Electronic Colloquium on Computational Complexity, Report No. 41 (2015)

Dual Polynomials for Collision and Element Distinctness

Mark Bun*
mbun@seas.harvard.edu

Justin Thaler†
jthaler@fas.harvard.edu

Abstract

The approximate degree of a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ is the minimum degree of a real polynomial that approximates f to within error $1/3$ in the ℓ_∞ norm. In an influential result, Aaronson and Shi (J. ACM 2004) proved tight $\tilde{\Omega}(n^{1/3})$ and $\tilde{\Omega}(n^{2/3})$ lower bounds on the approximate degree of the Collision and Element Distinctness functions, respectively. Their proof was non-constructive, using a sophisticated symmetrization argument and tools from approximation theory.

More recently, several open problems in the study of approximate degree have been resolved via the construction of dual polynomials. These are explicit dual solutions to an appropriate linear program that captures the approximate degree of any function. We reprove Aaronson and Shi's results by constructing explicit dual polynomials for the Collision and Element Distinctness functions.

*Harvard University, School of Engineering and Applied Sciences. Supported by an NDSEG Fellowship and NSF grant CNS-1237235.

†Yahoo Labs. Parts of this work were performed while the author was a Research Fellow at the Simons Institute for the Theory of Computing. Supported in part by a Research Fellowship from the Simons Institute for the Theory of Computing.

ISSN 1433-8092

1 Introduction

The ε-approximate degree of a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ is the least degree of a real polynomial that approximates f to within error ε in the ℓ_∞ norm. Approximate degree is a fundamental measure of the complexity of a Boolean function, and has wide-ranging applications in theoretical computer science. For example, approximate degree upper bounds underly several of the best known algorithms for PAC learning [24], agnostic learning [22,23], learning in the presence of irrelevant information [25,31], and differentially private data release [19,44]. Meanwhile, lower bounds on approximate degree imply many optimal lower bounds on quantum query complexity, circuit complexity, and communication complexity (see for example [4,10–12,16,33–35,40]).

In an influential result, Aaronson and Shi proved tight $\tilde{\Omega}(n^{1/3})$ and $\tilde{\Omega}(n^{2/3})$ lower bounds on the approximate degree of the Collision and Element Distinctness functions [4].[1] The Collision lower bound matched an earlier $O(n^{1/3})$ upper bound due to Brassard et al. [15], while the lower bound for Element Distinctness was later shown to be tight by Ambainis [8].

The Collision lower bound subsequently found many applications and extensions in quantum complexity theory; Aaronson recently provided a retrospective overview of these developments [3]. Moreover, the $\tilde{\Omega}(n^{2/3})$ lower bound for Element Distinctness remains the best known approximate degree lower bound for any function in AC^0.

Aaronson and Shi proved their lower bound for Collision with a symmetrization argument. This style of argument proceeds in two steps. First, a polynomial p on n variables (which is assumed to approximate the target function f) is transformed into a polynomial q on $m < n$ variables in such a way that $\deg(q) \leq \deg(p)$. Second, a lower bound on $\deg(q)$ is proved, typically by applying Markov-Bernstein type inequalities from approximation theory. Aaronson and Shi's proof of the Collision lower bound is a particularly sophisticated application of this style of argument.

The lower bound for Element Distinctness follows from a reduction to the lower bound for Collision. This reduction is discussed in Section 5.

The Method of Dual Polynomials. Despite the many applications of approximate degree in theoretical computer science, significant gaps remain in our understanding of this complexity measure, and there are many simple functions whose approximate degree remains unknown. The slow nature of progress can be attributed in part to the limitations of symmetrization arguments. At an intuitive level, the process of symmetrization is inherently lossy: by turning a polynomial p on n variables into a polynomial q on $m < n$ variables, information about p is necessarily thrown away. Hence, several works have identified that an important research direction is to develop techniques beyond symmetrization for lower bounding the approximate degree of Boolean functions [2,17,38].

The last few years have seen significant progress toward this goal. In particular, a series of works has proved new approximate degree lower bounds for important classes of functions by constructing explicit *dual polynomials*, which are dual solutions to a certain linear program capturing the approximate degree of any function. These polynomials act as certificates of the high approximate degree of a function. Moreover, strong LP duality implies that the technique is lossless, in contrast to symmetrization. That is, for any function f and any ε, there is always some dual polynomial ϕ that witnesses a tight approximate degree lower bound for f; the challenge is to construct ϕ.

This "method of dual polynomials" was recently used to resolve the approximate degree of the AND-OR tree [17,37], closing a long line of incrementally larger lower bounds [6,21,28,38,39]. It has also been used

[1] Aaronson established a lower bound of $\tilde{\Omega}(n^{1/5})$ for the Collision function in a paper that appeared in STOC 2002 [1], and Shi improved it to the tight $\tilde{\Omega}(n^{1/3})$ in a FOCS paper that same year [39]. A joint journal paper appeared in 2004 [4]. The proof was simplified and extended to the "small range" case by Kutin [26]. Ambainis [6] independently extended Aaronson and Shi's lower bound to the small range case, using different techniques than Kutin.

to establish several "hardness amplification" results for approximate degree [18, 32, 43], and to prove new *threshold degree* lower bounds for several important classes of functions, including the intersection of two majorities [29, 38] and AC^0 [32]. The latter result represented the first superlogarithmic improvement over Minsky and Papert's seminal $\Omega(n^{1/3})$ lower bound from 1969 on the threshold degree of an AC^0 function. We also note that dual polynomials have recently been used to resolve several longstanding open problems in communication complexity, where they yield explicit distributions under which various communication problems are hard (see the survey of Sherstov [33]).

Contribution and motivation. We reprove Aaronson and Shi's results by constructing explicit dual polynomials for the Collision and Element Distinctness functions.[2] First, we give a direct construction of a dual polynomial for Collision. In Section 2.5, we give an overview of the ideas that go into this construction. We then show how to turn any dual polynomial ψ for Collision into a dual polynomial φ for Element Distinctness. We construct $\varphi(x)$ by averaging $\psi(y)$ over a carefully constructed set of extensions from each x to a longer input y.

We have two main motivations for reproving Aaronson and Shi's lower bound in this manner. First, only a handful of techniques are currently known for the construction of dual polynomials, especially for the case where $\varepsilon = \Theta(1)$. To date, dual polynomials have been constructed only for symmetric functions [17, 41] and a handful of highly structured block-composed functions [17, 18, 32, 36–38] (a block-composed function $F\colon \{-1,1\}^{M \cdot N} \to \{-1,1\}$ is a function of the form of the form $F = g(f(x_1), \ldots, f(x_M))$ for some $g\colon \{-1,1\}^M \to \{-1,1\}$ and $f\colon \{-1,1\}^N \to \{-1,1\}$). The Collision and Element Distinctness functions fall into neither category; our constructions of dual polynomials for these problems introduce several new techniques that we hope will prove useful in future applications.

A second motivation is to shed new light on the Collision lower bound itself. The earlier symmetrization-based proof [4, 26], while shorter than ours, is non-constructive and relies on Markov-Bernstein inequalities from approximation theory. In contrast, our proof is constructive and entirely elementary. We also believe that our analysis illuminates some of the more miraculous aspects of the earlier symmetrization-based proof – see Section 2.6 for further discussion of this point.

Related work on quantum query complexity. Aaronson and Shi's original motivation for studying the approximate degree of the Collision function was to understand its quantum query complexity (recall that approximate degree provides a *lower bound* on quantum query complexity [10]. However, it is known that the lower bound is not always tight [7]). Subsequent to Aaronson and Shi's work, other methods were developed for quantum query complexity [5, 7, 9, 20, 42, 47], and it is now known that one of these methods, called the *negative-weights adversary method* [20], is always tight.

The negative-weights adversary method for lower bounding quantum query complexity is closely analogous to the method of dual polynomials for approximate degree: the former is characterized by a semidefinite program, and a solution to this semidefinite program is known as an *adversary matrix*. A recent line of work, similar in spirit to our own, has proved or reproved optimal quantum query complexity lower bounds for several functions by constructing explicit adversary matrices. In particular, Belovs and Rosmanis [13] constructed an optimal adversary matrix for the Collision function in the "large range" case (note that the dual polynomial that we construct applies even in the "small range" case), and Belovs and Špalek constructed an optimal adversary matrix for the Element Distinctness function [14].

Very recently, Zhandry [46] (improving on work of Yuen [45]) has also proved a tight lower bound of $\Omega(N^{1/3})$ on the quantum query complexity of finding a collision in a randomly chosen function.

[2] Like Kutin's simplification and refinement of Aaronson and Shi's original proof of the Collision lower bound, our construction yields a dual polynomial for the Collision function even in the "small-range" case.

2 Preliminaries

2.1 Notation

For any positive integer n, we denote the set $\{1, \ldots, n\}$ by $[n]$, and the set $\{0, 1, \ldots, n\}$ by $[[n]]$. For a function $f : D \to \mathbb{R}$, define the L_1 norm $\|f\|_1 = \sum_{x \in D} |f(x)|$. For any subset $S \subseteq [n]$, we let $\chi_S \colon \{-1,1\}^n \to \{-1,1\}$ denote the parity function on S, i.e., $\chi_S(x) = \prod_{i \in S} x_i$.

2.2 Approximate Degree and its Dual Characterization

Let $D \subseteq \{-1,1\}^n$, and let $f \colon D \to \{-1,1\}$ be a partial Boolean function defined on D. A real polynomial $p \colon \{-1,1\}^n \to \{-1,1\}$ is said to ε-approximate f if

1. $|p(x) - f(x)| \leq \varepsilon$ for all $x \in D$, and
2. $|p(x)| \leq 1 + \varepsilon$ for all $x \in \{-1,1\}^n$.

