Electronic Colloquium on Computational Complexity, Report No. 15 (2019)

QMA Lower Bounds for Approximate Counting

William Kretschmer*

Abstract

We prove a query complexity lower bound for QMA protocols that solve approximate counting: estimating the size of a set given a membership oracle. This gives rise to an oracle A such that $\mathsf{SBP}^A \not\subset \mathsf{QMA}^A$, resolving an open problem of Aaronson [2]. Our proof uses the polynomial method to derive a lower bound for the SBQP query complexity of the AND of two approximate counting instances. We use Laurent polynomials as a tool in our proof, showing that the "Laurent polynomial method" can be useful even for problems involving ordinary polynomials.

1 Introduction

Among counting complexity classes, the complexity class SBP captures approximate counting: estimating a #P function within a constant multiplicative factor. Despite having a definition in terms of counting complexity, SBP is known to lie between two interactive proof classes. In particular, Bohler et. al. [4], who defined SBP, showed that $\mathsf{MA} \subseteq \mathsf{SBP} \subseteq \mathsf{AM}$. Thus, under plausible derandomization assumptions [10], one would have $\mathsf{NP} = \mathsf{MA} = \mathsf{SBP} = \mathsf{AM}$.

In this work, we study the relation between SBP and QMA. The containment $\mathsf{SBP} \subseteq \mathsf{QMA}$ would follow trivially if SBP collapses to MA, but it is unclear whether quantum Merlin makes proving this containment any easier. In the relativized world, Aaronson [2] recently asked whether there might exist an oracle A relative to which $\mathsf{SBP}^A \not\subset \mathsf{QMA}^A$. He noted that exhibiting such an oracle is equivalent to ruling out a black box QMA protocol for approximate counting. We formally define the approximate counting problem as follows:

Problem 1. *The approximate counting problem* $\mathsf{ApxCount}_{N,w}$ *is: given a membership oracle for a set* $A \subseteq [N] = \{1, 2, \ldots, N\}$ *promised that either* $|A| \leq w$ *("no" instance) or* $|A| \geq 2w$ *("yes" instance), determine which of these is the case.*

More generally, one can consider the problem of distinguishing $|A| \leq w$ or $|A| \geq (1+\epsilon)w$ where ϵ is an arbitrary constant, or may even depend on N and w. However, we restrict our attention to fixed ϵ because SBP precisely captures approximate counting in the case where A is the set of accepting paths of a nondeterministic polynomial-time Turing machine (and so $|A|$ is a #P function), w is an FP function, and $\epsilon = 1$. Thus, an SBP-QMA oracle separation would follow if for some function $w(N)$, any QMA protocol for $\mathsf{ApxCount}_{N,w}$ requires either a $(\log N)^{\omega(1)}$-size witness, or else $(\log N)^{\omega(1)}$ queries.

To prove such a lower bound, we study the query complexity of $\mathsf{ApxCount}_{N,w}$ in the context of the complexity class SBQP, a quantum analogue of SBP first defined by Kuperberg [8]. SBQP is

*University of Texas at Austin. Email: kretsch@cs.utexas.edu. Supported by a Simons Investigator award.

ISSN 1433-8092

in some sense the smallest "natural" complexity class that contains both SBP and QMA. Indeed, just as $\mathsf{MA}^A \subseteq \mathsf{SBP}^A$ for any oracle A, so is $\mathsf{QMA}^A \subseteq \mathsf{SBQP}^A$ for any oracle A.

Note that one cannot hope to prove a nontrivial SBQP query complexity lower bound for approximate counting, as the SBQP query complexity of $\mathsf{ApxCount}_{N,w}$ is $O(1)$ for any N and w. Instead, we use the observation that SBP is not obviously closed under intersection[1]. In this light, we consider the analogous intersection problem $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ wherein we are given a pair of sets $A_0, A_1 \subseteq [N]$ and asked to determine whether both sets have size at least $2w$, or whether one of the sets has size at most w[2]. Because QMA is closed under intersection, a QMA protocol for $\mathsf{ApxCount}_{N,w}$ that receives a witness of size $(\log N)^{O(1)}$ and makes $(\log N)^{O(1)}$ queries implies (via in-place amplification) the existance of an SBQP algorithm for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ that makes $(\log N)^{O(1)}$ queries.

Our main result is that no such SBQP algorithm exists. Specifically, we show that any SBQP algorithm for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ requires $\Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$ queries. We also modify this argument to show that $\Omega(w)$ queries are necessary when $N = 2^{\Omega(w)}$. This in turn shows that any QMA protocol for $\mathsf{ApxCount}_{N,w}$ that receives a witness of size m and makes T queries must satisfy $m \cdot T = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$ (or $m \cdot T = \Omega(w)$ when $N = 2^{\Omega(w)}$). Our proof uses the celebrated polynomial method of Beals et. al. [3]: for an algorithm that makes T queries, we construct a bivariate polynomial $p(x, y)$ of degree at most $2T$ that equals the probability that the algorithm accepts on a random $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ instance where $|A_0|$ and $|A_1|$ are of fixed size. We then show that if the algorithm is an SBQP algorithm that correctly solves $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$, then any such polynomial must have large degree.

