sufficiently powerful to handle the case of Mod_q versus exact-of-ACC(p) for any two distinct primes q, p.

We begin by showing, for any odd prime p, that a polynomial of degree $o(\sqrt{n})$ can p^k-represent parity for at most a fraction $1/2 + o(1)$ of the input settings. Furthermore, this holds for k depending on the number of inputs in an arbitrary manner (it is this case that generalizes Smolensky's theorem). It would be most satisfying to be able to prove this theorem with the p^k replaced by m where m is any odd number. But as will be clear from the proof, we make essential use of the fact that the modulus is a prime power.

Since we are proving lower bounds for parity, there is no need for using complex inputs; we can work directly with the parity basis, in which the inputs y_i, $1 \leq i \leq n$, take on the values $\{-1, 1\}$, and the degree of a polynomial is identical with its weight. The proof of the theorem is similar in spirit to that of Aspnes et al. [ABFR] (for parity versus perceptrons). A key difference is that we are working in a finite ring ($\mathbf{Z}/p^k\mathbf{Z}$) with divisors of zero, and this requires much greater care in solving the linear equations as in the proof of [ABFR]. But the main idea is the same: we show that if we can "approximately" p^k-represent parity well with a low-degree polynomial t, then we can weakly p^k-represent it with a polynomial of degree less than $n/2$. By the main theorem (with q=2), this is a contradiction.

The theorem described above is given after the following analogue of Lemma 2.1 of [ABFR]. The lemma is used to construct a polynomial which is zero on all the points of disagreement between parity and t.

Lemma 11 *Let m be any odd number. Let $S \subseteq \{-1,1\}^n$ be such that $||S|| < \sum_{i=0}^{l} \binom{n}{i}$ where $l < n/2$. Then there is a degree l integer polynomial $w : \{-1,1\}^n \to \mathbf{Z}$ obeying the following properties:*

(i) For any $\mathbf{y} \in S$, we have $w(\mathbf{y}) = 0$.

(ii) There is a $\mathbf{y} \in \{-1,1\}^n$ with $\prod_{i=1}^{n} y_i = -1$, such that $w(\mathbf{y}) \not\equiv 0\ (m)$.

Proof: Write down an arbitrary degree l polynomial w_1 with rational coefficients. The number of coefficients in w_1 is one more than the upper bound on $||S||$. Thus, for each $\mathbf{y} \in S$, we can set $w_1(\mathbf{y}) = 0$ (note that here we are using equality over the rationals), and regard this as a set of linear homogeneous equations for the coefficients of w_1. There is a non-trivial solution for the coefficients, since there are more "variables" than equations. Since the equations are homogeneous, we can find an *integer* solution $w_2(\mathbf{y})$ by multiplying w_1 times a least common denominator. For the same reason we can divide w_2 by a greatest common divisor to obtain the degree l integer polynomial $w(\mathbf{y})$. We now prove that w has the desired properties. Property (i) is obvious. For property (ii), first note that some coefficient of $w(\mathbf{y})$ is not divisible by m, since we have divided by a GCD. We claim that this implies property (ii). For suppose property (ii) did not hold, that is for every $\mathbf{y}$ with $\prod_{i=1}^{n} y_i = -1$, we have $w(\mathbf{y}) \equiv 0\ (m)$. Since w has degree $< n/2$, by Lemma 4 (with $q = 2$ and $r = 1$), we conclude that all the coefficients of w are zero mod m, which is a contradiction. ∎

Theorem 12 *Let p be an odd prime, and let $t : \{-1,1\}^n \to \mathbf{Z}$ be an integer polynomial of degree $o(\sqrt{n})$. Then for sufficiently large n, for any integer k,*

$$||\{\mathbf{y} \in \{-1,1\}^n | \prod_{i=1}^{n} y_i = -1 \quad \textit{iff} \quad t(\mathbf{y}) \not\equiv 0\ (p^k)\}|| \leq 2^n(\frac{1}{2} + o(1)).$$