The ε-approximate degree of f, denoted $\widetilde{\deg}_\varepsilon(f)$, is the minimum degree of an ε-approximation for f. We use $\widetilde{\deg}(f)$ to denote $\widetilde{\deg}_{1/3}(f)$, and refer to this quantity without qualification as the *approximate degree* of f. The choice of $1/3$ is arbitrary, as $\widetilde{\deg}(f)$ is related to $\widetilde{\deg}_\varepsilon(f)$ by a constant factor for any constant $\varepsilon \in (0,1)$.

Given a partial Boolean function f, let p be a real polynomial that attains the smallest ε subject to the constraints above, over all polynomials of degree at most d. Since we work over $x \in \{-1,1\}^n$, we may assume without loss of generality that p is multilinear with the representation $p(x) = \sum_{|S| \leq d} c_S \chi_S(x)$, where the coefficients c_S are real numbers. Then p is an optimum of the following linear program.

min	ε	
such that	$\left\lvert f(x) - \sum_{\lvert S\rvert \leq d} c_S \chi_S(x) \right\rvert \leq \varepsilon$	for each $x \in D$
	$\left\lvert \sum_{\lvert S\rvert \leq d} c_S \chi_S(x) \right\rvert \leq 1 + \varepsilon$	for each $x \in \{-1,1\}^n \setminus D$
	$c_S \in \mathbb{R}$	for each $\lvert S\rvert \leq d$
	$\varepsilon \geq 0$	

The dual linear program is as follows.

max	$\sum_{x \in D} \phi(x) f(x) - \sum_{x \in \{-1,1\}^n \setminus D} \lvert\phi(x)\rvert$	
such that	$\sum_{x \in \{-1,1\}^n} \lvert\phi(x)\rvert = 1$	
	$\sum_{x \in \{-1,1\}^n} \phi(x) \chi_S(x) = 0$	for each $\lvert S\rvert \leq d$
	$\phi(x) \in \mathbb{R}$	for each $x \in \{-1,1\}^n$

Strong LP-duality thus implies the following dual characterization of approximate degree:

Theorem 1. *Let $f : D \to \{-1,1\}$ be a partial Boolean function. Then $\widetilde{\deg}_\varepsilon(f) > d$ if and only if there is a polynomial $\phi \colon \{-1,1\}^n \to \mathbb{R}$ such that*

$$\sum_{x \in D} f(x)\phi(x) - \sum_{x \in \{-1,1\}^n \setminus D} |\phi(x)| > \varepsilon \cdot \sum_{x \in \{-1,1\}^n} |\phi(x)|, \tag{1}$$

and

$$\sum_{x \in \{-1,1\}^n} \phi(x)\chi_S(x) = 0 \text{ for each } |S| \leq d. \tag{2}$$

If ϕ satisfies Eq. (1), we say that ϕ has *correlation greater than* ε with f. If ϕ satisfies Eq. (2), we say ϕ has *pure high degree* d. We refer to any feasible solution ϕ to the dual linear program as an (ε, d)-*dual polynomial* for f.

2.3 The Collision and Element Distinctness Functions

Let $[N] = \{1, \ldots, N\}$, and fix a triple of positive integers n, N, R such that $R \geq N$, and $n = N \cdot \log_2 R$. For simplicity throughout, we assume that R is a power of 2. The Collision and Element Distinctness functions are typically thought of as *properties* of functions mapping $[N]$ to $[R]$. However, it will be convenient for us to think of them instead as functions on the Boolean hypercube $\{-1, 1\}^n$. To this end, given an input $x \in \{-1, 1\}^n$, we interpret x as the evaluations of a function g_x mapping $[N] \to [R]$. That is, we break x up into N blocks, each of length $\lceil \log_2 R \rceil$, and regard each block x_i as the binary representation of $g_x(i)$.

Definition 2 (Collision Function)**.** *A function $g_x \colon [N] \to [R]$ is said to be k-to-1 if for every $i \in [N]$, there exists exactly $k-1$ values $j \neq i$ such that $g_x(i) = g_x(j)$. Let $T_k := \{x \in \{-1,1\}^n : g_x$ is k-to-$1\}$ (clearly, T_k is non-empty only if $k|N$). The Collision function, which we denote by $\mathrm{Col}_{N,R}$, is the partial Boolean function defined on $T_1 \cup T_2 \subseteq \{-1,1\}^n$ such that $\mathrm{Col}_{N,R}(x) = 1$ if and only if $x \in T_1$. That is, $\mathrm{Col}_{N,R}$ is the partial Boolean function corresponding to the property that g_x is a 1-to-1 function, with the promise that g_x is either 1-to-1 or 2-to-1.*

Definition 3 (Element Distinctness Function)**.** *The* Element Distinctness function*, denoted $\mathrm{ED}_{N,R}$, is the total Boolean function defined such that $\mathrm{ED}_{N,R}(x) = 1$ if and only if g_x is 1-to-1. That is, $\mathrm{ED}_{N,R}$ is the total Boolean function corresponding to the property that g_x is 1-to-1.*

Let $B \subset \{-1,1\}^n$ denote the set of inputs x such that g_x is neither 1-to-1 nor 2-to-1. Then an (ε, d)-dual polynomial ϕ for $\mathrm{Col}_{N,R}$ has the following properties (cf. Section 2.2):

1. $\sum_{x \in T_1} \phi(x) - \sum_{x \in T_2} \phi(x) - \sum_{x \in B} |\phi(x)| > \varepsilon \cdot \sum_{x \in \{-1,1\}^n} |\phi(x)|$.
2. $\sum_{x \in \{-1,1\}^n} \phi(x) \chi_S(x) = 0$ for all $|S| \leq d$.

Similarly, an (ε, d)-dual polynomial for $\mathrm{ED}_{N,R}$ satisfies:

1. $\sum_{x \in T_1} \phi(x) - \sum_{x \notin T_1} \phi(x) > \varepsilon \cdot \sum_{x \in \{-1,1\}^n} |\phi(x)|$.
2. $\sum_{x \in \{-1,1\}^n} \phi(x) \chi_S(x) = 0$ for all $|S| \leq d$.

2.4 Overview of the Symmetrization-Based Proof of the Collision Lower Bound

Kutin's simplified proof of the Collision lower bound [26] proceeds in two steps. The first step is a symmetrization step, which establishes the following remarkable result (we state this result slightly informally in this overview).

Lemma 4 (Informal version of Lemma 5)**.** *Call a triple (m, a, b)* valid *if $a|m$ and $b|(N-m)$. For any triple (m, a, b), let $R_{m,a,b}$ denote the set of inputs $x \in \{-1,1\}^n$ such that $g_x \colon [N] \to [R]$ maps m of its inputs to $[R]$ in an a-to-1 manner, and maps the remaining $N - m$ of its inputs to $[R]$ in a b-to-1 manner. Then there is a trivariate polynomial P of total degree at most d such that for every valid triple (m, a, b), it holds that $P(m, a, b) = \mathbb{E}_{x \in R_{m,a,b}}[p(x)]$.*

Note in the above lemma that the sets $R_{m,a,b}$ are not uniquely determined; for instance $R_{m,1,1} = R_{0,a,1} = R_{N,1,b} = T_1$ for every triple (m, a, b).

The second step of Kutin's proof argues that if p is a $(1/3)$-approximating polynomial for the Collision function, then P must have degree $\Omega(N^{1/3})$. Hence by Lemma 4, p must have degree $\Omega(N^{1/3})$ as well.

In more detail, the second step of Kutin's proof proceeds via a case analysis. Four cases are considered.

- The first is: $P(N/2, 1, 2) \geq 1/2$, and $|P(N/2, 1, b)| \leq 2$ for all $b \in [N^{2/3}]$. In this case, Kutin is able to apply Markov's inequality from approximation theory to conclude that the degree of P in its third variable is $\Omega(N^{1/3})$.

- The second is: $P(N/2, 1, 2) \geq 1/2$, and $|P(N/2, 1, b)| > 2$ for some $b \in [N^{2/3}]$. In this case, Kutin is able to apply Bernstein's inequality from approximation theory to conclude that the degree of P in its first variable is $\Omega(N^{1/3})$.

- The third is: $P(N/2, 1, 2) < 1/2$, and $|P(N/2, a, 2)| \leq 2$ for all $a \in [N^{2/3}]$. In this case, Kutin is able to apply Markov's inequality to conclude that the degree of P in its second variable is $\Omega(N^{1/3})$.

- The fourth is: $P(N/2, 1, 2) < 1/2$, and $|P(N/2, a, 2)| > 2$ for some $a \in [N^{2/3}]$. In this case, Kutin is able to apply Bernstein's inequality to conclude that the degree of P in its first variable is $\Omega(N^{1/3})$.

A key technical complication that must be dealt with in the argument above is that $|P(m, a, b)|$ may be much larger than 1 for *invalid* triples (m, a, b). This may seem like a minor technicality, but in fact it is a central issue: if $P(m, a, b)$ were bounded for all invalid triples, then it would be possible to argue that the total degree of P is $\Omega(N^{1/2})$, which would imply a (false) lower bound of $\Omega(N^{1/2})$ on the approximate degree of $\mathrm{Col}_{N,R}$.