In our view, the proof of this degree lower bound (Theorem 11) is of independent mathematical interest. At a high level, from this polynomial $p(x, y)$, we take a parametric curve through the xy plane to construct a univariate Laurent polynomial $q(t)$ of the same degree[3]. Crucially, we leverage the symmetries of the problem to view this Laurent polynomial as an ordinary univariate polynomial of the same degree. Finally, we appeal to classical results in approximation theory to argue that this univariate polynomial must have large degree. We find this application of Laurent polynomials surprising, particularly because the recent result of Aaronson on the BQP query complexity of approximate counting in the QSamples+queries model *also* used Laurent polynomials, albeit for an entirely different reason [2]. For Aaronson's result, Laurent polynomials are fundamentally necessary just to describe the acceptance probability of the algorithm, while in our case ordinary polynomials suffice. This suggests that the "Laurent polynomial method" may prove to be useful even for problems involving ordinary polynomials.

[1] There even exists an oracle relative to which SBP is *not* closed under intersection [7], and SBP's closure or non-closure under intersection in the unrelativized world remains an open problem.

[2] As a technicality, we typically assume that both sets satisfy the $\mathsf{ApxCount}_{N,w}$ promise, though strictly speaking only the smaller set needs to satisfy the promise on a "no" instance of $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$.

[3] A Laurent polynomial $q(t)$ can contain both positive and negative integer powers of t. Formally, we can write $q(t) = q_+(t) + q_-(1/t)$ where q_+ and q_- are ordinary polynomials. We follow the convention that the degree of a Laurent polynomial $q(t)$ is $\deg q = \max\{\deg q_+, \deg q_-\}$. This is for consistency with the definition of degree for multivariate polynomials, as in the polynomial $q_+(x) + q_-(y)$ (i.e. viewing $x = t$ and $y = 1/t$ as separate indeterminates).

2 Preliminaries

Though SBP and SBQP can be defined in terms of counting complexity functions (as above), for our purposes it is easier to work with the following equivalent definitions (see Böhler et. al. [4]):

Definition 2. *The complexity class* SBP *consists of the languages L for which there exists a probabilistic polynomial time algorithm M and a polynomial σ with the following properties:*

1. *If $x \in L$, then $\Pr[M(x)\ accepts] \geq 2^{-\sigma(|x|)}$.*
2. *If $x \notin L$, then $\Pr[M(x)\ accepts] \leq 2^{-\sigma(|x|)}/2$.*

The complexity class SBQP *is defined analogously, wherein the classical algorithm is replaced with a quantum algorithm.*

A classical (respectively, quantum) algorithm that satisfies the above promise for a particular language will be referred to as an SBP (respectively, SBQP) algorithm throughout this paper. Using this definition, a tight query complexity relation between QMA protocols and SBQP algorithms follows from the procedure of Marriott and Watrous [9], which shows that one can exponentially improve the soundness and completeness errors of a QMA protocol without increasing the witness size (see Aaronson [1] for a proof of the following lemma):

Lemma 3 (Guessing lemma)**.** *Suppose V^A is a* QMA *verifier for some problem and that V^A makes T queries to an oracle A, receives an m-qubit witness, and has soundness and completeness errors $1/3$. Then there is an* SBQP *algorithm Q^A for the same problem that receives no witness and makes $O(m \cdot T)$ queries.*

Because we study oracle intersection problems, it is often convenient to think of an algorithm as having access to *two* oracles, wherein the first bit in the oracle register selects the choice of oracle. As a consequence, we need a slight generalization of a now well-established fact in quantum complexity: that the acceptance probability of a quantum algorithm with an oracle can be expressed as a polynomial in the bits of the oracle string.

Lemma 4 (Symmetrization with two oracles)**.** *Suppose Q^{A_0,A_1} is a quantum algorithm that makes T queries to a pair of oracles $A_0, A_1 \subseteq [N]$. Then there exists a bivariate real polynomial $p(x,y)$ of degree at most $2T$ such that:*

$$p(x,y) = \mathop{\mathbb{E}}_{\substack{|A_0|=x,\\ |A_1|=y}} \left[\Pr[Q^{A_0,A_1}\ accepts]\right]$$

for all $x, y \in [N]$.

Proof. We can equivalently view the oracles as strings in $\{0,1\}^N$ such that the algorithm makes queries to a single oracle $A = A_0|A_1$ which is the concatenation of the two oracles. Then, Lemma 4.2 of Beals et. al. [3] tells us that there exists a real polynomial $r(A)$ of degree at most $2T$ such that $r(A) = r(A_0, A_1) = \Pr[Q^{A_0,A_1}\ \text{accepts}]$ for any $A \in \{0,1\}^{2N}$ that is a string of $\{0,1\}$ variables. We then apply the symmetrization lemma of Minsky and Papert [11] to symmetrize r, first with respect to A_0, then with respect to A_1:

$$p_0(x, A_1) = \mathop{\mathbb{E}}_{|A_0|=x} r(A_0, A_1) = \mathop{\mathbb{E}}_{|A_0|=x} \left[\Pr[Q^{A_0,A_1}\ \text{accepts}]\right]$$

$$p(x,y) = \mathop{\mathbb{E}}_{|A_1|=y} p_0(x, A_1) = \mathop{\mathbb{E}}_{\substack{|A_0|=x,\\ |A_1|=y}} \left[\Pr[Q^{A_0,A_1} \text{ accepts}]\right]$$

