Electronic Colloquium on Computational Complexity, Report No. 116 (2026)

On CC^0 Lower Bounds for AND via Torus Polynomials

Vaibhav Krishan* Jayalal Sarma†

July 7, 2026

Abstract

We explore the torus polynomial approximation based approach towards a long-standing question: whether AND can be computed by CC^0 circuits - the class of constant-depth polynomial size circuits containing MOD_m gates for some natural number m. Bhrushundi, Hosseini, Lovett and Rao (ITCS 2019) introduced torus polynomial approximations as an approach for proving lower bounds against ACC^0 - a class containing CC^0 where the circuits are also allowed AND, OR and NOT gates.

We show how lower bounds for torus polynomials approximating AND can be used to make progress on this question. Using lower bounds on the degree of symmetric torus polynomials approximating AND, proved by Krishan and Vishwanathan (ITCS 2026), we prove size lower bounds for *symmetric* CC^0-circuits computing AND. More precisely, we prove that any depth h symmetric CC^0 circuit requires $2^{\widetilde{\Omega}(n^{1/O(h)})}$ size to compute AND.

A key ingredient in our proof is an argument that we can construct symmetric torus polynomials to approximate symmetric CC^0 circuits. Our construction exhibits an explicit correspondence between the symmetry of the circuit and that of the polynomial. Using this, we also establish lower bounds for weaker notions of circuit symmetry. Lower bounds for symmetric CC^0 circuits were also independently established by Pago (ICALP 2026) using different techniques.

In the *asymmetric* regime, we establish degree upper bounds for depth three circuits of the form $\mathsf{MOD}_p \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ where $m = pq$ is a semiprime. This circuit class is a special case of the *constant degree hypothesis*, introduced by Barrington, Straubing and Thérien (Information and Computation, 1990), where m could be an arbitrary composite number. We argue that improved lower bounds for asymmetric torus polynomials approximating AND imply size lower bounds for semiprime m and hence progress on the constant-degree hypothesis.

*The Institute of Mathematical Sciences, Chennai, India. Email: vaibhavk@imsc.res.in
†Indian Institute of Technology Madras (IIT Madras), Chennai, India. Email: jayalal@cse.iitm.ac.in

ISSN 1433-8092

Contents

1 Introduction

The polynomial method has proven to be a powerful tool for tackling fundamental questions in theoretical computer science. In particular, studying polynomial approximations for Boolean functions has led to advances in cryptography, quantum computing and circuit complexity (see [Aar08, BT21] and references therein). In circuit complexity, considering the degree of polynomials approximating Boolean functions has led to remarkable progress in proving circuit lower bounds for explicit functions, a notoriously difficult quest in the area.

Landmark results in this direction were proved by Razborov [Raz87], and independently by Smolensky [Smo87]. They proved that for any prime p, MOD_p[1] cannot be computed by constant-depth polynomial size circuits that use $\mathsf{AND}, \mathsf{OR}, \mathsf{NOT}$ and MOD_q gates for a prime $q \neq p$. In addition, they used the same technique to prove that such circuits cannot compute $\mathsf{MAJORITY}$. One main technical step in these arguments is that any function in the circuit class can be approximated by a low-degree polynomial over the finite field $\mathbb{F}_q$. This is complemented by an argument that the candidate function requires a high degree for any polynomial that approximates it.

Since then, polynomial approximations have been used extensively for various lower bounds against constant-depth circuits [AKV20, HRRY19, LSS+21, OSS19], and even for constructing *pseudo-random generators* fooling such circuits [HS19]. However, the method has hit a roadblock in showing superpolynomial lower bounds when MOD_6 gates are allowed. More generally, it is unclear whether low-degree polynomials over fields/rings can approximate functions in ACC^0, the class of constant-depth polynomial size circuits that use MOD_m gates for some natural number m.

[1] A MOD_p gate outputs 1 if and only if the sum of its inputs is not divisible by p.

Recently, Bhrushundi, Hosseini, Lovett and Rao [BHLR19] introduced a new notion of polynomial approximations, called *torus polynomial approximations* (see also [Kri21]) [2], specifically towards approximating functions in ACC^0. They showed that all functions in ACC^0 have low-degree torus polynomial approximations, establishing it as a plausible method for proving lower bounds against ACC^0. However, strong lower bounds are not known on the degree required for approximating any explicit function under this notion of approximation. In particular, the authors in [BHLR19] conjectured that low-degree torus polynomials cannot approximate MAJORITY.

The authors in [BHLR19] further showed that if we restrict to symmetric polynomials[3], then MAJORITY requires $\Omega\left(\sqrt{\frac{n}{\log n}}\right)$ degree to approximate it within an error of $\frac{1}{20n}$. In a subsequent work, Krishan and Vishwanathan [KV26] showed that this restriction on symmetry is too strong. The authors proved, in [KV26, Theorem 7], that the same degree lower bound holds even for the AND function. They use this lower bound to argue that studying symmetric torus polynomials is not suitable for resolving the MAJORITY vs ACC^0 question. In this work, we use symmetric torus polynomials towards another dual frontier - whether constant-depth polynomial size circuits can compute AND with only MOD gates. More precisely:

Question 1. *Does* CC^0 *contain the* AND *function?*

where $\mathsf{CC}^0 = \bigcup_{m\in\mathbb{N}} \mathsf{CC}^0[m]$, and $\mathsf{CC}^0[m]$ denotes the class of constant-depth polynomial size circuits comprising MOD_m gates.

It is widely believed, especially since the work in [BST90, BBR94], that $\mathsf{AND} \notin \mathsf{CC}^0$. This question is well-understood when m is a power of a prime p. Using classical results from [Raz87, Smo87], one can obtain small-degree polynomials over $\mathbb{F}_p$ representing[4] such circuits. On the other hand, AND requires a large degree to represent over any $\mathbb{F}_p$, hence proving the lower bound.

However, progress on the conjecture when m is composite is limited to slightly superlinear lower bounds [CGPT06] in the general setting, or strong lower bounds in highly restricted settings [KW25, Pag26a]. In fact, a seemingly simple case of Question 1 remains open, posed by Barrington, Straubing and Thérien [BST90], known as the *constant degree hypothesis*. They conjectured that any circuit of the form $\mathsf{MOD}_p \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ requires exponential size to compute AND. Here, p is a prime, and $\mathsf{AND}_{O(1)}$ denotes an AND gate of constant fan-in.

Our Results:

We contribute to the program of proving circuit lower bounds using torus polynomial approximations against symmetric CC^0 circuits. A circuit is said to be *symmetric* if each permutation of its variables can be extended to an automorphism of the underlying DAG. Kawałek and Wieß [KW25]

[2] A polynomial P is a torus polynomial that ε-approximates a Boolean function f, if for every $x \in \{0,1\}^n$, the fractional part of $P(x)$ is ε-close to $\frac{f(x)}{2}$.

[3] A polynomial is symmetric if monomials of the same degree share the same coefficient.

[4] A polynomial P represents a function over $\mathbb{F}_p$ if for each $x \in \{0,1\}^n$, $P(x) = f(x) \mod p$.

initiated the study of symmetric CC^0 circuits, and proved a special case of the constant degree hypothesis. Their lower bound applied to symmetric $\mathsf{MOD}_p \circ \mathsf{MOD}_q \circ \mathsf{AND}_{O(1)}$ circuits for two primes p and q. Soon thereafter, Pago [Pag26a] extended the lower bound to symmetric CC^0 circuits.