2.5 Overview of Our Construction for the Collision Function

Like Kutin's proof, our construction also makes essential use of Lemma 4. Whereas Kutin used Lemma 4 to reduce to a setting where Markov-Bernstein inequalities could be applied in a non-constructive manner, we instead use Lemma 4 to argue that the dual polynomial ϕ that we construct has pure high degree $\Omega(N^{1/3})$.

In more detail, we present our construction in two stages, in order to highlight distinct ideas that go into the proof. In the first stage, we construct a simpler dual polynomial $\phi\colon \{-1,1\}^n \to \{-1,1\}$ that exhibits an $\Omega(\sqrt{\log N / \log\log N})$ lower bound on the approximate degree of $\mathrm{Col}_{N,R}$. The second stage constructs a dual polynomial ψ exhibiting the optimal $\Omega(N^{1/3})$ lower bound.

Overview of the first stage. Let $H_k \subseteq \{-1,1\}^n$ denote the set of inputs of Hamming weight n. The symmetrization-based proof of the Collision lower bound from [4, 26] carries the strong intuition that the sets T_k should play the same role that H_k plays in Nisan and Szegedy's seminal symmetrization-based lower bound for the OR function [28]. We direct the interested reader to Aaronson's lecture notes [27] for a detailed explanation of this intuition. The construction of our simpler dual witness ϕ instantiates this intuition in the dual setting.

Recall that a dual polynomial ϕ witnessing the fact that $\widetilde{\deg}(\mathrm{Col}_{N,R}) \geq d$ must satisfy two properties: (1) it must have correlation greater than ε with $\mathrm{Col}_{N,R}$, and (2) it must have pure high degree at least d. We define ϕ in a way that mimics the structure of known dual witnesses for symmetric functions, even though ϕ is not itself symmetric. Specifically, our construction ensures that the analysis establishing Properties (1) and (2) becomes similar to the analyses of known dual polynomials for the OR function [17, 41].

In more detail, our prior work [17] built on work of Špalek [41] to give a dual witness γ for the fact that $\widetilde{\deg}_\varepsilon(\mathrm{OR}_n) = \Omega(\sqrt{n})$ for any constant $\varepsilon < 1$; moreover, γ places non-zero weight only on sets H_k, for

values of k equal (up to scaling factors) to perfect squares. The pure high degree of γ is shown to be equal to (at least) the number of sets H_k upon which γ places non-zero weight.

Call an input $x \in \{-1,1\}^n$ *valid* if it is in $R_{m,a,b}$ for some valid triple (m,a,b). By analogy with γ, the dual witness ϕ that we construct in Stage 1 places weight only on inputs $x \in T_k$ for divisors k of N that are also (up to scaling factors) perfect squares. In particular, our definition of ϕ ensures that:

$$\phi(x) = 0 \text{ for all invalid inputs } x. \tag{3}$$

We are able to combine Eq. (3) with Lemma 4 and a basic combinatorial identity (cf. Lemma 9) to show that the pure high degree of ϕ is at least $|S|$, where S denotes the set of T_k's upon which ϕ places non-zero weight. Moreover, our definition of ϕ is carefully chosen to ensure that its correlation with $\mathrm{Col}_{N,R}$ is large: the precise calculation is closely analogous to the analysis from [17, 41] showing that γ is well-correlated with the OR function [17, 41].

Overview of the second stage. In the second stage, we construct a dual polynomial ψ that exhibits the optimal $\Omega(N^{1/3})$ lower bound. Rather than only weighting inputs in T_k for some some divisors k of n, ψ weights inputs in $R_{m,a,b}$ for many valid triples (m,a,b). There are two key ideas that go into the construction of ψ.

The first idea is to define ψ as the sum of two simpler dual polynomials ψ_1 and ψ_2, each with pure high degree $\Omega(N^{1/3})$ – then the sum ψ also has pure high degree $\Omega(N^{1/3})$ (see Lemma 15). The first polynomial ψ_1 places a large constant fraction (close to $1/2$) of its L_1 mass on T_1, whereas ψ_2 places a large constant fraction of its L_1 mass on T_2. Neither ψ_1 nor ψ_2 is well-correlated with $\mathrm{Col}_{N,R}$ in the sense of Eq. (3). However, they each place a constant fraction of their L_1 mass on $R_{N/2,2,1}$, and they are designed so that their values exactly cancel out on inputs in $R_{N/2,2,1}$. This allows us to show that $\psi = \psi_1 + \psi_2$ satisfies Eq. (3), even though ψ_1 and ψ_2 individually do not.

The second idea goes into the construction of ψ_1 and ψ_2 themselves. Specifically, we think of ψ_1 and ψ_2 as each being constructed in a two-step process. We focus on ψ_1 in this discussion, since the construction of ψ_2 is similar. Very roughly speaking, in the first step, we consider a "polynomial" ψ' of pure high degree $\Omega(N^{1/3})$ that places a large constant fraction of its L_1 mass on T_1; the construction of ψ' is closely related to our construction of the simpler dual polynomial ϕ from Stage 1.

The reason we place the term "polynomial" in quotes above is that there is an important technical caveat to our construction of ψ': we think of ψ' as placing weight on sets $R_{N/2,a,1}$ for many *invalid* triples $(N/2,a,1)$, in addition to some valid ones. Of course, if $(N/2,a,1)$ is invalid, then $R_{N/2,a,1} = \emptyset$, so ψ' cannot place non-zero weight on the set. To address this issue, in Step 2, we add to ψ' a bunch of polynomials $\psi''_{N/2,a,1}$, each of pure high degree $\Omega(N^{1/3})$. For each invalid triple $(N/2,a,1)$, $\psi''_{N/2,a,1}$ is specifically constructed to cancel out the weight that ψ' "places" on $R_{N/2,a,1}$.

Analogously to how our constructions of ϕ and ψ' were closely related to the dual witness for OR constructed in our earlier work [17], our construction of $\psi''_{N/2,a,1}$ is closely related to a dual witness η for the Majority function, MAJ, that we constructed in the same work. Each $\psi''_{N/2,a,1}$ places additional non-zero mass on (non-empty) sets of the form $R_{m,a,1}$ for some $a \neq 1$ and $m \in [N]$, but we are able to show that the total mass placed on such sets is small, using an analysis closely related to the analysis of η from [17]. Hence we are able to show that $\psi_1 = \psi' + \sum_{\text{invalid triples } (N/2,a,1)} \psi''_{N/2,a,1}$ still places a large constant fraction of its L_1 mass on T_1.

2.6 Discussion

On Kutin's second step. Our construction of the optimal dual witness ψ for the Collision function mimics the second step of Kutin's symmetrization argument in three important ways described below. We find this

mimicry to be somewhat surprising – in our earlier work [17], we constructed an optimal dual polynomial for symmetric Boolean functions that bore little relation to Paturi's well-known symmetrization-based proof of the same result [30]. We believe that this mimicry sheds new light, or at least gives a new perspective, on why Kutin's proof takes the structure that it does.

Recall that the second step of Kutin's proof (cf. Section 2.4) proceeds via a case analysis. The first "branch" in the case analysis depends on whether the expected value of the assumed n-variate approximation p to $\mathrm{Col}_{N,R}$ on the set $R_{N/2,2,1}$ is large or small. This is mimicked in our construction of ψ as a sum of two dual polynomials ψ_1 and ψ_2, both of which individually place a lot of weight on $R_{N/2,2,1}$, but whose sum places *zero* weight on $R_{N/2,2,1}$.

The second "branch" in Kutin's case analysis depends on whether $|P(N/2, a, 1)|$ or $|P(N/2, 2, b)|$ is small for all $a, b \leq N^{2/3}$. He needs to consider this second branch because $P(m, a, b)$ is not guaranteed to be bounded for invalid triples (m, a, b).

This branch is mimicked in our construction of ψ_1 (respectively, ψ_2) as the sum of a single "polynomial" ψ' that tries to place weight on sets $R_{N/2,a,1}$ for invalid triples $(N/2, a, 1)$ (respectively, $(N/2, 2, b)$), and many other polynomials $\psi''_{N/2,a,1}$ (respectively, $\psi''_{N/2,2,b}$), one for each invalid triple $(N/2, a, 1)$ (respectively, $(N/2, 2, b)$). In our dual setting, the reason we need to incorporate the polynomials $\psi''_{N/2,a,1}$ is to cancel out the weight that ψ' tries to place on invalid sets $R_{N/2,a,1}$.

Finally, recall that Kutin applied Markov's inequality from approximation theory in two of the four cases considered in his analysis, and Bernstein's inequality in the other two cases. Markov's inequality underlies Nisan and Szegedy's standard symmetrization-based proof that the approximate degree of OR is $\Omega(\sqrt{n})$ [28], while Berstein's inequality underlies Paturi's proof that the approximate degree of MAJ is $\Omega(n)$ [30]. This is mimicked in our construction of ψ_1 and ψ_2 as the sum of ψ' and the $\psi''_{N/2,a,1}$ and $\psi''_{N/2,2,b}$ polynomials: the construction of ψ' is closely analogous to the dual witness for OR from [17], while the construction of the $\psi''_{N/2,a,1}$ and $\psi''_{N/2,2,b}$ polynomials is based on the dual witness for MAJ from [17].

On the first step, or why k-to-1 inputs matter. As noted by several authors (e.g., [2, Slide 36]), the most miraculous element of the symmetrization-based proof of the Collision lower bound is the first step (cf. Lemma 4). The crux of this step is to establish, roughly speaking, that for any n-variate polynomial p of total degree d, the function $P(k) := \mathbb{E}_{x \in T_k}[p(x)]$ is a polynomial in k of degree at most d. Why should this hold? More basically, why should inputs that are k-to-1 even play a prominent role in the proof?