Proof: Let Δ denote the set of "disagreements" between t and parity:

$$\begin{aligned}\Delta = \{\mathbf{y} \in \{-1,1\}^n \mid & (\prod_{i=1}^{n} y_i = -1 \quad AND \quad t(\mathbf{y}) \equiv 0\ (p^k)) \\ & \qquad\qquad OR \\ & (\prod_{i=1}^{N} y_i = 1 \quad AND \quad t(\mathbf{y}) \not\equiv 0\ (p^k))\}.\end{aligned}$$

It suffices to show that $||\Delta||$ is "large," that is, at least a fraction $1/2 - o(1)$ of all input settings. Let d denote the degree of t (by hypothesis, $d = o(\sqrt{n})$). Suppose

$$||\Delta|| < \sum_{i=1}^{n/2-d-1} \binom{n}{i}. \tag{5}$$

By Lemma 11, there is a degree $n/2-d-1$ polynomial w such that for each $\mathbf{y} \in \Delta$, $w(\mathbf{y}) \equiv 0\ (p^k)$, but there is some $\mathbf{y}$ with $\prod_{i=1}^n y_i = -1$ such that $w(\mathbf{y}) \not\equiv 0\ (p)$. Since, by the definition of Δ, we set $w(\mathbf{y}) = 0$ for any odd-parity $\mathbf{y}$ with $t(\mathbf{y}) \equiv 0\ (p^k)$, it follows that $\prod_{i=1}^n y_i = -1$ and $w(\mathbf{y}) \not\equiv 0\ (p)$ implies $t(\mathbf{y}) \not\equiv 0\ (p^k)$. We know that a $\mathbf{y}$ exists with $\prod_{i=1}^n y_i = -1$ and $w(\mathbf{y}) \not\equiv 0\ (p)$. For such a $\mathbf{y}$, it is easy to see that $t(\mathbf{y})w(\mathbf{y}) \not\equiv 0\ (p^k)$. (It is worth pointing out that the previous sentence is the *one* point in the proof where we use the fact that we are working modulo a prime power!) Thus the integer polynomial $t(\mathbf{y})w(\mathbf{y})$ has degree $n/2-1$, it is not the zero function mod p^k, but it is 0 mod p^k whenever $\prod_{i=1}^n y_i = 1$. In other words, it weakly p^k-represents the parity function and has degree $< n/2$. By Theorem 5, this is a contradiction. Thus $||\Delta||$ cannot obey (5). Using the fact that $d = o(\sqrt{n})$, it is easy to show that this implies that $||\Delta|| \geq 2^n(1/2 - o(1))$, which proves the theorem. ∎

To obtain the lower bound for exact-of-ACC(p) circuits, we need the part of Smolensky's theorem [Sm] that shows how to approximate an ACC(p) circuit with a low-degree polynomial mod p. This is given in the following version of this lemma, which is easily seen to follow from the proof of lemma 2 in [Sm] (the statement below follows more closely that of lemma VIII.3.3 in [Str]).

Lemma 13 (Smolensky): *Let p be prime, and let $r : \mathbf{N} \to \mathbf{N}$ be such that $r(n) = o(n^{1/4d})$. Let $\{C_n\}$ denote a family of Boolean functions computed by ACC(p)-type circuits of depth d and size $2^{r(n)}$. Then there exists a family of polynomials t_n over $\mathbf{Z}/p\mathbf{Z}$ such that the degree of t_n is $o(n^{1/4})$, and $t_n(\mathbf{y}) \equiv C_n(\mathbf{y}) \quad (p)$ for a set of $\mathbf{y} \in \{-1,1\}^n$ of cardinality at least $2^n(1-2^{-r(n)})$.*

The lower bound for exact-of-ACC(p) circuits is expressed in terms of a tradeoff between the number of ACC(p) subcircuits and the size of those subcircuits. The theorem says that if both the number of subcircuits and the size of the subcircuits are too small, than the exact-of-ACC(p) circuits cannot compute parity (and indeed cannot even approximate it).