Building on the lower bound in [KV26, Theorem 7], on the degree required to approximate AND by symmetric torus polynomials, we prove a size lower bound for *symmetric* CC^0 circuits computing AND. More formally,

Theorem 1.1. *Any symmetric* CC^0 *circuit of depth* h *requires size* $s(n) = 2^{\widetilde{\Omega}(n^{1/O(h)})}$ *to compute* AND_n.

We use the degree of symmetric torus polynomials approximating a function within $\frac{1}{20n}$ error as the measure for the lower bound. More precisely, we construct low-degree symmetric torus polynomials that approximate symmetric CC^0 circuits within $\frac{1}{20n}$ error. We follow the construction from [BHLR19, Corollary 20] to obtain the torus polynomials approximating CC^0 circuits. Our main contribution is to show that the construction produces symmetric torus polynomials if we start with symmetric circuits. Then, we use the degree lower bound from [KV26, Theorem 7] to prove a size lower bound for symmetric CC^0 circuits computing AND.

In an independent line of work towards establishing size lower bounds against symmetric circuits, Kawałek and Wieß [KW25] and Pago [Pag26a] use the period of symmetric functions as the measure. A symmetric function f over n variables has period b if $f(i) = f(i+b)$ for each $0 \leq i \leq n-b$, where $f(i)$ denotes the function's value at points with Hamming weight i. In both [KW25, Pag26a], the main technical argument is to establish that respective symmetric functions computed by the circuits have small period. In contrast, AND has period exactly n, which leads to lower bounds on the size of the circuit.

We show that the degree of approximation by symmetric torus polynomials subsumes the period of a function as a measure. Formally, consider the measure $\mu(f)$ for a Boolean function f, defined as the minimum degree of a symmetric torus polynomial that $\frac{1}{20n}$-approximates f. We show that for any function with a small period, $\mu(f)$ is small. Note that our statement assumes that the period has only a few distinct prime divisors. This is indeed true for the period of symmetric CC^0 circuits as described in [Pag26a] (see [Pag26b, Lemma 3.3] for the exact expression of the period).

Theorem 1.2. *Consider any periodic symmetric function* $f : \{0,1\}^n \to \{0,1\}$ *with period* m*, where* m *has* $O(1)$ *distinct prime divisors. Then, there exists a symmetric torus polynomial of degree* $d = O(m \log^{O(1)}(n))$ *approximating* f *within* $\frac{1}{20n}$ *error. In other words,* $\mu(f) = \widetilde{O}(m)$*, where* $\widetilde{O}$ *hides some polylog factors.*

We note that AND has period n but has $\widetilde{O}(\sqrt{n})$-degree symmetric torus polynomial approximations (see [KV26, Remark 22]). Hence, there is a quadratic separation between the period and the degree. It is conceivable that the degree measure captures a larger class of functions than periodic

functions. Therefore, studying torus polynomial approximations, and in particular symmetric torus polynomials, may enable progress beyond studying the period of functions.

However, note that the size lower bound we prove in Theorem 1.1 is quantitatively weaker than the corresponding lower bound in [Pag26a]. The exponent in our result diminishes as the depth of the circuit grows. On the other hand, in [Pag26a], they prove a size lower bound of the form $2^{\Omega(n^{1/r})}$ for symmetric $\mathsf{CC}^0[m]$ circuits, where r denotes the number of distinct prime divisors of m. Crucially, their lower bound does not depend on the circuit's depth. Theorem 1.2 shows that, in principle, a similar dependence on r can be obtained through degree as a measure as well. To see this, we note that in [Pag26a], they proved that any size-s symmetric $\mathsf{CC}^0[m]$ circuits computes a symmetric function with period $O(s^r)$. Moreover, the period has $O(1)$ many distinct prime divisors. Hence, we can combine Theorem 1.2, and the degree lower bound from [KV26], to obtain a size lower bound of the form $2^{\widetilde{\Omega}(n^{1/2r})}$.

Finally, our approach provides an advantage when considering relaxed notions of symmetries for CC^0 circuits. Pago [Pag26a] also studied CC^0 circuits in a more general setting, which they called *nested block symmetry*, where the automorphism group of the circuit extends a particular subgroup of S_n. Using our techniques, we can derive lower bounds for such circuits as well; see Theorem 3.1 for the statement. We note that our approach establishes a correspondence between the circuit's symmetry and that of the approximating polynomial. This makes it easier to establish the lower bound for restricted notions of symmetry of the circuit without introducing an additional measure (compared to [Pag26a]).

Asymetric Toroidal Approximations: Motivated by the above discussion, we propose the degree of torus polynomial approximations as a measure to study the CC^0 vs AND question. In the special case when the polynomial is symmetric, we get a size lower bound for symmetric CC^0 circuits. Next, we show that progress on lower bounds for general torus polynomials approximating AND can lead to progress on an important case of the constant degree hypothesis. For simplicity, we state the result for $\mathsf{MOD}_2 \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ circuits when m is an even semiprime, see Theorem 4.1 for a more general statement.

Theorem 1.3. *Consider any $\mathsf{MOD}_2 \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ circuit C of polynomial size, where m is an even semiprime. Then, there exists a degree-$O(\log(n))$ torus polynomial approximating C within $\frac{1}{\Omega(n)}$ error.*

This result suggests a potential approach for proving that $\mathsf{MOD}_2 \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ circuits, for an even semiprime n, require superpolynomial size to compute AND_n. Using [KV26, Theorem 6], we get that any torus polynomial approximating AND within $\frac{1}{\Omega(n)}$ error must have degree $\Omega(\log(n))$. The authors note a gap between this lower bound and the known upper bound for AND, and suggest bridging this gap in [KV26, Open Problem 4], hinting at a possibility of improving the lower bound. Using Theorem 1.3, improving the degree lower bound to $\omega(\log(n))$ implies the superpolynomial size lower bound we are aiming for. In the appendix, we present evidence that there is room for improvement in the lower bound argument from [KV26, Theorem 6]. Hence, we

conjecture that the lower bound is indeed stronger, see Conjecture 4.1 for a more general statement.

Conjecture 1.1. *Any torus polynomial that approximates* AND_n *within* $\frac{1}{20n}$ *error requires degree* $\omega(\log(n))$.

Organization of the Paper: We introduce some notation and preliminary statements required for our results in Section 2. Section 3 is dedicated to proving Theorem 1.1, and its extension to nested block symmetry. In Section 4, we describe our proposed approach for making progress on the constant degree hypothesis by improving lower bounds for torus polynomials approximating AND. Finally, in Section 5, we prove Theorem 1.2 about symmetric torus polynomials approximating periodic functions.

2 Preliminaries

We start by formally defining the notion of symmetry for circuits.

Definition 2.1 (Symmetric Circuits)**.** *For a subgroup* $\Gamma \leq S_n$*, a circuit* C *over variables* $(x_1, \ldots, x_n)$ *is called to be* Γ*-symmetric if it satisfies the following property. For each permutation* $\pi \in \Gamma$*, there is a permutation* π' *over the gates in* C *that extends* π*, that is* $\pi'(x_i) = \pi(x_i)$ *for each* $i \in [n]$*, such that the following holds: A gate* g *is connected to another gate* g' *in* C *if and only if* $\pi'(g)$ *is connected to* $\pi'(g')$ *in* C*.*

If $\Gamma = S_n$*, the circuit is called symmetric.*

The circuits we consider are composed of MOD_m gates, defined as follows.

$$\mathsf{MOD}_m(x_1, \ldots, x_n) = \begin{cases} 1 & \sum_i x_i \mod m = 0 \\ 0 & \sum_i x_i \mod m \neq 1 \end{cases}$$

Torus Polynomials: The main tool for our results is approximation by torus polynomials, introduced in [BHLR19]. We define it formally below.