We provide some partial intuition for this in Section 6. Specifically, we explain that there is an (asymptotically) optimal approximation p for $\mathrm{Col}_{N,R}$ such that k-to-1 inputs correspond to constraints that are made tight by the solution corresponding to p in the primal linear program of Section 2.2. Hence, complementary slackness suggests that there should be a corresponding dual witness ψ that places weight only on inputs that are k-to-1, or nearly so, justifying the prominent role that k-to-1 inputs play in both the symmetrization-based proof and our new dual proof.

2.7 Formal Statement of Lemma 4

Following Kutin [26], we define a special collection of functions which are a-to-1 on one part of the domain and b-to-1 on the other part. For $N > 0$, recall that a triple of numbers (m, a, b) is *valid* if $a|m$ and $b|(N - m)$. For each valid triple (m, a, b), we define

$$g_{m,a,b}(i) = \begin{cases} \lceil i/a \rceil & \text{if } 1 \leq i \leq m \\ R - \lfloor (N-i)/b \rfloor & \text{if } m < i \leq n. \end{cases}$$

Moreover, for each valid triple (m, a, b), we define a set $R_{m,a,b}$ that is the orbit of $g_{m,a,b}$ under the

automorphism group $S_N \times S_R$. Namely,

$$x \in R_{m,a,b} \iff \exists \sigma \in S_N, \tau \in S_R\colon \quad \tau \circ g_x \circ \sigma = g_{m,a,b}.$$

Note that the sets $R_{m,a,b}$ are not uniquely determined; for instance $R_{m,1,1} = R_{0,a,1} = R_{N,1,b} = T_1$ for every m, a, b.

Lemma 5. *Let $p(x)$ be a real polynomial over $\{-1,1\}^n$ of degree d. There is a trivariate polynomial P of degree at most d with the property that for all valid triples (m,a,b),*

$$P(m,a,b) = \mathbb{E}_{x \in R_{m,a,b}}[p(x)].$$

The statement of Lemma 5 differs slightly from the corresponding lemma in Kutin's work [26] (Lemma 7 below). Lemma 5 follows by combining Kutin's formulation with the following simple lemma from [18].

Lemma 6 ([18])**.** *Let p be a polynomial over $\{-1,1\}^n$. Consider the map $T\colon \{-1,1\}^m \to \{0,1\}^{N \cdot R}$ defined by $T_{ij}(x) = 1$ if $g_x(i) = j$, and $T_{ij}(x) = 0$ otherwise. Then there is a polynomial $q\colon \{0,1\}^{N \cdot R} \to \mathbb{R}$ with $\deg q \leq \deg p$, such that $q(T(x)) = p(x)$ for all $x \in \{-1,1\}^n$.*

Lemma 7 ([26])**.** *Let $q(t)$ be any degree d polynomial in the variables t_{ij}. For a valid triple (m,a,b), define $Q(m,a,b)$ by*

$$P(m,a,b) = \mathbb{E}_{x \in R_{m,a,b}}[q(T(x))].$$

Then P is a degree d polynomial in m, a, b.

3 An $\Omega(\sqrt{\log N / \log\log N})$ Lower Bound for the Collision Function

The following lemma is a refinement of [17, Proposition 14], which was used there to construct a dual polynomial for OR.

Lemma 8. *There exists a constant $\zeta > 0$ such that for all $\delta \in (0,1)$ and $L \geq 1$, there is an explicit $\omega : \{1, \ldots, L\} \to \mathbb{R}$ with*

1. $\omega(1) \geq \frac{1-\delta}{2}$
2. $-\omega(2) \geq \frac{1-\delta}{2}$
3. $\sum_{k=1}^{L} |\omega(k)| = 1$
4. *For every polynomial $p : \{1, \ldots, L\} \to \mathbb{R}$ of degree $d \leq \zeta\sqrt{\delta L}$, we have $\sum_{k=1}^{L} p(k)\omega(k) = 0$.*

The proof will make use of the following simple combinatorial identity, a simple proof of which can be found in [29, Appendix A].

Lemma 9. *For any $L > 0$, let $q\colon \mathbb{R} \to \mathbb{R}$ be a univariate polynomial of degree strictly less than L. Then $\sum_{k=0}^{L} (-1)^k \binom{L}{k} q(k) = 0$.*

Proof of Lemma 8. Let $c = \lceil 16/\delta \rceil$. Let $m = \lfloor \sqrt{(L-1)/c} \rfloor$ and define the set

$$T = \{1\} \cup \{ci^2 \colon 0 \leq i \leq m\}.$$

Note that $|T| = \Omega(\sqrt{L/c})$. Define the function $\hat{\omega} : \{0, 1, \ldots, L\} \to \mathbb{R}$ by

$$\hat{\omega}(k) = \binom{L}{k} \frac{c^m (m!)^2}{L!} \prod_{j \in [[L]] \setminus T} (k - j) = \begin{cases} \dfrac{c^m (m!)^2}{\prod_{j \in T \setminus \{k\}} (k - j)} & \text{if } k \in T, \\ 0 \quad \text{otherwise.} \end{cases}$$

It is easy to check that $\hat{\omega}(0) = 1$.

For $k = 1$, we have $|\hat{\omega}(k)| = \frac{c^m (m!)^2}{\prod_{i=1}^m (ci^2 - 1)}$. Notice that:

$$\begin{aligned} 1 \leq \frac{c^m (m!)^2}{\prod_{i=1}^m (ci^2 - 1)} &= \prod_{i=1}^m \frac{i^2}{i^2 - 1/c} \\ &= \prod_{i=1}^m \left(1 + \frac{1}{ci^2 - 1}\right) \\ &\leq \exp\left(\sum_{i=1}^m \frac{2}{ci^2}\right) \\ &\leq e^{8/c} \leq 1 + \frac{3}{2}\delta. \end{aligned}$$

On the other hand, for $k = c\ell^2$ with $\ell > 0$, $|\hat{\omega}(k)|$ equals:

$$\begin{aligned} \frac{c^m (m!)^2}{(c\ell^2 - 1) \prod_{i \in [[m]] \setminus \{\ell\}} |c\ell^2 - ci^2|} &= \frac{(m!)^2}{(c\ell^2 - 1) \prod_{i \in [[m]] \setminus \{\ell\}} (i + \ell)|i - \ell|} \\ &= \frac{2(m!)^2}{(c\ell^2 - 1)(m + \ell)!(m - \ell)!} \\ &\leq \frac{2}{c\ell^2 - 1}, \end{aligned}$$

where the last inequality follows because

$$\frac{(m!)^2}{(m + \ell)!(m - \ell)!} = \frac{m}{m + \ell} \cdot \frac{m - 1}{m + \ell - 1} \cdot \ldots \cdot \frac{m - \ell + 1}{m + 1}$$

is a product of factors that are each smaller than 1. Thus, the total contribution of terms excluding 0 and 1 to the ℓ_1 mass of $\hat{\omega}$ is at most

$$\sum_{i=1}^m \frac{2}{ci^2 - 1} < \sum_{i=1}^\infty \frac{4}{ci^2} < \frac{8}{c} \leq \frac{\delta}{2}.$$

Now let $b = 0$ if $|T|$ is even, and $b = 1$ otherwise, and define $\omega \colon \{1, \ldots, L\} \to \mathbb{R}$ via:

$$\omega(k) = (-1)^{k+b} \hat{\omega}(k - 1) / \|\hat{\omega}\|_1.$$

Then

$$-\omega(2) \geq \omega(1) \geq \frac{1}{1 + |\hat{\omega}(1)| + \delta/2} \geq \frac{1}{2 + 2\delta} \geq \frac{1 - \delta}{2}.$$

This yields the first two claims about ω. The third claim follows immediately from the definition. Finally, let p be a polynomial of degree strictly less than $|T| - 1$. Then

$$\sum_{k=1}^L p(k)\omega(k) = \sum_{k=0}^{L-1} (-1)^k \cdot (-1)^b \cdot \binom{L}{k} \cdot \frac{c^m (m!)^2}{L! \|\hat{\omega}\|_1} \cdot p(k+1) \prod_{j \in [[L]] \setminus T} (k - j) = \sum_{k=0}^{L-1} (-1)^k \binom{L}{k} q(k), \quad (4)$$

where

$$q(k) = \frac{(-1)^b c^m (m!)^2}{L! \|\hat{\omega}\|_1} p(k+1) \prod_{j \in [[L]] \setminus T} (k-j)$$

is a polynomial of degree less than L. Since $q(L) = 0$, the right hand side of Eq. (4) is zero by Lemma 9. This gives the last claim. □

Our prior work [17], building on work of Špalek [41], obtained a dual polynomial γ for OR_L by setting the total weight of γ on inputs in H_k (the set of inputs of Hamming weight k) to be $\omega(k+1)$. In that work, the first three properties of ω ensured that γ had high correlation with OR, while the fourth ensured that it had pure high degree $\Omega(\sqrt{L})$.

Analogously, our dual polynomial ϕ for $\mathrm{Col}_{N,R}$ below sets the total weight of ϕ on T_k to be $\omega(k)$. Then again, the first three properties of ω ensure that ϕ is well-correlated with $\mathrm{Col}_{N,R}$, and the fourth ensures that it has pure high degree $\Omega(\sqrt{L})$. However, there is the complication that T_k must be non-empty, i.e., k must divide N, for every k in the support of ω. To handle this complication, we take N large enough so that all $k = 1, 2, \ldots, L$ divide N, yielding an $\Omega(\sqrt{\log N / \log\log N})$ lower bound.