□

We now state some useful facts from approximation theory that will be useful in our proofs. We start with the Markov brothers' inequality:

Lemma 5 (Markov)**.** *Let p be a real polynomial of degree d, and suppose that:*

$$\max_{x,y\in[a,b]} |p(x) - p(y)| \leq H.$$

Then for all $x \in [a,b]$, the derivative p' satisfies:

$$|p'(x)| \leq \frac{H}{b-a} d^2.$$

This lemma has a useful consequence:

Corollary 6. *Let p be a real polynomial of degree d, and suppose that $|p(x)| \leq 1$ for all integers $x \in \{0, 1, \ldots, k\}$. If $\max_{x\in[0,k]} |p(x)| \geq 1.001$, then $d = \Omega(\sqrt{k})$.*

Proof. Without loss of generality, we may scale p by some constant and choose x so that $|p(x)| = 1.001$ is the maximum absolute value of $p(x)$ on $[0,k]$. By the mean value theorem, there exists some $x^* \in [\lfloor x \rfloor, \lceil x \rceil]$ such that $|p'(x^*)| \geq 0.001$. Applying the previous lemma, we find that:

$$0.001 \leq \frac{2 \cdot 1.001}{k} d^2$$
$$\sqrt{\frac{0.001}{2.002}} k \leq d.$$

□

Put another way, if a polynomial is bounded at all integers $\{0, 1, \ldots, k\}$ and has degree $o(\sqrt{k})$, then the polynomial satisfies a marginally weaker bound on all of $[0,k]$. We might wonder whether we can still assume some nontrivial bound when d is not so much smaller than k. Indeed we can:

Lemma 7 (Coppersmith and Rivlin [6])**.** *Let p be a real polynomial of degree $d \leq k$, and suppose that $|p(x)| \leq 1$ for all integers $x \in \{0, 1, \ldots, k\}$. Then there exist constants a, b that do not depend on d or k such that for all $x \in [0,k]$, we have:*

$$|p(x)| \leq a \cdot \exp\left(bd^2/k\right).$$

We will also use a bound as stated by Paturi [12] that bounds a polynomial in terms of its degree and a bound on a nearby interval:

Lemma 8. *Let p be a real polynomial of degree d, and suppose that $|p(x)| \leq 1$ for all $|x| \leq 1$. Then for all x with $|x| \leq 1 + \mu$, we have:*

$$|p(x)| \leq \exp\left(2d\sqrt{2\mu + \mu^2}\right).$$

Setting $\mu = \frac{1}{k}$ and performing some computation gives rise to the following:

Corollary 9. *Let p be a real polynomial of degree d, and suppose that $|p(x)| \leq 1$ for all $|x| \leq 1$. If $|p(1 + 1/k)| \geq 1.001$, then $d = \Omega(\sqrt{k})$.*

Finally, we state a useful fact about Laurent polynomials:

Lemma 10 (Symmetric Laurent polynomials)**.** *Let $\ell(x)$ be a real Laurent polynomial of degree d that satisfies $\ell(x) = \ell(1/x)$. Then there exists a real polynomial q of degree d such that $\ell(x) = q(x + 1/x)$.*

Proof. $\ell(x) = \ell(1/x)$ implies that the coefficients of the x^i and x^{-i} terms are equal for all i, as otherwise $\ell(x) - \ell(1/x)$ would not equal the zero polynomial. Thus, we may write $\ell(x) = \sum_{i=0}^{d} a_i \cdot (x^i + x^{-i})$ for some coefficients a_i. So, it suffices to show that $x^i + x^{-i}$ can be expressed as a polynomial in $x + 1/x$ for all $0 \leq i \leq d$.

We prove by induction on i. The case $i = 0$ corresponds to constant polynomials. For $i > 0$, by the binomial theorem, observe that $(x + 1/x)^i = x^i + x^{-i} + r(x)$ where r is a degree $i - 1$ real Laurent polynomial satisfying $r(x) = r(1/x)$. By the induction assumption, r can be expressed as a polynomial in $x + 1/x$, so we have $x^i + x^{-i} = (x + 1/x)^i - r(x)$ is expressed as a polynomial in $x + 1/x$. □

3 Main Result

3.1 Lower Bound for SBQP

Our results hinge on the following theorem, which uses Laurent polynomials to prove a degree lower bound for bivariate polynomials that satisfy a particular set of bounds at points in the plane:

Theorem 11. *Let w and N be integers with $0 < w < 2w \leq N$. Let $R_x = [2w, N] \times [0, w]$ and $R_y = [0, w] \times [2w, N]$ be disjoint rectangles in the plane, and let $L = R_x \cup R_y$. Let $p(x, y)$ be a real polynomial of degree d with the following properties:*

1. *$p(2w, 2w) \geq 2$.*
2. *$0 \leq p(x, y) \leq 1$ for all $(x, y) \in L \cap \mathbb{Z}^2$.*

Then $d = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$.