Definition 2.2 (Torus Polynomials ([BHLR19]))**.** *Fix a natural number* n*, and a Boolean function* $f : \{0,1\}^n \to \{0,1\}$*. Consider a polynomial* $P \in \mathbb{R}[x_1, \ldots, x_n]$*. We define* P *as a* torus polynomial *approximating* f *within an error of* ε*, for some* $0 \leq \varepsilon < \frac{1}{4}$*, if the following holds for some integer function* $Z : \{0,1\}^n \to \mathbb{Z}$*: For each* $a \in \{0,1\}^n$*,* $P(a)$ *belongs to the interval* $P(a) \in [Z(a) + f(a)/2 - \varepsilon, Z(a) + f(a)/2 + \varepsilon]$*.*

In other words, the fractional part *of* $P(a)$ *is at most* ε *away from* $\frac{f(a)}{2}$*.*

The proof of our main result crucially relies on symmetric torus polynomials, in which the approximating polynomial is symmetric, as formally defined below.

Definition 2.3 (Symmetric Polynomial)**.** *A polynomial $P \in \mathbb{R}[x_1, \ldots, x_n]$ is symmetric if for any permutation $\pi \in S_n$, $P(x) = P(\pi \circ x)$ syntactically. In other words, monomials of the same degree share the same coefficient.*

In general, P is Γ-symmetric, for some subgroup $\Gamma \leqslant S_n$, if $P(x) = P(\pi \circ x)$ for any $\pi \in \Gamma$. In other words, monomials within the same orbit under Γ share the same coefficient.

In [KV26, Theorem 7], the authors proved a lower bound on the degree of symmetric torus polynomials approximating the AND function. We state this result, a key component in our proof, below.

Theorem 2.1 ([KV26])**.** *Any symmetric torus polynomial that approximates AND_n within $\frac{1}{20n}$ error must have degree $\widetilde{\Omega}(\sqrt{n})$.*

3 Size Lower Bounds for Symmetric CC^0 Circuits

The proof starts by converting a symmetric CC^0 circuit to a symmetric torus polynomial that approximates it within $\frac{1}{20n}$ error. Then, we use Theorem 2.1 [KV26], which allows us to prove the size lower bound. Formally, the conversion result is as follows.

Lemma 3.1. *Consider any symmetric CC^0 circuit of depth h and size s. Then, there exists a symmetric torus polynomial P that approximates it within $\frac{1}{20n}$ error, such that $deg(P) \leq \log^{O(h)}(s)$.*

Proof Sketch. We start with a CC^0 circuit C as considered in the statement. Our plan for finding a torus polynomial that approximates this circuit proceeds in three steps.

1. Convert the CC^0 circuit to a layered circuit such that AND gates appear only at the bottom layer, and each layer consists of MOD_q gates for a single prime power q. The conversion is well-known, see [BT94] for reference. We argue that this conversion can be performed while preserving symmetry, as shown in Lemma 3.2.

2. Next, we use the torus polynomial construction from [BHLR19] to obtain a torus polynomial that approximates C within $\frac{1}{20n}$ error. As our argument depends crucially on this construction, we describe it in Lemma 3.6. This yields a torus polynomial of degree $\log^{O(h)}(s)$.

3. Finally, we argue that the conversion produces a symmetric polynomial when starting with a symmetric circuit. We prove this formally in Lemma 3.6.

At the end, we obtain a symmetric torus polynomial that approximates the circuit, with the claimed degree. □

Using Lemma 3.1, we can complete the proof of our main result as follows.

Proof of Theorem 1.1. Consider a symmetric CC^0 circuit C, with depth h, size s and AND gates of fan-in at most $d = O(1)$ at the bottom. Then, there exists a symmetric torus polynomial P, of degree at most $\log^{O(h)}(s)$ that approximates C within $\frac{1}{20n}$ error. Now, assume that C computes AND_n. Using [KV26, Theorem 7], we get that $\log^{O(h)}(s) \geq \widetilde{\Omega}(\sqrt{n})$. Hence, $s(n) = 2^{\widetilde{\Omega}(n^{1/O(h)})}$. □

3.1 Converting CC^0 circuits to Layered Form

To start the proof of Lemma 3.1, we describe the procedure to convert a symmetric CC^0 circuit into a symmetric *layered modular circuit*. We define layered modular circuits below.

Definition 3.1 (Layered Modular Circuit)**.** *A circuit is called a layered modular circuit if it satisfies the following conditions:*

- *Each gate in the circuit is a* MOD_q *gate for some prime power* q*, except possibly the bottom layer which can consist of* AND *gates of constant fan-in.*
- *Each layer of the circuit uses* MOD_q *gates for a single* q*.*
- *Different* MOD_q *gates use prime powers* q *for distinct primes.*

We use a well-known procedure for this conversion, see [BT94] for reference. Our main contribution is to argue that the conversion preserves symmetry.

Lemma 3.2. *Consider a* CC^0 *circuit of size* s *and depth* h*. Then, it can be transformed to an equivalent layered modular circuit of size* $\mathsf{poly}(s)$ *and depth* $O(h)$*. Crucially, if the original circuit is symmetric, the produced circuit is symmetric as well.*

Proof. Broadly speaking, the conversion involves the following three steps (see [BT94]):

- Simulate a MOD_m gate as an AND gate over MOD_q gates, where each MOD_q uses a prime-power factor q of m, and no two MOD_q gates use a power of the same prime.
- Convert this to a layered circuit by introducing dummy MOD gates.
- Bring all AND gates to the bottom by switching AND and MOD gates at each layer.

Now, we argue that each of these steps can be performed while maintaining symmetry. This suffices for the proof of the statement.

First, consider a MOD_m gate g, where $m = \prod_{i=1}^{f} q_i$ such that each q_i is a prime power for a distinct prime, and denote its inputs by $(g_1, \ldots, g_s)$. The gate g is replace by an AND gate g_a over MOD_{q_i} gates g_{q_i}, each of which take $(g_1, \ldots, g_s)$ as their inputs. To see how symmetry is preserved, choose any permutation $\pi \in S_n$, which extends to a permutation π' over the gates in the original circuit. Apply the transformation described above to $\pi'(g)$, which is connected to

$(\pi'(g_1), \ldots, \pi'(g_s))$. As $\pi'(g)$ is also a MOD_m gate, it is replaced by an AND gate g'_a over MOD_{q_i} gates g'_{q_i}, each of which take $(\pi'(g_1), \ldots, \pi'(g_s))$ as their input. Hence, the permutation π'' over the gates in the new circuit maps $\pi''(g_a) = g'_a$, $\pi''(g_{q_i}) = g'_{q_i}$ and $\pi''(g_i) = \pi'(g_i)$. It is easy to check that π'' is as required.