Theorem 10. *Let $N = L!$ for some L. For $\delta > 0$, there exists an explicit $(1 - 2\delta, d)$ dual polynomial ϕ for $\mathrm{Col}_{N,R}$ with $d = \Omega(\sqrt{\delta L}) = \Omega(\sqrt{\delta \log N / \log\log N})$.*

Proof. First, notice that $k|N$ for all $k \in [L]$, so $T_k \neq \emptyset$ for every such k. Define $\phi(x) = \omega(k)/|T_k|$ if x is in T_k for some $k \in [L]$, and $\phi(x) = 0$ otherwise, where ω is obtained by applying Lemma 8. Note that $\phi(x)$ is well-defined since $|T_k| \neq 0$ for all $k \in [L]$, and each $x \in \{-1, 1\}^n$ is in T_k for at most one value of k.

We check:

$$\sum_{x \in T_1} \phi(x) - \sum_{x \in T_2} \phi(x) = \omega(1) - \omega(2) \geq 1 - \delta,$$

where the inequality holds by Parts 1 and 2 of Lemma 8. Moreover,

$$\sum_{x \in B} |\phi(x)| = \sum_{k=3}^{L} |\omega(k)| \leq \delta,$$

where the inequality holds by combining Parts 1-3 of Lemma 8. Thus,

$$\sum_{x \in T_1} \phi(x) - \sum_{x \in T_2} \phi(x) - \sum_{x \in B} |\phi(x)| \geq 1 - 2\delta.$$

Second,

$$\sum_{x \in T_1 \cup T_2 \cup B} |\phi(x)| = \sum_{k=1}^{L} |\omega(k)| = 1,$$

where the final equality holds by Part 3 of Lemma 8.

Finally, let $d = \zeta\sqrt{\delta L}$ where ζ is as in the statement of Lemma 8, and let $S \subseteq [n]$ with $|S| \leq d$. We must show that $\sum_{x \in T_1 \cup T_2 \cup B} \phi(x)\chi_S(x) = 0$. Note that:

$$\sum_{x \in T_1 \cup T_2 \cup B} \phi(x)\chi_S(x) = \sum_{k=1}^{L} \sum_{x \in T_k} \phi(x) \cdot \chi_S(x) = \sum_{k=1}^{L} \sum_{x \in T_k} (\omega(k)/|T_k|) \cdot \chi_S(x) = \sum_{k=1}^{L} \omega(k) \cdot \mathbb{E}_{x \in T_k}[\chi_S(x)],$$

where the first equality holds because $\phi(x) = 0$ if x is not in T_k for some $k \in [L]$.

By Lemma 5, there is a trivariate polynomial P of total degree at most d such that $P(m, a, b) = \mathbb{E}_{x \in R_{m,a,b}}[\chi_S(x)]$ for all valid triples (m, a, b). In particular, since $k|N$ for all $k \in [L]$, $q(k) := P(N, k, 1)$

is a *univariate* polynomial in k such that $q(k) = E_{x \in T_k}[\chi_S(x)]$ for all $k \in [L]$. Hence, Part 4 of Lemma 8 implies that

$$\sum_{k=1}^{L} \omega(k) \cdot \mathbb{E}_{x \in T_k}[\chi_S(x)] = 0.$$

□

4 An $\Omega(N^{1/3})$ Lower Bound for the Collision Function

The following lemma is a refinement of [18, Proposition 10], which constructed an explicit dual polynomial for MAJ.

Lemma 11. *There exists a constant $\rho > 0$ for which the following holds. Let $\delta \in (0,1)$, $N > 0$ an even integer, and $k \in [N]$. Then there is an explicit $\eta_k : [[N]] \to \mathbb{R}$ such that*

1. *η_k is supported on $\{2k, 4k, \ldots, 2\lfloor N/2k \rfloor k\} \cup \{N/2\}$*
2. *$\eta_k(N/2) > \frac{1-\delta}{2}$*
3. *$\sum_{r=0}^{N} |\eta_k(r)| = 1$*
4. *For every polynomial $p : \{0, \ldots, N\} \to \mathbb{R}$ of degree $d \leq \rho\sqrt{\delta}N/k$, we have $\sum_{r=0}^{N} p(r)\eta_k(r) = 0$.*

Proof. Throughout the proof, we assume for simplicity that $N/2$ is not a multiple of $2k$. The analysis when $N/2$ is a multiple of $2k$ is similar.

Let $c = \lceil \frac{10}{\sqrt{\delta}} \rceil$ and $t = 2\lfloor N/(4k) \rfloor k$ and define the set

$$S = \{t \pm 2c\ell k \colon 0 \leq \ell \leq \lfloor t/(2ck) \rfloor\}.$$

Note that $|S| = \Omega(N/ck)$. We claim that $\pi_S(i) := \prod_{j \in S, j \neq i} |j - i|$ is minimized at $i = t$. Notice that translating all points in S by a constant does not affect $\pi_S(i)$, and scaling all points in S by a constant does not affect $\operatorname{argmin}_i \pi_S(i)$. Thus, it is enough to show that $\pi_{S^*}(i)$ is minimized at $i = 0$ for the set $S^* = \{\pm\ell \colon \ell \leq t\}$. In this case, $\pi_{S^*}(i)$ takes the simple form $(t-i)!(t+i)!$, and we see that for all $i \in S^*$,

$$\frac{\pi_{S^*}(0)}{\pi_{S^*}(i)} = \frac{(t!)^2}{(t-i)!(t+i)!} = \frac{t}{t+|i|} \cdot \frac{t-1}{t+|i|-1} \cdot \cdots \cdot \frac{t-|i|+1}{t+1}$$

is a product of terms smaller than 1, so $\pi_{S^*}(i)$ is indeed minimized at $i = 0$.

Now let $T = S \cup \{t - 2k, N/2\}$ and define the function

$$\hat{\eta}(r) = \binom{N}{r} \frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{N!} \prod_{j \in [[N]] \setminus T} (r-j) = \frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{\prod_{j \in T \setminus \{r\}}(r-j)}$$

where $h = \lfloor t/2ck \rfloor$. The normalization is chosen so that $|\hat{\eta}(t)| = 1$.

The reason that we include *both* $(r - (t-2k))$ and $(r - (N/2))$ in the denominator of $\hat{\eta}$ is to ensure that the rate of decay of $\hat{\eta}(r)$ is at least quadratic as r moves away from t. This will ultimately allow us to show that a large fraction of the ℓ_1 mass of $\hat{\eta}$ comes from the point $r = N/2$.

For $r = t - 2k$, the mass $|\hat{\eta}(r)|$ is

$$\begin{aligned}\frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{2k(N/2-t+2k)\prod_{\ell=1}^{h}(2ck\ell-2k)(2ck\ell+2k)} &= \frac{(N/2-t)}{N/2-t+2k}\prod_{\ell=1}^{h}\left(1+\frac{1}{(c\ell)^2-1}\right)\\ &\le \frac{1}{2}\exp\left(\sum_{\ell=1}^{h}\frac{2}{c^2\ell^2}\right)\\ &\le \frac{1}{2}\exp\left(\frac{\pi^2}{3c^2}\right)\\ &< \frac{1+\delta}{2},\end{aligned}$$

where the first inequality holds because $N/2 - t \le 2k$, combined with the fact that $\prod_{\ell=1}^{h}(1+a_\ell) \le \exp(\sum_{\ell=1}^{h} a_\ell)$ for nonnegative a_ℓ.

For $r = N/2$, we get

$$\begin{aligned}|\hat{\eta}(r)| &= \frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{(N/2-t)(N/2-t+2k)\prod_{\ell=1}^{h}(2ck\ell+(N/2-t))(2ck\ell-(N/2-t))}\\ &= \frac{2k}{N/2-t+2k}\prod_{\ell=1}^{h}\left(\frac{(2ck\ell)^2}{(2ck\ell)^2-(N/2-t)^2}\right) \ge \frac{1}{2}.\end{aligned}$$

Now we analyze the remaining summands, and show that their total contribution is much smaller than 1. Recall that the choice $i = t$ minimizes $\pi_S(i)$, and that $\pi_S(t) = (2ck)^{2h}(h!)^2$. Therefore,

$$|\hat{\eta}(t+2ck\ell)| = \frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{\prod_{j\in T\setminus\{t+2ck\ell\}}|t+2ck\ell-j|} \le \frac{2k(N/2-t)}{|2ck\ell+2k||2ck\ell-(N/2-t)|} \le \frac{1}{c^2\ell^2-1},$$

where the final inequality exploits the fact that $N/2 - t < 2k$. Similarly,

$$|\hat{\eta}(t-2ck\ell)| = \frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{\prod_{j\in T\setminus\{t+2ck\ell\}}|t-2ck\ell-j|} \le \frac{2k(N/2-t)}{|2ck\ell-2k||2ck\ell+(N/2-t)|} \le \frac{1}{c^2\ell^2-1}.$$

We can use this quadratic decay to bound the total ℓ_1 mass of the points outside of $\{t-k, t, N/2\}$:

$$\sum_{j\in S\setminus\{t\}}|\hat{\eta}(j)| \le \sum_{\ell\neq 0}\frac{1}{(c^2\ell^2-1)} \le \frac{2}{c^2-1}\cdot\frac{\pi^2}{6} < \frac{\delta}{2}.$$

Now let $\eta_k(r) = (-1)^{r+h+N/2}\hat{\eta}(r)/\|\hat{\eta}\|_1$. Since $\hat{\eta}$ is supported on $T \subseteq \{2k, 4k, \ldots, 2\lfloor N/2k\rfloor k\} \cup \{N/2\}$, the function η_k is as well, giving the first claim. Moreover,

$$\eta_k(N/2) \ge \frac{1/2}{(1/2+\delta/2)+1/2+\delta/2} \ge \frac{1-\delta}{2}.$$

This yields the second claim about η_k. The third claim follows immediately from the definition. Finally, let p be a polynomial of degree strictly less than $|T|$ (where $|T| \ge \rho N/k$ for a constant ρ). Then

$$\sum_{r=0}^{N} p(r)\eta_k(r) = \sum_{r=0}^{N} p(r)\frac{(-1)^{r+h+N/2}}{\|\hat{\eta}\|_1}\binom{N}{r}\frac{(2ck)^{2h}(h!)^2(2k)(N/2-t)}{N!}\prod_{j\in[[N]]\setminus T}(r-j)$$
$$= \sum_{r=0}^{N}(-1)^r\binom{N}{r}q(r)$$

for a polynomial q of degree strictly less than N. This is equal to zero by Lemma 9, giving the final claim. □

We obtain our dual polynomial ψ for the $\mathrm{Col}_{N,R}$ as a linear combination of two simpler functions ψ_1 and ψ_2. These functions have the following properties.