Theorem 14 *Fix any integer d, any real number $0 < \epsilon < 1/4d$, and any odd prime p. Let $r : \mathbf{N} \to \mathbf{N}$ be such that $r(n) = o(n^{1/4d})$. Consider a family $\{C_n\}$ of exact-of-ACC(p) circuits of depth $d+1$, where for each n the number of ACC(p) subcircuits in C_n is $\leq 2^{r(n)-n^\epsilon}$, and the size of each ACC(p) subcircuit is $\leq 2^{r(n)}$. Then, for all sufficiently large n, C_n agrees with parity for at most a fraction $1/2 + o(1)$ of the input settings.*

Proof: Let c_i, $1 \leq i \leq 2^{r(n)-n^\epsilon}$, denote the ACC($p$) subcircuits of C_n. Since c_i has depth d and size less than $2^{r(n)}$ where $r(n) = o(n^{1/4d})$, it follows from Lemma 13 that for some polynomial $t_i : \{-1,1\}^n \to \mathbf{Z}$ of degree $o(n^{1/4})$, $c_i(\mathbf{y}) \not\equiv t_i(\mathbf{y})\ (p)$ for at most $2^{n-r(n)}$ many $\mathbf{y}$'s. It is convenient to express this as a probability over $\mathbf{y}$ chosen uniformally at random from $\{-1,1\}^n$:

$$Pr(c_i(\mathbf{y}) \not\equiv t_i(\mathbf{y})\ (p)) \leq 2^{-r(n)}.$$

Since there are $\leq 2^{r(n)-n^\epsilon}$ subcircuits, it follows that,

$$\begin{aligned} Pr((\exists i) c_i(\mathbf{y}) \not\equiv t_i(\mathbf{y})\ (p)) &\leq \frac{2^{r(n)-n^\epsilon}}{2^{r(n)}} \\ &= o(1). \end{aligned}$$

So for at least $2^n(1-o(1))$ input settings, all subcircuits c_i agree with their polynomials t_i. On these input settings, C_n can be simulated by an exact-of-Mod$_p$-of-AND circuit with $2^{r(n)-n^\epsilon}$ Mod$_p$-of-AND subcircuits whose bottom fan-in is $o(n^{1/4})$ (corresponding to the degree of t_i). We now apply the method of Theorem 7. Compose each polynomial t_i with the modulus-amplifying polynomial $Q_{n^{1/4d}}$. This gives a polynomial $t'_i = Q_{n^{1/4d}} \circ t_i$ of degree $n^{1/4d} \cdot o(n^{1/4}) = o(\sqrt{n})$ which equals t_i modulo $p^{n^{1/4d}}$. Since $2^{r(n)-n^\epsilon} = o(p^{n^{1/4d}})$, we can add up the t'_i's to obtain a polynomial t of degree $o(\sqrt{n})$ which $p^{n^{1/4d}}$-represents C_n on $2^n(1-o(1))$ input settings. But such a polynomial cannot agree with parity on more than $2^n(1/2+o(1))$ settings, by Theorem 12. A simple probabilistic argument shows that this implies that C_n can agree with parity on at most $2^n(1/2+o(1))$ settings. ∎

It is certainly natural to conjecture that Theorem 14 holds for Mod$_q$ versus exact-of-ACC(p) circuits, where q and p are any two distinct primes (or even relatively prime). However, the current technique does seem to single out $q = 2$ as being special, and this requires a bit of explanation. The problem is in proving a good lower bound on the number of disagreements, which in the case of $q = 2$ is the quantity on the right hand side of equation 5. It is critical that approximately 1/2 (or, at any rate, a constant fraction) of the binomial coefficients are summed. For $q = 2$ it is fortunate that the lower bound we have on the weak degree is about $n/2$. Since $n/2$ also happens to be the average weight of a monomial, the lower bound on $||\Delta||$ is close to half the sum of the binomial distribution over randomly chosen monomials. However, with Mod$_q$ (for odd prime q) the situation is very different. In this case (for N *Boolean* inputs) we can do no better than sum $\binom{N}{i}$ up to the weak rp-Mod-degree, approximately $N/2(q-1)$. But this is well below the average weight, which, for N Boolean inputs, is still $N/2$. That is, the best lower bound we can get on $||\Delta||$ is at most,

$$\sum_{i=0}^{N/2(q-1)} \binom{N}{i},$$

which, by the Chernoff bound, is exponentially smaller than 2^N. It is not possible to get around this problem by obtaining a better lower bound on the rp-Mod-degree. The lower bound obtained in this paper is fairly tight, as is evident from the following upper bound.