We remark that L gets this name because it looks like the letter "L", albeit with the bottom left corner missing (see Figure 1, shaded regions). The proof idea is as follows. First, we argue that either $d = \Omega(\sqrt{w})$, or else p satisfies a marginally weaker bound on the rectangle R_x by applying the Markov brothers' inequality (via Corollary 6) to horizontal and vertical lines through R_x. In the latter case, we show that taking an appropriate curve that passes through R_x and the point $(2w, 2w)$ gives rise to a univariate Laurent polynomial ℓ of degree d. We use Lemma 10 for symmetric Laurent polynomials to reinterpret this as an ordinary polynomial q of degree d. We then show that q is bounded on a large interval and grows quickly outside that interval, which implies (by Corollary 9) that q has degree $\Omega(\sqrt{N/w})$.

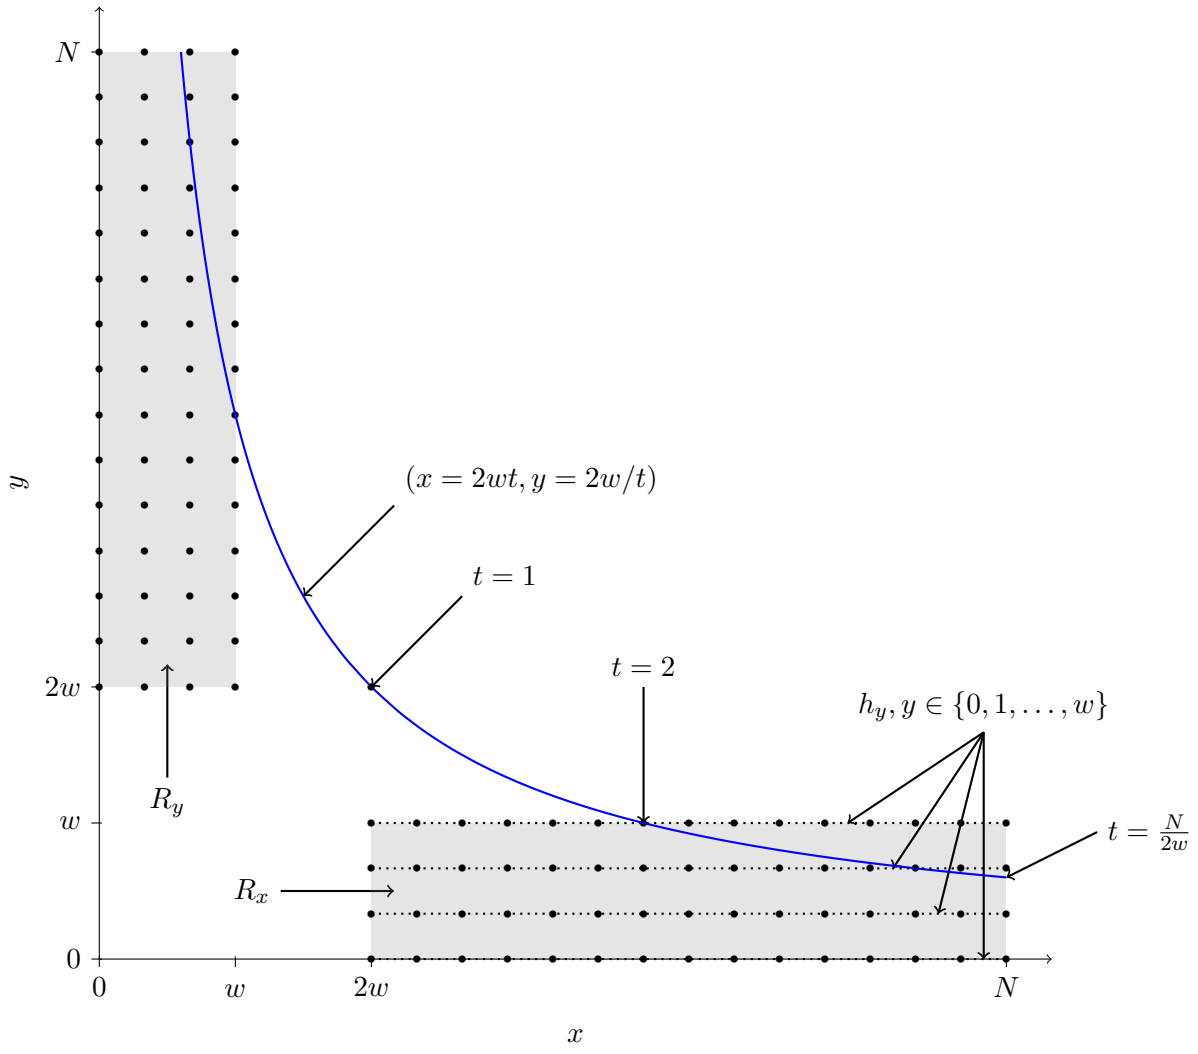

Figure 1: Diagram of Theorem 11. The lattice points $L \cap \mathbb{Z}^2$ where $0 \leq p(x, y) \leq 1$ are plotted. Not shown: vertical lines $v_x = x \times [0, w]$ through R_x for each $x \in [2w, N]$ (there are infinitely many such lines).