Lemma 12. *Let $N > 0$ be an integer multiple of 4. For $R \geq N$, there exist explicit $\psi_1, \psi_2 \colon \{-1,1\}^n \to \mathbb{R}$ and $d = \Omega(\delta^{1/3}N^{1/3})$ such that*

1. $\sum_{x\in T_1}\psi_1(x) > \frac{1-\delta}{2}$.
2. $-\sum_{x\in T_2}\psi_2(x) > \frac{1-\delta}{2}$.
3. $\|\psi_1\|_1 = \|\psi_2\|_1 = 1$.
4. $\sum_{x\in T_2}|\psi_1(x)| = \sum_{x\in T_1}|\psi_2(x)| = 0$.
5. *ψ_1, ψ_2 have pure high degree at least d.*
6. $\sum_{x\in T_1}\psi_1(x) = \sum_{x\in R_{N/2,2,1}}\psi_2(x)$.
7. $\sum_{x\in R_{N/2,2,1}}\psi_1(x) = \sum_{x\in T_2}\psi_2(x)$.
8. *ψ_1 and ψ_2 are each constant on each set $R_{m,a,b}$ when (m,a,b) is valid.*

Together, they yield the desired dual polynomial for $\mathrm{Col}_{N,R}$.

Theorem 13. *Let $N > 0$ be an integer multiple of 4. For $R \geq N$, there exists an explicit $(1-6\delta, d)$-dual polynomial ψ for $\mathrm{Col}_{N,R}$ for $d = \Omega(\delta^{1/3}N^{1/3})$.*

Remark 14. *The dependence of the lower bound Theorem 13 on both parameters δ and N for $\frac{1}{N} \leq \delta \leq \frac{1}{10}$, is tight up to a logarithmic factor in the size of the range. We show this in Appendix A by constructing an explicit approximating polynomial for $\mathrm{Col}_{N,R}$ of the appropriate degree, by building on the ideas underlying the quantum query algorithm of Brassard et al. [15].*

Proof of Theorem 13, assuming Lemma 12. Let $a = \sum_{x\in T_1}\psi_1(x)$ and let $b = \sum_{x\in T_2}|\psi_2(x)| = -\sum_{x\in T_2}\psi_2(x)$, where ψ_1 and ψ_2 are as in Lemma 12. Let $\psi(x) = a\psi_1(x) + b\psi_2(x)$. By Property 5 of Lemma 12 and Lemma 15 below, ψ also has pure high degree at least d. So we need only show that ψ has correlation at least $1-6\delta$ with $\mathrm{Col}_{N,R}$. To this end, note that

1. $\sum_{x\in T_1}\psi(x) = a^2 > \frac{(1-\delta)^2}{4}$. This inequality uses Properties 1 and 3 of Lemma 12.
2. $-\sum_{x\in T_2}\psi(x) = b^2 > \frac{(1-\delta)^2}{4}$. This inequality uses Properties 2 and 3 of Lemma 12.

3. $\sum_{x\in B}|\psi(x)| \le a\sum_{x\in B\setminus R_{N/2,2,1}}|\psi_1(x)| + b\sum_{x\in B\setminus R_{N/2,2,1}}|\psi_2(x)| \le (a+b)\delta$. Here, the first inequality exploits the fact that

$$\sum_{x\in R_{N/2,2,1}}|\psi(x)| = \sum_{x\in R_{N/2,2,1}}|a\cdot\psi_1(x)+b\cdot\psi_2(x)| = 0. \tag{5}$$

The last equality in Eq. (5) holds because, for all $x \in R_{N/2,2,1}$,

$$\begin{aligned}
&a\cdot\psi_1(x)+b\cdot\psi_2(x)\\
&=\left(\sum_{x'\in T_1}\psi_1(x')\right)\cdot\psi_1(x)+\left(-\sum_{x'\in T_2}\psi_2(x')\right)\psi_2(x)\\
&=\left(\sum_{x'\in R_{N/2,2,1}}\psi_2(x')\right)\cdot\psi_1(x)+\left(-\sum_{x'\in R_{N/2,2,1}}\psi_1(x')\right)\psi_2(x)\\
&=\left(\sum_{x'\in R_{N/2,2,1}}\psi_2(x')\right)\left(\frac{1}{|R_{N/2,2,1}|}\sum_{x'\in R_{N/2,2,1}}\psi_1(x')\right)+\left(-\sum_{x'\in R_{N/2,2,1}}\psi_1(x')\right)\left(\frac{1}{|R_{N/2,2,1}|}\sum_{x'\in R_{N/2,2,1}}\psi_2(x')\right)\\
&=0,
\end{aligned}$$

where the second equality exploited Properties 6 and 7 of Lemma 12, and the last equality exploited Property 8.

Thus, the correlation of ψ with $\mathrm{Col}_{N,R}$ is

$$\begin{aligned}
\sum_{x\in T_1}\psi(x)-\sum_{x\in T_2}\psi(x)-\sum_{x\in B}|\psi(x)| &\ge a^2+b^2-(a+b)\delta\\
&\ge \frac{1}{2}-2\delta \ge (1-6\delta)\cdot\|\psi\|_1,
\end{aligned}$$

where the final inequality holds because $\|\psi\|_1 \le a^2+b^2+(a+b)\delta \le \frac{1}{2}+\delta$. □

Lemma 15. *Let $\psi_1,\psi_2 : \{-1,1\}^n \to \{-1,1\}$ each have pure high degree at least d. Then $\psi = \psi_1+\psi_2$ also has pure high degree at least d.*

Proof. Let $S \subseteq [n]$ with $|S| \le d$. Then

$$\sum_{x\in\{-1,1\}^n}\psi(x)\chi_S(x) = \sum_{x\in\{-1,1\}^n}\psi_1(x)\chi_S(x)+\sum_{x\in\{-1,1\}^n}\psi_2(x)\chi_S(x) = 0.$$

□

Proof of Lemma 12. Let ζ be the constant from Lemma 8, let ρ be the constant from Lemma 11, and let $\delta' = 1/2$. Set $K = 2(\rho N/\zeta)^{2/3}(\delta'/\delta)^{1/3}$. Let $d = \frac{1}{2}\rho^{1/3}\zeta^{2/3}(\delta')^{1/6}\delta^{1/3}N^{1/3} = \Omega(\delta^{1/3}N^{1/3})$, noting that $d \le \zeta(\delta/8)^{1/2}K^{1/2}$ and $d \le \rho(\delta')^{1/2}N/k$ for every $k \le K$. Let $\omega\colon \{1,\dots,K\} \to \mathbb{R}$, with correlation constant $\delta/8$, and $\eta_3,\dots,\eta_K\colon \{1,\dots,N\} \to \mathbb{R}$, with correlation constant δ', be as in the conclusions of those lemmas.

We start by defining a function $\Psi(m,k)$ as follows.

$$\Psi(m,k) = \omega(k)\cdot\mathbb{1}_{m=N/2} - \frac{\omega(k)}{\eta_k(N/2)}\mathbb{1}_{k\ge 3}\cdot\eta_k(m).$$

Here,

$$\mathbb{1}_{m=N/2} = \begin{cases} 1 & \text{if } m = N/2, \\ 0 & \text{otherwise,} \end{cases} \qquad \textit{and} \qquad \mathbb{1}_{k\geq 3} = \begin{cases} 1 & \text{if } k \geq 3, \\ 0 & \text{otherwise.} \end{cases}$$

We first show how to use Ψ to construct the polynomial ψ_1. Analogously to our construction of ϕ, we want ψ_1 to place a total weight of $\Psi(m,a)$ on each set $R_{m,a,1}$. Recall from our overview in Section 2.5 that we think of $\psi_1 = \psi' + \sum_{\text{invalid triples } (N/2,a,1)} \psi''_{N/2,a,1}$, where ψ' looks like the simpler "first stage" dual polynomial ϕ from our informal overview (which we constructed in Section 3) and each $\psi''_{N/2,a,1}$ cancels out the weight ϕ places on values of k the do not divide N. This structure underlies our construction of Ψ, where we add multiples of the polynomials $\eta_k(m)$ to cancel out the weight $\omega(k)$ places on invalid triples.