Second, to ensure that each layer contains MOD_q gates for a single prime power q, we modify the transformation from the first step itself. Instead of replacing the MOD_m gate with an AND over MOD_{q_i} gates, we replace it with a depth $f+1$ gadget. The top gate in this gadget is an AND gate. Now, we order q_is in their increasing order, and form the next f layers using MOD_{q_i} gates in the i^{th} layer. In the first layer, the leftmost MOD_{q_1} gate has fan-in s, its j^{th} input being a path consisting of $\mathsf{MOD}_{q_2}, \ldots, \mathsf{MOD}_{q_f}$ gates with g_j as the input to the bottom MOD_{q_f} gate. The remaining $f-1$ gates on the right have fan-in 1 each, taking a single gate from the second layer as input, as described inductively.

In the i^{th} layer, we describe the leftmost MOD_{q_i} gate after ignoring the gates already introduced by the layers above. The leftmost MOD_{q_i} gate has fan-in s, its j^{th} input being a path consisting of $\mathsf{MOD}_{q_{i+1}}, \ldots, \mathsf{MOD}_{q_f}$ gates with g_j as the input to the bottom MOD_{q_f} gate. All the other gates to its right have fan-in 1, each taking a single gate from the $i+1^{\text{th}}$ layer as its input. This completes the description of the gadget. Clearly, the gadget is equivalent to $\mathsf{MOD}_m(g_1, \ldots, g_s)$, and size $O(f^2 s) = O(s)$. The argument for why it preserves symmetry is similar to the argument for the first step.

Finally, for the third step, we need to bring AND gates to the bottom. This requires multiple applications of a well-known operation [BT94] which switches from $\mathsf{AND} \circ \mathsf{MOD}_q$ to $\mathsf{MOD}_q \circ \mathsf{AND}$. Consider an AND gate, over MOD_q gates $(g_1, \ldots, g_\ell)$, with $(h_{i,j_1}, \ldots, h_{i,j_{f'}})$ as the inputs for g_i. Here, we assume without loss of generality that each g_i has the same fan-in, by inserting "dummy" 0 inputs wherever necessary. Now, represent each MOD_q gate g_i as a polynomial P_i over $\mathbb{Z}_q$ such that $g_i(h_{i,j_i}, \ldots, h_{i,j_{f'}}) = P_i\left(\sum_{j=j_1}^{j_{f'}} h_{i,j}\right)$. Then, write the expression for $\mathsf{AND}(g_1, \ldots, g_\ell)$ as $\prod_{i=1}^{\ell} P_i$. Expand the product to get a single polynomial P, such that each monomial of P becomes an AND gate, feeding into a single MOD_q gate. This completes the description of the operation to switch from $\mathsf{AND} \circ \mathsf{MOD}_q$ to $\mathsf{MOD}_q \circ \mathsf{AND}$.

We argue that each such operation preserves symmetry. Consider an AND gate g, over MOD_q gates $(g_1, \ldots, g_\ell)$, with $(h_{i,j_1}, \ldots, h_{i,j_{f'}})$ as the inputs for g_i. As the original circuit is symmetric, for any permutation $\pi \in S_n$, we have an extended permutation π' such that $\pi'(g)$ takes $(\pi'(g_1), \ldots, \pi'(g_\ell))$ as its inputs, with $(\pi'(h_{i,j_1}), \ldots, \pi'(h_{i,j_{f'}}))$ as the inputs for $\pi'(g_i)$. Now, the polynomial P_i obtained for g_i is the same as P'_i obtained for $\pi'(g_i)$. Hence, for each monomial obtained in $\prod_{i=1}^{\ell} P_i$ of the form $\prod_{(i,j)\in S} h_{i,j}$ for some subset $S \subseteq [\ell] \times [j_1, \ldots, j_{f'}]$, there is a corresponding monomial $\prod_{(i,j)\in S} \pi'(h_{i,j})$. Therefore, we can find a new permutation π'', that maps the AND gate corresponding to $\prod_{(i,j)\in S} h_{i,j}$ with the AND gate corresponding to $\prod_{(i,j)\in S} \pi'(h_{i,j})$. It is easy to verify that the new permutation π'' suffices to prove symmetry for the new circuit.

Finally, one can easily verify that the size of the transformed circuit is at most $O(\mathsf{poly}(s))$, and the depth is $O(h)$. This completes the proof. □

3.2 Torus Polynomial Approximations for Layered CC^0 Circuits

As the next step in the proof of Lemma 3.1, we describe the procedure to obtain a torus polynomial approximating a layered circuit.

Lemma 3.3. *Consider a layered modular circuit of depth h and size s. Then, there exists a torus polynomial P that approximates it within $\frac{1}{20n}$ error, such that $deg(P) \leq \log^{O(h)}(s)$.*

The proof of this statement is implicit in [BHLR19]; however, we reproduce the details of the proof as they are crucial for our argument. As the first step, given a layered modular circuit C, we construct an integer polynomial Q, such that the binary expansion of $Q(a)$ contains the value of $C(a)$ at a fixed index. Formally, we prove the following.

Lemma 3.4. *Consider a layered circuit of depth h and size s. Then for any $e \geq 1$, there exists an integer polynomial Q of degree $(e\log(s))^{O(h)}$ such that the following holds for some $\ell \geq e$:*

$$\forall a \in \{0,1\}^n, Q(a) = C(a)2^\ell + E(a) \mod 2^{\ell+e}$$

where $E(a) \leq 2^{\ell-e}$.

For now, assume that the statement is true. We can complete the proof of Lemma 3.3 as follows.

Proof of Lemma 3.3. For a circuit C as per the statement, construct a polynomial Q for $e = \lceil \log(10n) \rceil$ using Lemma 3.4. As per the construction, for some $\ell \geq e$, we have $Q(a) = C(a)2^\ell + E(a) \mod 2^{\ell+e}$ with $E(a) \leq 2^{\ell-e}$. Define the torus polynomial $P(x) = \frac{Q(x)}{2^{\ell+1}} \mod 1$. Then, we have

$$P(a) = \frac{C(a)}{2} + \frac{E(a)}{2^{\ell+1}} \mod 1$$

The error for P is at most $\frac{E(x)}{2^{\ell+1}} \leq 2^{-e-1} \leq \frac{1}{20n}$. Finally, P has the same degree as Q, which is at most $(e\log(s))^{O(h)} = \log^{O(h)}(s)$. This completes the proof. □

Now, we prove Lemma 3.4. The proof idea is from [BHLR19], based on an induction on the circuit depth. For a depth-1 circuit, the statement is as follows.

Lemma 3.5 ([BHLR19])**.** *For a MOD_q gate, where q is a prime power, there exists an integer polynomial Q of degree $O(qe)$, and an integer $\ell = O(qe\log(n))$, such that for each $a \in \{0,1\}^n$:*

$$Q(a) = \mathsf{MOD}_q(a)2^\ell + E(a) \mod 2^{\ell+e}, 0 \leq E(a) \leq 2^{\ell-e}$$

We present the proof in the appendix for completeness, as it is not present in the published version of [BHLR19]. The inductive proof now proceeds as follows.

Proof of Lemma 3.4. The base case is a depth 1 circuit, which is covered in Lemma 3.5. Now, as the inductive hypothesis, assume that the statement is true for circuits of depth h. For the inductive step, consider a circuit of depth $h+1$.