Proposition 15 *The rp-Mod-degree of the Mod$_q$ function is at most $\lfloor N/q \rfloor + 1$.*

Proof: Let $x = \sum_{i=1}^{N} x_i$, and let

$$t(x_1, ..., x_N) = \prod_{j=0}^{\lfloor N/q \rfloor} (x - jq),$$

which is obviously of the degree given in the proposition. Choose m relatively prime to q such that for any $x_1, ..., x_N$, $|t(x_1, ..., x_N)| < m$. Then clearly $t(x_1, ..., x_N) \equiv 0\ (m)$ if and only if $x \equiv 0\ (q)$. ∎

6 Oracle Separations of $\mathrm{C_{=}P}^{\mathrm{Mod}_q\mathrm{P}}$

The lower bounds on exact-of-Mod_p^+ and exact-of-ACC(p) circuits imply oracle separations of some well-known polynomial-time counting classes. The exact gate corresponds to the class $\mathrm{C_{=}P}$ [Wa 86] and the Mod_p^+ circuits to the classes $\mathrm{Mod}_p\mathrm{P}$ [He 90], [BeiGil 92].

Theorem 16 *For any pair of distinct primes q, p, there is an oracle A such that $\mathrm{Mod}_q\mathrm{P}^A \not\subseteq \mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}^A}$.*

Proof: (Sketch): We use a standard reduction of the oracle separation problem to a circuit problem, as in [FSS] (see also [Has]). Briefly, define the test language $L = \{1^n |$ the number of strings in A of length n is $\not\equiv 0\ (q)\}$, which is clearly in $\mathrm{Mod}_q\mathrm{P}^A$. Then we diagonalize against $\mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}^A}$-machines. Here we use the circuit lower bounds from Corollary 8 to show that there is sufficient room to diagonalize (that is, in stage i of the diagonalization, n_i can be chosen so that strings of that length can be added to A in such a way that the i-th $\mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}^A}$-machine does not determine $1^{n_i} \in L$ correctly). ∎

We remark that Toda's theorem $\mathrm{PH} \subseteq \mathrm{PP}^{\oplus\mathrm{P}}$ cannot be improved to $\mathrm{PH} \subseteq \mathrm{C_{=}P}^{\oplus\mathrm{P}}$ relative to all oracles:

Proposition 17 *There is an oracle A such that $\mathrm{NP}^A \not\subseteq \mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}^A}$, for all prime p.*

Proof: The proof is similar to that of Theorem 16, where now we appeal to Corollary 10. ∎

The above proposition does not hold relative to a random oracle; in fact, relative to a random oracle $\mathrm{C_{=}P}$ contains the entire polynomial hierarchy [Ta]. But the results of section 5 yield random oracle separations of $\oplus\mathrm{P}$ from $\mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}}$. From Theorem 12 and standard techniques for random oracles (see [BenGil 81] and [Has]), we find,

Theorem 18 *Relative to a random oracle A, for any odd prime p, $\oplus\mathrm{P}^A \not\subseteq \mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{P}^A}$.*

It is known that relative to a random A, $\mathrm{Mod}_p\mathrm{P}^{\mathrm{PH}^A} \subseteq \mathrm{Mod}_p\mathrm{P}^A$ [Ta]. Using this, the above theorem gives the following stronger result (another proof follows from Theorem 14). Let $\mathrm{Mod}_p\mathrm{PH}$ denote the closure of P under the operations $\mathcal{C} \mapsto \mathrm{NP}^{\mathcal{C}}$ and $\mathcal{C} \mapsto \mathrm{Mod}_p\mathrm{P}^{\mathcal{C}}$, which is the polynomial-time analogue of $\mathrm{ACC}(p)$.

Theorem 19 *Relative to a random oracle A, for any odd prime p, $\oplus\mathrm{P}^A \not\subseteq \mathrm{C_{=}P}^{\mathrm{Mod}_p\mathrm{PH}^A}$.*

7 Discussion and Open Problems

The most important contribution of this paper has been to extend existing techniques and develop new ones for proving lower bounds on the degree of polynomials, over $\mathbf{Z}/m\mathbf{Z}$ where m is not a prime power, representing modular functions. Such lower bounds can be used to obtain lower bounds on circuit size, as is well known and is illustrated here in the results of section 5.