Proof of Theorem 11. We assume that $3w < N$, as otherwise $\sqrt{N/w} = O(1)$ and the theorem holds trivially.

Let $y \in \{0, 1, \ldots, w\}$ be an integer, and consider a horizontal line segment $h_y = [2w, N] \times y$ that passes through R_x (see Figure 1, dotted horizontal lines). The restriction of p to h_y gives rise to a univariate polynomial $p_y(x)$ of degree d. By the assumed bounds on $p(x, y)$ at lattice points in L, $|p_y(x)| \leq 1$ for all integers $x \in \{2w, 2w + 1, \ldots, N\}$. By the assumption $3w < N$, the interval $[2w, N]$ has length at least w. So, we may apply Corollary 6 to p_y to conclude that either $d = \Omega(\sqrt{w})$, or else $|p(x, y)| < 1.001$ for all $(x, y) \in h_y$.

Now, we use the bounds along the horizontal integer lines through R_x to get bounds along vertical lines. Let $x \in [2w, N]$ (*not necessarily* an integer), and consider a vertical line segment $v_x = x \times [0, w]$ that passes through R_x. The restriction of p to v_x gives rise to a univariate polynomial $p_x(y)$ of degree d. The intersection of v_x with the h_y's gives a bound $|p_x(y)| < 1.001$ for all integers $y \in \{0, 1, \ldots, w\}$. So, we may apply Corollary 6 to $p_x/1.001$ to conclude that either $d = \Omega(\sqrt{w})$, or else $|p(x, y)| < 1.001^2$ for all $(x, y) \in v_x$. Because every point (x, y) in the rectangle R_x lies on some v_x, we conclude that $|p(x, y)| < 1.001^2$ for all $(x, y) \in R$.

Observe that if $p(x, y)$ satisfies the statement of the theorem, then so does $p(y, x)$. This is because the constraints in the statement of the theorem are symmetric in x and y (in particular, because R_x and R_y are mirror images of one another along the line $x = y$; see Figure 1). As a result, we may assume without loss of generality that p is symmetric, i.e. $p(x, y) = p(y, x)$. Else, we may replace p by $\frac{p(x,y)+p(y,x)}{2}$ because the set of polynomials that satisfy the inequalities in the statement of the theorem are closed under convex combinations.

Consider the parametric curve $(x = 2wt, y = 2w/t)$ as it passes through R_x (see Figure 1). We can view the restriction of $p(x, y)$ to this curve as a Laurent polynomial $\ell(t) = p(2wt, 2w/t)$ of degree d. The bound of $p(x, y)$ on all of R_x implies that $|\ell(t)| < 1.001^2$ when $t \in [2, \frac{N}{2w}]$ and that $\ell(1) \geq 2$ (see Figure 1). Moreover, the condition that $p(x, y)$ is symmetric implies that $\ell(t) = \ell(1/t)$.

By Lemma 10 for symmetric Laurent polynomials, $\ell(t)$ can be viewed as a degree d polynomial $q(t+1/t)$. Under the transformation $s = t+1/t$, q satisfies $|q(s)| < 1.001^2$ for $s \in [2+1/2, \frac{N}{2w} + \frac{2w}{N}]$ and $q(2) \geq 2$. Note that the length of the interval $[2 + 1/2, \frac{N}{2w} + \frac{2w}{N}]$ is $\Theta(N/w)$ because $w < N$. By an appropriate affine transformation of q, we can conclude from Corollary 9 with $k = \Theta(N/w)$ that $d = \Omega(\sqrt{N/w})$. □

Theorem 11 implies an SBQP query complexity lower bound for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$:

Theorem 12. *Let Q^{A_0,A_1} be an SBQP algorithm for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ that makes T queries to membership oracles A_0 and A_1. Then $T = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$.*

Proof. Since Q is an SBQP algorithm, we may suppose that Q accepts with probability at least 2α on a "yes" instance and with probability at most α on a "no" instance. Using Lemma 4, take $p(x, y)$ to be the polynomial of degree at most $2T$ that satisfies:

$$p(x, y) = \mathop{\mathbb{E}}_{\substack{|A_0|=x,\\ |A_1|=y}} \left[\Pr[Q^{A_0,A_1} \text{ accepts}]\right].$$

Define $L' = ([0, w] \times [0, w]) \cup ([0, w] \times [2w, N]) \cup ([2w, N] \times [0, w])$. The conditions on the acceptance probability of Q^A for all A_0, A_1 that satisfy the $\mathsf{ApxCount}_{N,w}$ promise imply that $p(x, y)$ satisfies these corresponding conditions:

1. $1 \geq p(x,y) \geq 2\alpha$ for all $(x,y) \in ([2w,N] \times [2w,N]) \cap \mathbb{Z}^2$.

2. $0 \leq p(x,y) \leq \alpha$ for all $(x,y) \in L' \cap \mathbb{Z}^2$.

In particular, the polynomial $\frac{1}{\alpha} \cdot p(x,y)$ satisfies the (weaker) conditions of Theorem 11, from which it follows that $T = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$. □

We remark that even though we could assume $p(x,y) \geq 2\alpha$ over a large region, Theorem 11 only needed $p(x,y) \geq 2\alpha$ at a single point: $(x,y) = (2w, 2w)$. We view this as expressing the intuition that the acceptance probability of an SBQP algorithm "should" be increasing in $|A_0|$ and $|A_1|$.