Now we construct and analyze the polynomial ψ_1. Define

$$\hat{\psi}_1(x) = \begin{cases} \Psi(m,k)/|R_{m,k,1}| & \text{if } x \in R_{m,k,1} \setminus T_1 \\ \Psi(N/2,1)/|T_1| & \text{if } x \in T_1 \\ 0 & \text{otherwise.} \end{cases}$$

Notice that $\hat{\psi}_1$ is well-defined, because any $x \notin T_1$ is in $R_{m,k,1}$ for at most one triple $(m,k,1)$. We collect several calculations with $\hat{\psi}_1$. First,

$$\sum_{x\in T_1} \hat{\psi}_1(x) = \Psi(N/2,1) = \omega(1) > \frac{1-\delta/8}{2},$$

$$-\sum_{x\in R_{N/2,2,1}} \hat{\psi}_1(x) = -\Psi(N/2,2) = -\omega(2) > \frac{1-\delta/8}{2},$$

and

$$\begin{aligned} \sum_{x\in B\setminus R_{N/2,2,1}} |\hat{\psi}_1(x)| &= \sum_{\{(m,k)\colon k|m\}\setminus\{(N/2,1),(N/2,2)\}} |\Psi(m,k)| \\ &= \sum_{k=3}^{K} \left|\frac{\omega(k)}{\eta_k(N/2)}\right| \sum_{i=1}^{\lfloor N/2k \rfloor} |\eta_k(2ki)| \\ &\leq 4\sum_{k=3}^{K} |\omega(k)| \\ &\leq \frac{\delta}{2}, \end{aligned}$$

where the penultimate inequality exploits Properties 2 and 3 of Lemma 11, and the final inequality exploits Properties 1-3 of Lemma 8.

Noting that $|\omega(1)| + |\omega(2)| \leq 1$, it follows that $\|\hat{\psi}_1\|_1 \leq 1 + \delta/2$. So setting $\psi_1 = \hat{\psi}_1/\|\hat{\psi}_1\|_1$, it is immediate that ψ_1 satisfies the first three properties in the statement of the lemma. ψ_1 also satisfies the fourth property, since for any $x \in T_2$, $\psi_1(x) = \Psi(N,2)/|T_2| = 0$.

Now we will show that $\hat{\psi}_1$, and hence ψ_1, has pure high degree at least d. We require two observations.

- Ψ is supported on (m,k) for which $k|m$. To see this, note first that for any $k \geq 3$, $\Psi(N/2,k) = \omega(k) - \frac{\omega(k)}{\eta_k(N/2)} \cdot \eta_k(N/2) = 0$. The claim now follows from Property 1 of Lemma 11, combined with the fact that $2|N$.
- $\Psi(m,1)$ is nonzero only for $m = N/2$, and hence $\sum_{x\in T_1} \hat{\psi}_1(x) = \Psi(N/2,1) = \sum_{m=1}^{N} \Psi(m,1)$.

Fix any $S \subseteq [n]$ with $|S| \leq d$. Let $Q(m,k)$ be a polynomial of degree at most $d-1$ in each variable such that, for all pairs (m,k) with $k|m$, $Q(m,k) = \mathbb{E}_{x \in R_{m,k,1}}[\chi_S(x)]$. The existence of such a bivariate polynomial Q is guaranteed by Lemma 5. Then the previous two observations together imply that:

$$\sum_{x \in \{-1,1\}^n} \hat{\psi}_1(x)\chi_S(x) = \sum_{m=1}^{N}\sum_{k=1}^{N} \Psi(m,k)Q(m,k). \tag{6}$$

We remark that a key point is the derivation of Eq. (6) is that we have no control over the evaluations $Q(m,k)$ when k does not divide m, yet this is rendered irrelevant because $\Psi(m,k) = 0$ for all such pairs.

The right hand side of Eq. (6) equals:

$$\sum_{k=1}^{K} \omega(k)Q(N/2,k) - \sum_{k=3}^{K} \frac{\omega(k)}{\eta_k(N/2)} \left(\eta_k(N/2)Q(N/2,k) + \sum_{i=1}^{\lfloor N/2k \rfloor} \eta_k(2ik)Q(2ik,k) \right). \tag{7}$$

The first sum in Eq. (7) is zero by Lemma 8 since $Q(N/2,k)$ is a polynomial of degree at most d in k. The second sum is also zero because for each fixed k, $Q(r,k)$ is a polynomial of degree at most d in the variable r, and hence the term in parentheses is zero by Lemma 11 (Parts 1 and 4). Thus $\hat{\psi}_1$ has pure high degree at least d.

The construction of ψ_2 is similar. This time, we let

$$\hat{\psi}_2(x) = \begin{cases} \Psi(m,k)/|R_{m,k,2}| & \text{if } x \in R_{m,k,2} \setminus T_2 \\ \Psi(N/2,2)/|T_2| & \text{if } x \in T_2 \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\hat{\psi}_2$ is well-defined, because any $x \notin T_2$ is in $R_{m,k,2}$ for at most one triple $(m,k,2)$. We define $\psi_2 = \hat{\psi}_2 / \|\hat{\psi}_2\|_1$. Showing that ψ_2 satisfies Properties 1-4 of the lemma follows from the same calculations we used for ψ_1.

To show that ψ_2 has pure high degree at least d, we require the following additional observations.

- Ψ is supported on pairs (m,k) for which $k|m$ and $2|(N-m)$. To see the latter property, note that if $\Psi(m,k) \neq 0$, then m is even (this holds because $N/2$ is even, which follows from our requirement that N is a multiple of 4), and hence $N-m$ is as well.
- $\Psi(m,2)$ is nonzero only for $m = N/2$. It follows that $\sum_{x \in T_2} \hat{\psi}_2(x) = \Psi(N/2,2) = \sum_{m=1}^{N} \Psi(m,2)$.

With these observations in hand, showing that ψ_2 has pure high degree d then follows from calculations analogous to the ones we used for ψ_1.

Finally, the fact that ψ_1 and ψ_2 satisfy Properties 6, 7, and 8 of the lemma follows from their definitions, combined with the fact that $R_{N/2,1,2} = R_{N/2,2,1}$. In fact, $\sum_{x \in T_1} \psi_1(x)$ equals $\Psi(N/2,1)$, while $\sum_{x \in R_{N/2,2,1}} \psi_2(x)$ also equals $\Psi(N/2,1)$, giving Property 6. Similarly, $\sum_{x \in T_2} \psi_2(x)$ equals $\Psi(N/2,2)$, while $\sum_{x \in R_{N/2,1,2}} \psi_1(x)$ also equals $\Psi(N/2,1)$. This completes the proof.

□

5 A Dual Polynomial for Element Distinctness

We first recall the reduction from Collision to Element Distinctness given in [4].[3] The reduction shows how to turn a polynomial p approximating $\mathrm{ED}_{M,R}$ into a polynomial q approximating $\mathrm{Col}_{N,R}$, with $N \approx M^2$ and $\deg q \leq \deg p$.

[3] While the reduction given in Aaronson and Shi's paper is stated in terms of quantum query algorithms, it is straightforward to rephrase the reduction in terms of approximating polynomials instead.

We illustrate the reduction for $N = M^2/12$. Let $p : \{-1,1\}^m \to \{-1,1\}$ be an $(1/6)$-approximation of $\mathrm{ED}_{M,R}$, with $m = M \log R$. Define a polynomial $q : \{-1,1\}^n \to \{-1,1\}$ for $n = N \log R$ by

$$q(y_1, \ldots, y_N) = \frac{1}{\binom{N}{M}} \sum_{1 \le i_1 < i_2 < \cdots < i_M \le N} p(y_{i_1}, y_{i_2}, \ldots, y_{i_M}).$$

That is, $q(y)$ is the expected value of $p(x)$ where x is the concatenation of a random subset of M of the blocks $y_1, \ldots, y_N$. To simplify notation, for a set $S = \{i_1, i_2, \ldots, i_M\}$, let $y|_S = (y_{i_1}, y_{i_2}, \ldots, y_{i_M})$. Note that $\deg q \le \deg p$. Moreover, since q is an average of values in $[-7/6, 7/6]$, it is always in $[-7/6, 7/6]$ itself. To finish arguing that q is a $(1/3)$-approximation to $\mathrm{Col}_{N,R}$, we consider two cases:

1. If $y \in T_1$, i.e., y is a 1-to-1 input, then $y|_S$ is always 1-to-1. Hence $p(y|_S) \in [5/6, 7/6]$ for every subset of indices, so $q(y) \in [2/3, 4/3]$.

2. If $y \in T_2$, i.e., y is a 2-to-1 input, then with high probability $y|_S$ is not 1-to-1. This follows from the "birthday bound":

$$\Pr_{|S|=M}[\mathrm{ED}(y|_S) = 1] \le \exp(-M^2/4N) \le \frac{1}{12}.$$

   Therefore, $q(y) \le (11/12)(-5/6) + (1/12)(7/6) \le -2/3$.

The construction we give in this section takes a dual view of the reduction above. Namely, we show how to transform a dual polynomial ψ for $\mathrm{Col}_{N,R}$ into a dual polynomial φ for $\mathrm{ED}_{M,R}$, with $M^2 \approx N$. In the primal reduction, we constructed $q(y)$ from $p(x)$ by averaging p over all subsets of size M. The right analogue in the dual reduction is to construct $\varphi(x)$ by averaging $\psi(y)$ over a carefully constructed set of *extensions* from x to a longer input y. In particular, $\varphi(x)$ averages $\psi(y)$ over all y for which x could have been produced by taking a subset of M blocks of y.

We give this reduction formally below.