Suppose the root is a MOD_q gate, with $(y_1, \ldots, y_t)$ as its inputs, $t \leq s$, where $y_i = C_i(x)$ for the sub-circuit C_i of depth h over the original inputs x. Construct a polynomial Q of degree $O(qe)$, such that $Q(\cdot) = \mathsf{MOD}_q(\cdot)2^{\ell} + E(\cdot) \mod 2^{\ell+e}$ with $\ell = O(qe\log(t))$ and $E(\cdot) \leq 2^{\ell-e}$, using Lemma 3.5. Each monomial in this polynomial is a product of at most $O(qe)$ many y_i's. The total number of such monomials is at most $t' = t^{O(qe)}$.

Now, each monomial can be seen as an AND over the corresponding sub-circuits C_i, which can be propagated to the leaves. Finally, we get a sub-circuit $D_j(A_j)$ corresponding to each monomial of Q, where A_j is a product of $O(qe)$ many variables. For each D_j, use the inductive hypothesis to obtain a polynomial Q_j, such that

$$Q_j(A_j(x)) = D_j(A_j(x))2^{\ell'} + E_j(A_j(x)) \mod 2^{\ell'+e'}$$

where $E_j(A_j(x)) \leq 2^{\ell'-e'}$. We choose the value for e' later.

Now, we have

$$\sum_j D_j(A_j(x)) = C(x)2^{\ell} + E(x) \mod 2^{\ell+e}$$

Substituting the values of D_j, we get

$$\begin{aligned}\sum_j Q_j(A_j(x)) &= \left(\sum_j D_j(A_j(x))\right)2^{\ell'} + E'(x) \mod 2^{\ell'+e'} \\ &= C(x)2^{\ell+\ell'} + 2^{\ell'}E(x) + E'(x) \mod 2^{\ell+\ell'+e}\end{aligned}$$

To ensure the correctness of the latter equality, we will choose e' such that $\ell' + e' \geq \ell + \ell' + e$.

Now, as per the claimed statement, we need to show that $E'(x) + 2^{\ell'}E(x) \leq 2^{\ell+\ell'-e}$. However, this does not follow directly, as it can be the case that $E(a) = 2^{\ell-e}$ and $E'(a) > 0$ for some $a \in \{0,1\}^n$. Nonetheless, note that $E'(x) \leq t'2^{\ell'-e'}$. Hence, if we choose $e' \geq \lceil \log(t') \rceil$, we get $E'(x) + 2^{\ell'}E(x) \leq 2^{\ell'-O(e\log(s))} + 2^{\ell+\ell'-e} \leq 2^{\ell+\ell'-e+1}$. Therefore, if we start the proof by constructing Q with $e+1$ when using Lemma 3.5, we can ensure $E'(x) + 2^{\ell'}E(x) \leq 2^{\ell+\ell'-e}$.

To finish the proof, first note that the choice of $e' = \lceil \log(t') \rceil$ gives us $e' = O(qe\log(s))$. This is sufficient to ensure $\ell' + e' \geq \ell + \ell' + e$. For the chosen e', the degree of each Q_j is $O\big((e\log(s))^{h-1}\big)$.

The final polynomial is $\sum_j Q_j(A_j(x))$, with each A_j being a product of $O(qe)$ many variables. Hence, the degree bound we obtain is $O(e\log(s))^h$, as claimed in the statement. This completes

the proof. □

3.3 Symmetric Circuits Lead to Symmetric Torus Polynomials

The final step in the proof of Lemma 3.1 is to prove that the above construction produces a symmetric torus polynomial for a symmetric CC^0 circuit. We actually prove a more general statement about Γ-symmetric circuits for an arbitrary subgroup $\Gamma \leqslant S_n$.

Lemma 3.6. *For a Γ-symmetric circuit, the procedure in Lemma 3.3 produces a Γ-symmetric torus polynomial approximation.*

Corollary 1. *For a symmetric layered modular circuit C, the procedure in Lemma 3.3 can be used to obtain a symmetric torus polynomial approximating the function computed by the circuit.*

The proof of Lemma 3.6 follows from a slightly more general statement. We describe how the construction behaves with respect to morphisms over circuits that preserve the edge connections. Informally, we show that the construction produces syntactically similar torus polynomials for circuits that are syntactically similar. The formal statement follows.

Lemma 3.7. *Consider two layered circuits $C = (G, W)$ and $C' = (G', W')$, with inputs $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$ respectively, of the same depth. In the i^{th} layer of both circuits, they use MOD_{q_i} gates for the same q_i. Then, for any map $\rho : G \to G'$ with $\rho(x_i) = y_i$, such that $(g, g') \in W$ if and only if $(\rho(g), \rho(g')) \in W'$, the following holds: If Lemma 3.3 produces $P(x)$ as a torus polynomial approximating C, then it produces $P(y)$ as the torus polynomial approximating C'.*

Proof. We argue this inductively based on the depths of the circuits C, C'. In the base case, with depth 1, we look at the proof of Lemma 3.5. The construction of the polynomial Q depends only on the fan-in of the MOD_q gate and the error parameter e. Hence, the base case follows.

For the inductive case, we look at the proof of Lemma 3.3. The polynomials Q, Q' we construct for the top MOD_q gate, in C, C' respectively, only depend on the fan-in. Hence, for each monomial in Q, we can find a corresponding monomial in Q' using the map ρ.

Now, consider a monomial M_j of Q and the corresponding monomial M'_j from Q' obtained using ρ. We consider the monomials as AND over their inputs, then push the AND gate to the bottom, to obtain circuits of form $D_j(A_j(x))$ and $D'_j(A'_j(y))$. Using induction hypothesis, the torus polynomials we obtain for $D_j(A_j(x))$ and $D'_j(A_j(y))$, denoted by $Q_j(A_j(x))$ and $Q'_j(A'_j(y))$ respectively, are equal. Therefore, we have $\sum_j Q_j(A_j(x)) = \sum_j Q'_j(A'_j(y))$, where the LHS and RHS are the torus polynomials produced for C and C' respectively. This completes the proof of the statement. □

Now, we can finish the proof of Lemma 3.6 as follows.

Proof of Lemma 3.6. Consider a Γ-symmetric layered circuit $C = (G, W)$ as per the statement. Any permutation $\pi \in \Gamma$ extends to a permutation $\pi' \in Aut(G)$, such that $(g, g') \in W$ if and only if $(\pi'(g), \pi'(g')) \in W$. Now, construct a torus polynomial $P(x)$ that approximates C using Lemma 3.3. Then, Lemma 3.7 implies that $P(\pi(x))$ is produced as the torus polynomial approximating $C_{\pi'} = (\pi'(G), W)$.

However, $C_{\pi'} = C$, hence, Lemma 3.3 would produce $P(x)$ as the torus polynomial approximating C. Therefore, $P(\pi(x)) = P(x)$, completing the proof that P is symmetric. □

3.4 Nested Block Symmetric Groups

Now, we consider a weaker notion of symmetry, studied by Pago [Pag26a], which they called nested block symmetry. Nested block symmetric groups are automorphism groups of rooted trees, defined formally as follows.

Definition 3.2 (Nested Block Symmetric Group)**.** *Consider a tree of depth h, such that each node at distance i from the root has exactly k_i many children, and exactly n leaves labeled with $x_1, \ldots, x_n$. Note that this implies $\prod_{i=1}^{h} k_i = n$. Denote this tree by $\mathcal{T}_n^{\mathbf{k}}$, where $\mathbf{k}$ denotes the tuple $(k_1, \ldots, k_h)$. The nested block symmetric group $\Gamma_n^{\mathbf{k}}$ is defined as the group of permutations $\pi \in S_n$ over the leaves such that some $\pi' \in Aut(\mathcal{T}_n^{\mathbf{k}})$ extends π.*

We prove a lower bound for $\Gamma_n^{\mathbf{k}}$-symmetric CC^0 circuits computing AND.