3.2 (Non)-Tightness of SBQP Lower Bound

In this section, we compare our SBQP query complexity lower bound for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ to known upper bounds. We find a gap between these bounds, particularly when N is much larger than w. This motivates an approach to improving our lower bounds for large N. In Theorem 13, we prove that this approach indeed gives a better lower bound.

The best upper bound we know of for SBQP query complexity is $O\left(\min\left\{w, \sqrt{N/w}\right\}\right)$, so our bound is at least tight when $N = O(w^2)$. The $O(\sqrt{N/w})$ upper bound follows from the BQP algorithm of Brassard, Høyer, and Tapp [5]. The $O(w)$ upper bound is in fact an SBP upper bound with the following algorithmic interpretation: first, guess $w+1$ items randomly from each of A_0 and A_1. Then, verify using the membership oracle that the first $w+1$ items all belong to A_0 and that the latter $w+1$ items all belong to A_1, accepting if and only if this is the case. This accepts with nonzero probability if and only if $|A_0| \geq w+1$ and $|A_1| \geq w+1$.

Can the gap between the lower and upper bounds be improved? On the upper bound side, it is tempting to combine Grover search or Brassard-Høyer-Tapp approximate counting with the classical verification to get an $O(\sqrt{w})$ algorithm, but this fails in general because both algorithms always have some nonzero chance of accepting when the number of marked items is nonzero. This suggests that perhaps the lower bound is not tight, at least when $N \gg w$.

Looking for improvements on the lower bound side, careful observation reveals that the main bottleneck in the proof of Theorem 11 is the bound on the growth of polynomials bounded at equally spaced points (Corollary 6), which breaks down completely when the polynomial has degree $\omega(\sqrt{w})$. One might observe that we used Corollary 6 to bound $p(x,y)$ on all of R_x, even though we really just need a bound on $p(x,y)$ at the points $(x = 2wt, y = 2w/t)$.

In fact, this leads to an approach for improving the lower bound, which we now describe. At a high level, we might hope to bound $p(x,y)$ on $(x = 2wt, y = 2w/t)$ by observing that the curve approaches the line $y = 0$ as t grows large. When N is large enough, we *can* still conclude a bound on $p(x,0)$ for $(x,0) \in R$ using Corollary 6, and intuitively $p(x,y)$ should be close to $p(x,0)$ as $y \to 0$. This intuition indeed works, and allows us to conclude an $\Omega(w)$ lower bound when $N = 2^{\Omega(w)}$. Our strategy for proving this improved lower bound is to show that if $d = o(w)$, then there exists some $\epsilon > 0$ that depends only on w such that $|p(x,y)| < 1.002$ whenever $y \leq \epsilon$ and $(x,y) \in R$. Then, the curve $(x = 2wt, y = 2w/t)$ lies in this region whenever $\frac{2w}{\epsilon} \leq t \leq \frac{N}{2w}$. It follows that the polynomial $q(s)$ as in the proof of Theorem 11 satisfies $|q(s)| < 1.002$ for all $s \in [\frac{2w}{\epsilon} + \frac{\epsilon}{2w}, \frac{N}{2w} + \frac{2w}{N}]$ and $q(2) \geq 2$. So long as N satisfies $\frac{N}{2w} \geq w^2 \cdot \frac{2w}{\epsilon}$, the length of this interval is $\frac{2w}{\epsilon} \cdot \Omega(w^2)$. This gives a contradiction: an appropriate affine transformation of q satisfies the statement of Corollary 9 with $k = \Omega(w^2)$ but has degree $d = o(w)$. We conclude that $d = \Omega(w)$.

Theorem 13. *Let w, N, L, $p(x,y)$, and d satisfy the statement of Theorem 11. If $N = 2^{\Omega(w)}$, then $d = \Omega(w)$.*

Proof. Similar to the proof of Theorem 11, we first bound $p(x,y)$ on horizontal lines through R_x, but we can assume a better lower bound because N is now assumed to be large. Exactly as before, we let $y \in \{0,1,\ldots,w\}$ be an integer, and we consider a horizontal line segment $h_y = [2w, N] \times y$ that passes through R_x. The restriction of p to h_y gives rise to a univariate polynomial $p_y(x)$ of degree d. By the assumed bounds on $p(x,y)$ at lattice points in L, $|p_y(x)| \leq 1$ for all integers $x \in \{2w, 2w+1, \ldots, N\}$. But now, we can assume $N \gg w^2$, and so Corollary 6 implies that either $d = \Omega(w)$, or else $|p(x,y)| < 1.001$ for all $(x,y) \in h_y \cap L$.

This time, instead of using Corollary 6 to bound $p(x,y)$ on vertical lines through R_x, we start with the bound of Coppersmith and Rivlin (Lemma 7). As before, we let $x \in [2w, N]$ (*not necessarily* an integer), and we consider a vertical line segment $v_x = x \times [0,w]$ that passes through R_x. The restriction of p to v_x gives rise to a univariate polynomial $p_x(y)$ of degree d. The intersection of v_x with the h_y's gives a bound $|p_x(y)| < 1.001$ for all integers $y \in \{0,1,\ldots,w\}$.