Once the problem was translated into the appropriate algebraic language, the proof of the main theorem was surprisingly simple: compute the Fourier transform of the polynomial assuming it has small degree, and find that the coefficients are zero (contradiction). Most of the necessary work was in arriving at the appropriate algebraic setting. In this instance, the appropriate setting was the ring of polynomials over the algebraic integers $\overline{\mathbf{Z}}$.

It should be noted that very few properties of $\overline{\mathbf{Z}}$ were used. Algebraic integers are most useful when computing with polynomials in *different* roots of unity, while here we only needed one (ζ was fixed throughout section 3). As mentioned in a footnote, it would have been sufficient to extend to $\mathbf{Z}[\zeta]$ (the ring of integer polynomials in ζ) rather than $\overline{\mathbf{Z}}$, since the coefficients that arise in encoding the roots of unity as Boolean inputs are all in $\mathbf{Z}[\zeta]$. In this sense, Lemma 4 is stonger than is necessary (in particular, the stronger hypothesis $c_{\mathbf{e}} \in \mathbf{Z}[\zeta]$ would suffice). This suggests that further investigation of the properties of polynomials in rings of algebraic integers may be fruitful.

The open problems which seem most amenable to attack are related to the power of exact-of-$\mathrm{ACC}(p)$ circuits. General exact-of-ACC circuits have been considered previously by Beigel, Tarui and Toda [BTT], who show that probabilistic exact-of-ACC circuits of quasipolynomial size can be simulated by circuits consisting of a symmetric gate over polylog fan-in AND's (called SYM^+ circuits). Thus any lower bounds for SYM^+ circuits are at least as hard as for probabilistic exact-of-ACC circuits. We conjecture that for any pair of distinct primes q, p, the Mod_q function requires exponential size exact-of-$\mathrm{ACC}(p)$ circuits. From the proof of Theorem 14, it would suffice to show that the Mod_q function cannot be well-approximated by an exact-of-Mod_p^+ circuit. This problem, in turn, reduces to showing that the Mod_q function cannot be approximately p^k-represented by a low-degree polynomial. More generally, we conjecture that if q and m are relatively prime, Mod_q cannot be m-represented by a low-degree polynomial for any more than a certain constant fraction of the inputs (note m is permitted to vary with the number of inputs; for constant m, the conjecture is proved [Bei 93]). In this paper we have succeeded in proving this only in the special case of $q = 2$ and even then it was necessary to assume m is a prime power. While many of these problems may appear to be of a technical nature, we believe that their resolution will reveal more widely applicable lower bound techniques for modular functions.

Acknowledgements: I wish to thank Arthur Chou for helpful conversations, and some anonymous referees for helpful suggestions on an earlier version of this paper. The hospitality of the Computer Science Department of Boston University, where part of this work was done during a sabbatical from Clark University, is gratefully acknowledged.

References

[ABFR] J Aspnes, R. Beigel, M. Furst, and S. Rudich, The expressive power of voting polynomials, in *Combinatorica*, **14(2)**, (1994), 1–14.

[Al] E. Allender, A note on the power of threshold circuits. In *Proceedings of the 30th Symposium on Foundations of Computer Science*, (1989), 580-584.

[Bar] D. Barrington, Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1, in *Journal of Computer and System Sciences* **38**, (1989), 150-164.

[BBR] D. M. Barrington, R. Beigel, and S. Rudich, Representing Boolean functions as polynomials modulo composite numbers, in *Computational Complexity* **4** (1994) 367–382.

[BDG] J.L. Balcázar, J. Díaz, and J. Gabarró, *Structural Complexity I.* Springer, 1987.

[Bei 92] R. Beigel, When do extra majority gates help? polylog(n) majority gates are equivalent to one, in *Computational Complexity* **4** (1994) 314–324.

[Bei 93] R. Beigel, The polynomial method in circuit complexity, *Proceedings of the 8th IEEE Conference on Structure in Complexity Theory* (1993) 82-95.

[BeiGil 92] R. Beigel and J. Gill, Counting classes: thresholds, parity, mods, and fewness in *Theoretical Computer Science* **103** (1992) 3-23.