Theorem 3.1. *Any $\Gamma_n^{\mathbf{k}}$-symmetric depth-h CC^0 circuit, with an underlying tree $\mathcal{T}_n^{\mathbf{k}}$, requires $2^{\widetilde{\Omega}(k^{1/O(h)})}$ size to compute* AND*, where $k = \max_{k' \in \mathbf{k}} k'$.*

Proof. To start the proof, we use Lemma 3.6 to construct a $\Gamma_n^{\mathbf{k}}$-symmetric torus polynomial P of degree $\log(s)^{O(h)}(n)$ approximating the circuit within an error of $\varepsilon = \frac{1}{20n}$. Denote the depth of $\mathcal{T}_n^{\mathbf{k}}$ as h'. Now, consider the case when $k = k_{h'}$, i.e. the maximum fan-in in the underlying tree $\mathcal{T}_n^{\mathbf{k}}$ appears at the layer just above the leaves. Then, we choose an arbitrary node above the leaves and set all variables that occur outside this node to 1. After applying this partial restriction, we notice that the group $\Gamma_n^{\mathbf{k}}$ collapses to the symmetric group S_k. Hence, we get a symmetric torus polynomial P' approximating AND over k variables. Therefore, using Theorem 2.1 for AND over k variables, we get the size lower bound.

Otherwise, if $k \neq k_{h'}$, we reduce the depth of $\mathcal{T}_n^{\mathbf{k}}$ by considering each node just above the leaves, and setting all variables appearing within such a node as equal to each other. After applying this partial restriction, we effectively reduce the depth of $\mathcal{T}_n^{\mathbf{k}}$ by one. Note that this transformation does not increase the degree of P', and the function remains AND over the remaining variables. Now, we can proceed inductively till we reach the level where k appears, and use our argument above to finish the proof. □

4 Towards Constant Degree Hypothesis for Semiprime Moduli

In this section, we propose an approach to prove size lower bounds for $\mathsf{MOD}_p \circ \mathsf{MOD}_m \circ \mathsf{AND}_{O(1)}$ circuits computing AND_n, where m has two distinct prime divisors with one of them being p. For the sake of simplicity, we state our results for $m = pq$ being a product of two primes. The approach is again based on torus polynomials: we construct a low-degree torus polynomial approximation for such circuits, albeit in a slightly modified sense. We consider a generalized version of torus polynomial approximation, also studied in [BHLR19], wherein the polynomial is close to αf for some $\alpha \in (0, 1)$. For clarity, we define it formally below.

Definition 4.1 (α-Torus Polynomial Approximation)**.** *For some $\alpha \in (0, 1)$, and a Boolean function $f : \{0, 1\}^n \to \{0, 1\}$, we define P as an α-torus polynomial approximating f within an error of ε, if the* fractional part *of $P(a)$ is at most ε away from $\alpha f(a)$ for each $a \in \{0, 1\}^n$.*

For our main result, we need the following statement, which follows from a simple modification of the proof for Lemma 3.5.

Lemma 4.1. *For a MOD_q gate, where q is a prime power, and any natural number $p \geq 2$, there exists an integer polynomial Q of degree $O(qe)$, and an integer $\ell = O(qe \log(n))$, such that for each $a \in \{0, 1\}^n$:*

$$Q(a) = \mathsf{MOD}_q(a)p^\ell + E(a) \mod p^{\ell+e}, 0 \leq E(a) \leq p^{\ell-e}$$

Now, we proceed with the proof of our main result.

Theorem 4.1. *For any $\mathsf{MOD}_p \circ \mathsf{MOD}_{pq} \circ \mathsf{AND}_{O(1)}$ circuit of size $s = \mathsf{poly}(n)$, there exists a $\frac{1}{p}$-torus polynomial of degree $O(\log(n))$ that approximates the circuit within $\frac{1}{20n}$ error.*

Proof. To start the proof, we rewrite the MOD_{pq} gate as an AND over MOD_p and MOD_q gates. Then, we transform the circuit such that the first two layers are MOD_p gates, then a layer of MOD_q gates, and finally AND gates of $O(1)$ fan-in at the bottom. Now, consider the top two layers of MOD_p gates as a circuit C with inputs $(y_1, \ldots, y_m)$, where each y_i is computed by a $\mathsf{MOD}_q \circ \mathsf{AND}_{O(1)}$ circuit. Construct a polynomial P, using Fermat-Euler theorem, such that $P(\cdot) = C(\cdot) \mod p$. The polynomial P has $O(1)$ degree, hence, it contains at most $t' = \mathsf{poly}(n)$ many monomials. Each of these monomials M_j can be thought of as an $\mathsf{AND}_{O(1)} \circ \mathsf{MOD}_q \circ \mathsf{AND}_{O(1)}$ circuit, which we convert to a $\mathsf{MOD}_q \circ \mathsf{AND}_{O(1)}$ circuit.

Now, for each MOD_q gate D_j with its inputs being $A_j(x)$, we use Lemma 4.1 to obtain a polynomial P_j such that $P_j(A_j(x)) = D_j(A_j(x))p^\ell + E_j(A_j(x)) \mod p^{\ell+1}$, where $E_j(A_j(x)) \leq p^{\ell-e}$. We will choose the value for e later. Then, we get $\sum_{j=1}^{t'} P_j(A_j(x)) = \left(\sum_{j=1}^{t'} D_j(A_j(x))\right)p^\ell + t'E(x) \mod p^{\ell+1} = C(x)p^\ell + t'E(x) \mod p^{\ell+1}$. Finally, we need to ensure that $t'E(x) \leq \frac{p^\ell}{20n}$, for which it suffices to choose a large enough $e = O(\log(n))$. The final polynomial as required by the statement is $\frac{\sum_{j=1}^{t'} P_j(A_j(x))}{p^{\ell+1}}$, which has degree $O(\log(n))$. □

This result suggests a potential approach for proving that $\mathsf{MOD}_p \circ \mathsf{MOD}_{pq} \circ \mathsf{AND}_{O(1)}$ circuits require superpolynomial size to compute AND_n. In [KV26, Theorem 6], the authors proved that any torus polynomial that approximates AND_n within an error smaller than ε requires degree more than $\log\left(\frac{1}{\varepsilon}\right) - 2$. For $\varepsilon = \frac{1}{20n}$, the degree required is $\Omega(\log(n))$. Their proof can be easily extended to prove that any α-torus polynomial that approximates AND_n within $\frac{1}{20n}$ error requires degree $\Omega(\log(n))$. The authors note a gap between this lower bound and the known upper bound, and suggest bridging this gap in [KV26, Open Problem 4], hinting at a possibility of improving the lower bound. Using Theorem 4.1, we note that improving the lower bound to $\omega(\log(n))$ proves the superpolynomial size lower bound we are aiming for. Hence, we conjecture that the lower bound is indeed stronger.

Conjecture 4.1. *For any $\alpha \in (0,1)$, any α-torus polynomial that approximates AND_n within $\frac{1}{20n}$ error requires degree $\omega(\log(n))$.*

5 Degree Upper Bounds for Periodic Functions

In this section, we establish an upper bound on the degree of symmetric torus polynomials approximating periodic functions. We prove Theorem 1.2 restated below.