Suppose for a contradiction that $d = o(w)$. Then Lemma 7 implies that $|p_x(y)| < 2^{o(w)}$ for all $(x,y) \in v_x$. From the Markov brothers' inequality (Lemma 5), we can assume that the derivative satisfies $|p'_x(y)| \leq \frac{2^{o(w)}}{w} \cdot o(w^2) \leq 2^{o(w)}$ for all $y \in [0,w]$. Because $|p_x(0)| < 1.001$, then by basic calculus, there exists $\epsilon = 2^{-o(w)}$ such that $p_x(y) < 1.002$ for all $y \leq \epsilon$. In particular, $|p(x,y)| < 1.002$ whenever $y \leq \epsilon$ and $(x,y) \in R$.

Recall that it sufficed to show $\frac{N}{2w} \geq w^2 \cdot \frac{2w}{\epsilon}$, or equivalently $N \geq \frac{4w^4}{\epsilon}$ to get a contradiction from the assumption $d = o(w)$. Because $\epsilon = 2^{-o(w)}$, this follows from the assumption that $N = 2^{\Omega(w)}$. We conclude that $d = \Omega(w)$. □

3.3 Lower Bound for QMA

We now prove two results about QMA complexity that follow from the SBQP lower bound of Theorem 12:

Corollary 14. *There exists an oracle A and a pair of languages L_0, L_1 such that:*

1. $L_0, L_1 \in \mathsf{SBP}^A$
2. $L_0 \cap L_1 \notin \mathsf{SBQP}^A$.
3. $\mathsf{SBP}^A \not\subset \mathsf{QMA}^A$.

Proof. For an arbitrary function $A : \{0,1\}^* \to \{0,1\}$ and $i \in \{0,1\}$, define $A_i^n = \{x \in \{0,1\}^n : A(i,x) = 1\}$. Define the unary language $L_i^A = \{1^n : |A_i^n| \geq 2^{n/2}\}$. Observe that as long as A satisfies the promise $|A_i^n| \geq 2^{n/2}$ or $|A_i^n| \leq 2^{n/2-1}$ for all $n \in \mathbb{N}$, then $L_i^A \in \mathsf{SBP}^A$. Intuitively, the oracles A that satisfy this promise encode a pair of $\mathsf{ApxCount}_{N,w}$ instances $|A_0^n|$ and $|A_1^n|$ for every $n \in \mathbb{N}$ where $N = 2^n$ and $w = 2^{n/2-1}$.

Theorem 12 tells us that an SBQP algorithm Q that makes $o(2^{n/4})$ queries fails to solve $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ on *some* pair (A_0, A_1) that satisfies the promise. Thus, one can construct an A such that $L_0, L_1 \in \mathsf{SBP}^A$ and $L_0 \cap L_1 \notin \mathsf{SBQP}^A$, by choosing (A_0^n, A_1^n) so as to diagonalize against all SBQP algorithms.

Because QMA^A is closed under intersection for any oracle A, and because $\mathsf{QMA}^A \subseteq \mathsf{SBQP}^A$ for any oracle A, it must be the case that either $L_0 \notin \mathsf{QMA}^A$ or $L_1 \notin \mathsf{QMA}^A$. □

We remark that this gives an alternative construction of an oracle relative to which SBP is not closed under intersection. To our knowledge, this is the first that uses the polynomial method directly.

Using the guessing lemma (Lemma 3), we can also place an explicit lower bound on the QMA complexity of $\mathsf{ApxCount}_{N,w}$:

Corollary 15. *Let V^A be* QMA *verifier for the $\mathsf{ApxCount}_{N,w}$ with soundness and completeness errors $1/3$. Suppose V^A receives a witness of length m and makes T queries to a set membership oracle A. Then $m \cdot T = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$.*

Proof. Running V a constant number of times with fresh witnesses to reduce the soundness and completeness errors, one obtains a verifier with soundness and completeness errors $1/6$ that receives an $O(m)$-length witness and makes $O(T)$ queries. Repeating twice with two oracles and computing the AND, one obtains a QMA verifier V'^{A_0,A_1} for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ with soundness and completeness errors $1/3$ that receives an $O(m)$-length witness and makes $O(T)$ queries. Applying the guessing lemma (Lemma 3) to V', there exists an SBQP algorithm Q^{A_0,A_1} for $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$ that makes $O(m \cdot T)$ queries. Theorem 12 tells us that $m \cdot T = \Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$. □

Alternatively, one can conclude that $m + T = \Omega\left(\min\left\{w^{1/4}, (N/w)^{1/4}\right\}\right)$. Furthermore, when $N = 2^{\Omega(w)}$, one can instead conclude that $m \cdot T = \Omega(w)$ and therefore $m + T = \Omega(\sqrt{w})$ using Theorem 13 in place of Theorem 11 in the proof of Theorem 12.