[BenGil 81] C. Bennett and J. Gill, Relative to a random oracle A, $\mathrm{P}^A \neq \mathrm{NP}^A \neq$co-$\mathrm{NP}^A$ with probability 1, *SIAM Journal on Computing,* **10(1)** (1981) 96-113.

[BRS] R. Beigel, N. Reingold, and D. Spielman, PP is closed under intersection, in *Proceedings of the 23rd ACM Symposium on the Theory of Computation*, (1991), 1-11. A journal version is to appear in *Journal of Computer and System Sciences*.

[BS] D. M. Barrington and H. Straubing, Complex polynomials and circuit lower bounds for modular counting, in *Computational Complexity* **4** (1994) 325-338.

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

[BTT] R. Beigel, J. Tarui, and S. Toda, On probabilistic ACC circuits with an exact-threshold output gate, in *Proceedings of the 3rd International Symposium on Algorithms and Computation*, (1992) 420-429.

[CGT] J.-Y. Cai, F. Green and T. Thierauf, On the correlation of symmetric functions, to appear in *Mathematical Systems Theory.*

[FFK] S. Fenner, L. Fortnow and S. Kurtz, Gap-definable counting classes, in *Proceedings of the Sixth Annual Conference on Structure in Complexity Theory,* IEEE Computer Society Press (1991) 30-42.

[FSS] M. Furst, J. B. Saxe, and M. Sipser, Parity, circuits, and the polynomial-time hierarchy, in *Mathematical Systems Theory,* **17** (1984) 13-27.

[GKRST] F. Green, J. Köbler, K. Regan, T. Schwentick, and J. Torán, The power of the middle bit of a #P function, in *Journal of Computer and System Sciences* **50** (1995) 456-467.

[Gre 91] F. Green, An oracle separating $\oplus$P from $\mathrm{PP}^{\mathrm{PH}}$, *Information Processing Letters* **37** (1991) 149-153.

[Gre 93] F. Green, On the power of deterministic reductions to $\mathrm{C}_=\mathrm{P}$, in *Mathematical Systems Theory* **26** (1993) 215-233.

[Has] J. Håstad, Computational limitations of small-depth circuits, the MIT press, Cambridge, 1987.

[He 90] U. Hertrampf, Relations among MOD-classes. In *Theoretical Computer Science* **74** (1990) 325-328.

[HMPST] A. Hajnal, W. Maass, P. Pudlák, M. Szegedy, and G. Turán, Threshold circuits of bounded depth, in *Proceedings 28th Annual IEEE Symposium on Foundations of Computer Science*, IEEE Computer Society Press (1987) 99-110.

[IR] K. Ireland and M. Rosen, A classical introduction to modern number theory, Second Edition, Springer-Verlag, New York, 1990.

[KP] M. Krause and P. Pudlák, On the computational power of depth 2 circuits with threshold and modulo gates, in *Proceedings of the Twenty-Sixth Annual ACM Symposium on the Theory of Computing,* ACM Press (1994) 48-57.

[Raz] A. A. Razborov, Lower bounds on the size of bounded depth networks over a complete basis with logical addition, *Matematicheskie Zametki* **41** (1987) 598-607. English translation in *Mathematical Notes of the Academy of Sciences of the USSR* **41** (1987) 333-338.

[Sm] R. Smolensky, Algebraic methods in the theory of lower bounds for Boolean circuit complexity, in *Proceedings of the 19th Annual ACM Symposium on Theory of Computing* (1987) 77-82.

[Str] H. Straubing, Finite Automata, Formal Logic and Circuit Complexity, Birkhäuser, Boston 1994.

[Ta] J. Tarui, Degree complexity of Boolean functions and its applications to relativized separations, in *Proceedings of the Sixth Annual Conference on Structure in Complexity Theory*, IEEE Computer Society Press (1991) 382-390.

[Tod] S. Toda. PP is as hard as the polynomial-time hierarchy. In *SIAM Journal on Computing* **20**, (1991) 865-877.

[Tsai] S.-C. Tsai, Lower bounds on representing Boolean functions as polynomials in Z_m, in *Proceedings of the Eighth Annual Conference on Structure in Complexity Theory*, IEEE Computer Society Press (1993) 96-101.

[Wa 86] K. Wagner, The complexity of combinatorial problems with succinct input representation. In *Acta Informatica* **23**, (1986), 325-356.