Theorem 1.2. *Consider any periodic symmetric function $f : \{0,1\}^n \to \{0,1\}$ with period m, where m has $O(1)$ distinct prime divisors. Then, there exists a symmetric torus polynomial of degree $d = O(m \log^{O(1)}(n))$ approximating f within $\frac{1}{20n}$ error. In other words, $\mu(f) = \widetilde{O}(m)$, where $\widetilde{O}$ hides some polylog factors.*

As the first step in the proof, we construct a symmetric layered modular circuit that computes MOD_m. The statement is a simple case of Lemma 3.2, for which we omit the proof.

Lemma 5.1. *There is a symmetric layered modular circuit that computes MOD_m. If m has $O(1)$ distinct prime divisors, then the circuit has size $O(n)$ and depth $O(1)$.*

Next, we describe a construction for torus polynomials approximating layered modular circuits. The proof is a simple modification of the proof for Lemma 3.3, which we omit.

Lemma 5.2. *Consider a layered modular circuit of depth $O(1)$ and size $O(n)$ that uses MOD_{q_i} gates for $i \in [h]$. Then, there exists a torus polynomial of degree $\prod_i q_i \cdot (e \log(n))^{O(1)}$ approximating the circuit within $\varepsilon = 2^{-e}$ error.*

Using this result, we can prove Theorem 1.2 as follows.

Proof of Theorem 1.2. First, we prove the statement for $f = \mathsf{MOD}_m$ function, where m has $r = O(1)$ distinct prime divisors. We use Lemma 5.1 to obtain a symmetric layered modular circuit computing MOD_m. Then, we use Lemma 5.2 to obtain a torus polynomial of degree $m \cdot (\log(n))^{O(1)}$

approximating MOD_m within $\frac{1}{20n}$ error. Using Lemma 3.6, we get that the polynomial is symmetric. This completes the MOD_m case.

Now, consider any symmetric function $f : \{0,1\}^n \to \{0,1\}$ of period m. Any such function can be represented as MOD_m^A for some $A \subseteq [m]$, defined as follows.

$$\mathsf{MOD}_m^A(x_1, \ldots, x_n) = \begin{cases} 1 & \sum_i x_i \in A \\ 0 & \sum_i x_i \notin A \end{cases}$$

We write MOD_m^A as a disjoint OR of $\mathsf{MOD}_m^{\{a\}}$ for each $a \in A$. Each $\mathsf{MOD}_m^{\{a\}}$ over n inputs can be rewritten as MOD_m over at most $2n$ inputs with some dummy 1 inputs. Then, for each $\mathsf{MOD}_m^{\{a\}}$ gate, we obtain a symmetric torus polynomial $P_{\{a\}}$ approximating it within $\frac{1}{20n^2}$ error using Lemma 5.2.

Finally, note that MOD_m^A is a disjoint OR over $\mathsf{MOD}_m^{\{a\}}$. Hence, in the sum $\sum_{a\in A} P_{\{a\}}$, at most one polynomial contributes a fractional part of $\frac{1}{2}$. Moreover, each $P_{\{a\}}$ contributes an error within $\frac{1}{20n^2}$, yielding a total error of $\frac{|A|}{20n^2} \leq \frac{1}{20n}$. Therefore, $\sum_{a\in A} P_{\{a\}}$ is a symmetric torus polynomial approximating MOD_m^A within an error of $\frac{1}{20n}$.

The degree of this polynomial is as claimed in the statement. This completes the proof. □

Discussion and Open Problems

One immediate direction is to make progress on Conjecture 4.1. In this direction, we present evidence that the proof of [KV26, Theorem 6] presented by the authors has room for improvement. Formally, the statement we prove improves the degree lower bound by an additive factor of 1. The improvement is more conceptual, even though the improvement on the degree bound is minor. It highlights that there is room for improvement in their proof, and further analysis may lead to a resolution of Conjecture 4.1. Note that we borrow heavily from the notation setup in [KV26] for brevity.

Theorem 5.1. *Any torus polynomial approximating* AND *within an error smaller than ε must have degree more than $\log(\frac{1}{\varepsilon}) - 1$.*

Proof. Fix some n, d, denote $f = \mathsf{AND}_n$, and choose any $\varepsilon < \frac{1}{2^{d+1}}$. We use the method described in [KV26, Theorem 8] for the proof. To prove the lower bound, we need to find a witness $\gamma \in nullspace(M(n,d))$ for each $Z : \{0,1\}^n \to \mathbb{Z}$ such that:

$$\left|\left\langle Z + \frac{f}{2}, \gamma \right\rangle\right| > \varepsilon \|\gamma\|_1 \tag{1}$$

First, we use [KV26, Lemma 9], which implies that $Z(a) = 0$ for any $a \in \{0,1\}^n$ with $|a| \leq d$. Now, choose any point with $|a| = d + 1$. Construct a vector $\gamma \in nullspace(M(n,d))$ using [KV26,

Construction 2] with $S_1 = \emptyset, S_2 = a$ and $I = [d+1]*$ as the input to the construction. For this vector, we have $\|\gamma\|_1 = 2^{d+1}$, $\gamma_a = 1$, and $\gamma_{a'} = 0$ for any point $|a'| \geq d+1$. Hence, $\left\langle Z + \frac{f}{2}, \gamma \right\rangle = Z(a)\gamma_a = Z_a$, whereas $\varepsilon\|\gamma\|_1 < 1$. Therefore, if $|Z(a)| \geq 1$, we have that inequality 1 is satisfied. This leaves us to consider $Z(a) = 0$ as the only remaining possibility. Note that this argument applies to each point a with $|a| = d+1$ independently to show that $Z(a) = 0$ for that point.

We continue with the argument inductively, stopping before we reach Hamming weight n, assuming that $Z(a') = 0$ for all points with $|a'| = i$ for some $i \geq d+1$. Then, consider a point a with $|a| = i+1$, and invoke [KV26, Construction 2] with $S_1 = a'', S_2 = a$ and $I = [d+1]^*$ for some $a'' \subseteq a$ of size $|a''| = |a| - (d+1)$. Again, as in the base case, if $|Z(a)| = 1$, then inequality 1 is satisfied. Hence, $Z(a) = 0$ is the only possibility that remains.

Finally, once we reach Hamming weight n with $a = 1^n$, we have $Z(a') = 0$ for any point with $|a'| < n$. Now, as the integer parts up to Hamming weight $n-1$ are 0, we can think of the torus polynomial as a real polynomial approximating the following function g:

$$g(a) = \begin{cases} 0 & a \neq 1^n \\ Z(a) + \frac{1}{2} & a = 1^n \end{cases}$$

However, any real polynomial that approximates g within ε error also approximates AND within at most ε error. As the real approximation degree of AND for ε error is $\Omega\left(\sqrt{n \log(\frac{1}{\varepsilon})}\right)$ [BT21], such a polynomial cannot exist. This completes the proof. □

The conceptual contribution of the proof above is to use multiple vectors to witness the lower bound. In [KV26], the authors had used a single witness for each Z to prove their result. Hence, this can be thought of as a step towards further improvements.