4 Discussion and Open Problems

The QMA lower bound for $\mathsf{ApxCount}_{N,w}$ is not optimal in general: when $w = O(1)$, there is no QMA protocol for $\mathsf{ApxCount}_{N,w}$ that receives a constant size witness and makes a constant number of queries for large N. Fundamentally, this shows that SBQP lower bounds cannot give optimal QMA lower bounds. However, our SBQP bounds themselves are not tight: can one improve the gap between the $\Omega\left(\min\left\{\sqrt{w}, \sqrt{N/w}\right\}\right)$ lower bound and $O\left(\min\left\{w, \sqrt{N/w}\right\}\right)$ upper bound for the SBQP query complexity of $\mathsf{AND}_2 \circ \mathsf{ApxCount}_{N,w}$? From Theorem 13, we know that the complexity must *eventually* reach $\Omega(w)$ (at least when N is exponentially large), so it seems reasonable to conjecture that the $\Omega(\sqrt{w})$ lower bound is not tight even for smaller values of N. On the other hand, we have also not made a serious attempt to improve the trivial $O(w)$-query SBP algorithm using Grover search (or similar techniques). Thus, it appears entirely possible that *neither* bound is tight, perhaps depending on N.

At a deeper level, we would like to know if there is any meaningful connection between our use of Laurent polynomials and their use by Aaronson [2] in studying the QSamples+queries model. One way we might hope to establish such a connection is to extend our proof to an SBQP lower bound in the QSamples+queries model by proving an analogue of Theorem 11 for Laurent polynomials. We remark that the argument in Theorem 11 that turns a symmetric bivariate polynomial $p(x, y) = p(y, x)$ into a univariate Laurent polynomial $\ell(t) = \ell(1/t)$ works just as well if $p(x, y)$ is a symmetric bivariate Laurent polynomial. Thus, essentially all that is needed is an analogue of Corollary 6 for Laurent polynomials bounded on a set of equally spaced points[4].

[4] A priori, it might appear that Theorem 11 also breaks down for large N: as t grows, the curve $(x = 2wt, y = 2w/t)$

5 Acknowledgements

I would like to thank Thomas Watson for suggesting the intersection approach to proving an SBP-QMA oracle separation. I am also grateful to Scott Aaronson who supported this research, introduced me to this particular problem, and provided valuable guidance. I would also like to thank Robin Kothari, Justin Thaler, and Patrick Rall for their helpful feedback on this writing.

References

[1] Scott Aaronson. "Impossibility of Succinct Quantum Proofs for Collision-freeness". In: *Quantum Info. Comput.* 12.1-2 (Jan. 2012), pp. 21–28.

[2] Scott Aaronson. *Quantum Lower Bound for Approximate Counting Via Laurent Polynomials*. 2018. eprint: arXiv:1808.02420.

[3] Robert Beals et al. "Quantum Lower Bounds by Polynomials". In: *J. ACM* 48.4 (July 2001), pp. 778–797. DOI: 10.1145/502090.502097.

[4] Elmar Böhler, Christian Glaßer, and Daniel Meister. "Error-bounded probabilistic computations between MA and AM". In: *Journal of Computer and System Sciences* 72.6 (2006), pp. 1043 –1076. DOI: 10.1016/j.jcss.2006.05.001.

[5] Gilles Brassard, Peter Høyer, and Alain Tapp. "Quantum Counting". In: *Proceedings of the 25th International Colloquium on Automata, Languages and Programming*. ICALP '98. Berlin, Heidelberg: Springer-Verlag, 1998, pp. 820–831.

[6] Don Coppersmith and T. J. Rivlin. "The Growth of Polynomials Bounded at Equally Spaced Points". In: *SIAM J. Math. Anal.* 23.4 (July 1992), pp. 970–983. DOI: 10.1137/0523054.

[7] Mika Göös et al. "Rectangles Are Nonnegative Juntas". In: *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*. STOC '15. Portland, Oregon, USA: ACM, 2015, pp. 257–266. DOI: 10.1145/2746539.2746596.

[8] Greg Kuperberg. "How Hard Is It to Approximate the Jones Polynomial?" In: *Theory of Computing* 11.6 (2015), pp. 183–219. DOI: 10.4086/toc.2015.v011a006.

[9] Chris Marriott and John Watrous. "Quantum Arthur—Merlin Games". In: *Comput. Complex.* 14.2 (June 2005), pp. 122–152. DOI: 10.1007/s00037-005-0194-x.

[10] P. B. Miltersen and N. V. Vinodchandran. "Derandomizing Arthur-Merlin games using hitting sets". In: *40th Annual Symposium on Foundations of Computer Science*. Oct. 1999, pp. 71–80. DOI: 10.1109/SFFCS.1999.814579.

[11] M. Minsky, S.A. Papert, and L. Bottou. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1969.

[12] Ramamohan Paturi. "On the Degree of Polynomials That Approximate Symmetric Boolean Functions". In: *Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing*. STOC '92. New York, NY, USA: ACM, 1992, pp. 468–474. DOI: 10.1145/129712.129758.

approaches $y = 0$, where any Laurent polynomial with negative degree terms is clearly unbounded. However, the $\Omega(\sqrt{N/w})$ part of the lower bound really only applies when $N = O(w^2)$, so for the ranges of N that we care about, it is sufficient to look at $t \in [1, 2w]$ over which $y \geq 1$.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il