In another direction, the proof technique of Theorem 3.1 works for any subgroup $\Gamma \leq S_n$ which simplifies to a large enough copy of the symmetric group under two restrictions: setting variables equal to each other, and setting variables as 1. In particular, consider Γ that collapses to $S_{n'}$ for some n' that is superpolynomial in $\log(n)$. Then, using our arguments, one obtains a superpolynomial lower bound on the size of Γ-symmetric CC^0 circuits computing AND. We leave it open to describe such subgroups $\Gamma \leq S_n$.

Open Problem 1. *Describe subgroups $\Gamma \leq S_n$ such that under the following restrictions:*

- *setting variables equal to each other,*
- *setting variables as 1,*

the group collapses to $S_{n'}$ for some $n' = \log^{\omega(1)}(n)$.

Acknowledgments: We would like to thank anonymous MFCS 2026 reviewers for their helpful and detailed comments.

A Appendix: Proof of Lemma 3.5

Proof. Let T be a polynomial such that $T(x) = \mathsf{MOD}_q(x) \mod q^{\ell'}$ for some ℓ' to be chosen later. Here, T is constructed by combining Fermat-Euler theorem and modulus amplifying polynomials [BT94], which has degree $O(\ell')$. Define $Q = \left(w\left\lceil \frac{2^\ell}{q^{\ell'}} \right\rceil + 2^\ell\right)T$ with $w = -q^{\ell'} \mod 2^e$.

The calculations are the same as those communicated to us by the authors in [BHLR19]. We reproduce them for the sake of completeness. Fix some x, and let $T(x) = aq^{\ell'} + \mathsf{MOD}_q(x)$. Now, we continue the calculation as follows:

$$\begin{aligned} Q(x) &= \left(w\left\lceil \frac{2^\ell}{q^{\ell'}} \right\rceil + 2^\ell\right)T(x) \\ &= \left(w\left(\frac{2^\ell}{q^{\ell'}} + 1 - \left\{\frac{2^\ell}{q^{\ell'}}\right\}\right)\right)\left(aq^{\ell'} + \mathsf{MOD}_q(x)\right) \\ &= \left(wa2^\ell + q^{\ell'}a2^\ell\right) + 2^\ell \mathsf{MOD}_q(x) + waq^{\ell'} - waq^{\ell'}\left\{\frac{2^\ell}{q^{\ell'}}\right\} + w\left\lceil \frac{2^\ell}{q^{\ell'}} \right\rceil \mathsf{MOD}_q(x) \end{aligned}$$

Now, note that $\left(wa2^\ell + q^{\ell'}a2^\ell\right) \mod 2^{\ell+e} = 0$. Moreover, $|aq^{\ell'}| \leq 2^{O(q\ell' \log(n))}$ and $w \leq 2^e$. Hence, if we ensure the following:

$$w\left(aq^{\ell'}\left(1 - \left\{\frac{2^\ell}{q^{\ell'}}\right\}\right) + \left\lceil \frac{2^\ell}{q^{\ell'}} \right\rceil\right) \leq 2^{\ell-e}$$

Then, $E(x) \leq 2^{\ell-e}$, as required by the statement. Choosing a large enough $\ell' = O(e)$ and $\ell = O(qe\log(n))$ suffices to ensure that the calculation goes through. □

References

[Aar08] Scott Aaronson. The polynomial method in quantum and classical computing. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, page 3, USA, 2008. IEEE Computer Society.

[AKV20] Noga Alon, Mrinal Kumar, and Ben Lee Volk. Unbalancing sets and an almost quadratic lower bound for syntactically multilinear arithmetic circuits. *Comb.*, 40(2):149–178, 2020.

[BBR94] David A. Mix Barrington, Richard Beigel, and Steven Rudich. Representing boolean functions as polynomials modulo composite numbers. *Comput. Complex.*, 4:367–382, 1994.

[BHLR19] Abhishek Bhrushundi, Kaave Hosseini, Shachar Lovett, and Sankeerth Rao. Torus polynomials: An algebraic approach to ACC lower bounds. In Avrim Blum, editor, *10th*

Innovations in Theoretical Computer Science Conference, ITCS 2019, San Diego, California, USA, January 10-12, 2019, volume 124 of *LIPIcs*, pages 13:1–13:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[BST90] David A Mix Barrington, Howard Straubing, and Denis Thérien. Non-uniform automata over groups. *Information and Computation*, 89(2):109–132, 1990.

[BT94] Richard Beigel and Jun Tarui. On ACC. *Comput. Complex.*, 4:350–366, 1994.

[BT21] Mark Bun and Justin Thaler. Guest column: Approximate degree in classical and quantum computing. *SIGACT News*, 51(4):48–72, January 2021.

[CGPT06] Arkadev Chattopadhyay, Navin Goyal, Pavel Pudlák, and Denis Thérien. Lower bounds for circuits with mod_m gates. In *47th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2006, Berkeley, California, USA, October 21-24, 2006, Proceedings*, pages 709–718. IEEE Computer Society, 2006.

[HRRY19] Pavel Hrubeš, Sivaramakrishnan Natarajan Ramamoorthy, Anup Rao, and Amir Yehudayoff. Lower bounds on balancing sets and depth-2 threshold circuits. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, Patras, Greece, July 9-12, 2019*, LIPIcs, pages 72:1–72:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[HS19] Prahladh Harsha and Srikanth Srinivasan. On polynomial approximations to AC. *Random Struct. Algorithms*, 54(2):289–303, 2019.

[Kri21] Vaibhav Krishan. Upper bound for torus polynomials. In Rahul Santhanam and Daniil Musatov, editors, *Computer Science - Theory and Applications - 16th International Computer Science Symposium in Russia, CSR 2021, Sochi, Russia, June 28 - July 2, 2021, Proceedings*, Lecture Notes in Computer Science, pages 257–263. Springer, 2021.

[KV26] Vaibhav Krishan and Sundar Vishwanathan. Lower bounds and separations for torus polynomials. In Shubhangi Saraf, editor, *17th Innovations in Theoretical Computer Science Conference, ITCS 2026, Bocconi University, Milan, Italy, January 27-30, 2026*, volume 362 of *LIPIcs*, pages 88:1–88:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026.

[KW25] Piotr Kawałek and Armin Weiß. Violating constant degree hypothesis requires breaking symmetry. In Olaf Beyersdorff, Michal Pilipczuk, Elaine Pimentel, and Kim Thang Nguyen, editors, *42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025, Jena, Germany, March 4-7, 2025*, LIPIcs, pages 58:1–58:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.

[LSS+21] Nutan Limaye, Karteek Sreenivasaiah, Srikanth Srinivasan, Utkarsh Tripathi, and S. Venkitesh. A fixed-depth size-hierarchy theorem for $\mathrm{AC}^0[\oplus]$ via the coin problem. *SIAM J. Comput.*, 50(4):1461–1499, 2021.

[OSS19] Igor Carboni Oliveira, Rahul Santhanam, and Srikanth Srinivasan. Parity helps to compute majority. In Amir Shpilka, editor, *34th Computational Complexity Conference, CCC 2019, New Brunswick, NJ, USA, July 18-20, 2019*, LIPIcs, pages 23:1–23:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[Pag26a] Benedikt Pago. Optimal lower bounds for symmetric modular circuits. In *53rd EATCS International Colloquium on Automata, Languages, and Programming ICALP 2026, London, UK, July 7-10, 2026*, 2026. To appear.

[Pag26b] Benedikt Pago. Optimal lower bounds for symmetric modular circuits, 2026. arXiv preprint arXiv:2604.04760.

[Raz87] Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 77–82. ACM, 1987.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il