Theorem 16. *Let $\psi : \{-1,1\}^n \to \{-1,1\}$ be a $(1-\delta, d)$-dual polynomial for $\mathrm{Col}_{N,R}$. Then ψ can be used to construct $\varphi : \{-1,1\}^m \to \{-1,1\}$ that is an $(1-2\delta, d)$-dual polynomial for $\mathrm{ED}_{M,R}$ when $M \ge 2\sqrt{N \log(2/\delta)}$.*

Corollary 17. *For any $\delta > 0$, there is an explicit $(1-\delta, d)$-dual polynomial for $\mathrm{ED}_{M,R}$ with $d = \Omega((\delta/\log(1/\delta))^{1/3} M^{2/3})$.*

Remark 18. *The dependence of Corollary 17 on δ is essentially tight for $\delta = O(M^{-2})$. See Appendix A for details.*

Proof of Theorem 16. Given a set $S = \{i_1, \ldots, i_M\} \subset [N]$ with $i_1 < i_2 < \cdots < i_M$ and a bit string $y = (y_1, \ldots, y_N) \in \{-1,1\}^n$, define the restriction of y to the set S, denoted by $y|_S \in \{-1,1\}^m$, to be the string of length $m = M \log R$ obtained by concatenating the blocks y_i for $i \in S$, i.e., $y|_S = (y_{i_1}, y_{i_2}, \ldots, y_{i_M})$. Given a bit string $x \in \{-1,1\}^m$, define the multiset of extensions of x, denoted by $\mathrm{ext}(x)$, to be the $\binom{N}{M} R^{N-M}$ strings $y \in \{-1,1\}^n$ where $y|_S = x$ for some $|S| = M$. Restrictions and extensions are related by the equivalence of the multisets:

$$\{(x,y) : x \in \{-1,1\}^m, y \in \mathrm{ext}(x)\} = \{(x,y) : y \in \{-1,1\}^n, x = y|_S \text{ for some } |S| = m\}.$$

For $x \in \{-1,1\}^m$, define the polynomial

$$\varphi(x) = \frac{1}{\binom{N}{M}} \sum_{y \in \mathrm{ext}(x)} \psi(y).$$

Let $\varphi(x) = 0$ for $x \notin \{-1,1\}^m$. We claim that φ is a good dual polynomial for the Element Distinctness function ED, which requires us to show

1. $\sum_{x\in\{-1,1\}^m} \varphi(x)\,\mathrm{ED}(x) > (1-2\delta)\cdot\sum_{x\in\{-1,1\}^m}|\varphi(x)|$
2. $\sum_{x\in\{-1,1\}^m} \varphi(x)\chi_S(x) = 0$ for all $|S|\le d$

To verify the first property, define

$$A(y) = \frac{1}{\binom{N}{M}}\sum_{|S|=M}\mathrm{ED}(y|_S).$$

We collect a few observations about A.

1. $|A(y)|\le 1$ for all y.
2. If $y\in T_1$, then $A(y)=1$.
3. If $y\in T_2$, then
$$\Pr_{|S|=M}[\mathrm{ED}(y|_S)=1]\le\exp(-M^2/4N).$$
Hence,
$$A(y)\le -1+2\exp(-M^2/4N)\le -1+\delta.$$

Therefore we get

$$\begin{aligned}
\sum_{x\in\{-1,1\}^m}\varphi(x)\,\mathrm{ED}(x) &= \frac{1}{\binom{N}{M}}\sum_{x\in\{-1,1\}^m}\sum_{y\in\mathrm{ext}(x)}\psi(y)\,\mathrm{ED}(x)\\
&= \frac{1}{\binom{N}{M}}\sum_{y\in\{-1,1\}^n}\sum_{|S|=M}\psi(y)\,\mathrm{ED}(y|_S)\\
&= \sum_{y\in\{-1,1\}^n}A(y)\psi(y)\\
&\ge \left(\sum_{y\in T_1}\psi(y)-\sum_{y\in T_2}\psi(y)-\sum_{y\in B}|\psi(y)|\right)-\delta\sum_{y\in T_2}|\psi(y)|\\
&\ge (1-2\delta)\sum_{y\in\{-1,1\}^n}|\psi(y)|\\
&= (1-2\delta)\sum_{y\in\{-1,1\}^n}\frac{1}{\binom{N}{M}}\sum_{|S|=M}|\psi(y)|\\
&\ge \frac{1-2\delta}{\binom{N}{M}}\sum_{x\in\{-1,1\}^m}\sum_{y\in\mathrm{ext}(x)}|\psi(y)|\\
&\ge (1-2\delta)\sum_{x\in\{-1,1\}^m}|\varphi(x)|.
\end{aligned}$$

For the second property, let T be a subset of $[N]$ with $|T| \leq d$. Then

$$\begin{aligned}\sum_{x\in\{-1,1\}^m} \varphi(x)\chi_T(x) &= \frac{1}{\binom{N}{M}} \sum_{x\in\{-1,1\}^m} \sum_{y\in \mathrm{ext}(x)} \psi(y)\chi_T(x) \\ &= \frac{1}{\binom{N}{M}} \sum_{|S|=M} \sum_{y\in\{-1,1\}^n} \psi(y)\chi_T(y|_S) \\ &= \frac{1}{\binom{N}{M}} \sum_{|S|=M} \sum_{y\in\{-1,1\}^n} \psi(y)\chi_{T|_S}(y) \\ &= 0,\end{aligned}$$

where $T|_S$ denotes the subset of T contained in the blocks specified by S.

□

6 On Complementary Slackness

Recalling that any bounded-error quantum query algorithm can be converted into an approximating polynomial [10], the collision-finding algorithm of Brassard, Høyer, and Tapp [15] yields an explicit, asymptotically optimal approximating polynomial for $\mathrm{Col}_{N,R}$. We describe this polynomial p below.

Recall that any approximating polynomial for $\mathrm{Col}_{N,R}$ represents a feasible solution to the primal linear program considered in Section 2.2. If the polynomial p were an *exactly* optimal ε-approximation for $\mathrm{Col}_{N,R}$, then complementary slackness would imply that the optimal dual polynomial ψ for $\mathrm{Col}_{N,R}$ is supported on the points corresponding to constraints made tight by p. That is, $\psi : \{-1,1\}^n \to \{-1,1\}$ is supported on $x \in \{-1,1\}^n$ for which $|p(x) - \mathrm{Col}(x)| = \varepsilon$. We refer to these as the *maximum-error points* of p.

While we do not know whether p is an exactly optimal approximating polynomial for $\mathrm{Col}_{N,R}$, we might still expect that an approximate version of complementary slackness might holds, in the sense that a "good" dual polynomial should place all or most of its weight on points that are "nearly" maximum-error points of p. Indeed, this intuition has proven accurate for all of the dual polynomials constructed in prior work, including for symmetric functions (see [17, Section 4.5]), block-composed functions (see [43, Section 1.2.4]), and the intersection of two majorities [38]. Below, we argue that k-to-1 inputs are nearly maximum-error points for p, which explains why our dual polynomials for collision are supported on inputs that are roughly k-to-1, in addition to why these inputs play a prominent role in the original symmetrization-based proof.

An asymptotically optimal approximation p for $\mathrm{Col}_{N,R}$. For a subset $S \subset [N]$, define $\mathsf{cross}_S : \{-1,1\}^n \to \mathbb{R}$ via:

$$\mathsf{cross}_S(x_1,\ldots,x_N) = |\{i \in S, j \notin S : x_i = x_j\}| = \sum_{i\in S, j\notin S} \mathrm{EQ}(x_i,x_j),$$

where EQ denotes the equality function. That is, $\mathsf{cross}_S(x)$ counts the number of cross-collisions between indices in S and indices outside of S. Notice that $\mathrm{EQ}(x_i,x_j)$ is a function of only $2 \cdot \log R$ variables, and hence $\mathsf{cross}_S(x_1,\ldots,x_N)$ is exactly computed by a polynomial of degree $2 \cdot \log R$.

In addition, for a subset $S \subset [N]$, define the function $\mathbb{I}_{\mathrm{ED},S}(x_1,\ldots,x_N)$ to be 1 if $x_i \neq x_j$ for all pairs $i,j \in S$ with $i \neq j$, and 0 otherwise. That is, $\mathbb{I}_{\mathrm{ED},S}$ indicates whether x is 1-to-1 on the indices in S. Notice that $\mathbb{I}_{\mathrm{ED},S}$ is a function of only $|S| \cdot \log R$ variables, and hence is exactly computed by a polynomial of degree $|S| \cdot \log R$.

For the remainder of the discussion, let $r = N^{1/3}$ – we focus on the quantity $\mathsf{cross}_S(x)$ when $|S| = r$. We will need the following simple observations.

1. If $x \in T_1$, i.e., x is a 1-to-1 input, then $\mathsf{cross}_S(x) = 0$ and $\mathbb{I}_{\mathrm{ED},S}(x) = 1$ for any S.
2. If $x \in T_2$, i.e., x is a 2-to-1 input, then $\mathbb{I}_{\mathrm{ED},S}(x) = 1 \implies \mathsf{cross}_S(x) = r$.
3. If $x \in T_2$, then, over the random choice of S, $\mathbb{I}_{\mathrm{ED},S}(x) = 0$ with probability at most $(N/2) \cdot (r/N)^2 \leq N^{-1/3}$.
4. For all $x \in \{-1, 1\}^n$, $\mathbb{I}_{\mathrm{ED},S}(x) = 1 \implies \mathsf{cross}_S(x) \leq N - r$.

Let $T_d : \mathbb{R} \to \mathbb{R}$ denote the degree-d Chebyshev polynomial of the first kind. This polynomial has the following properties:

- $T_d(x) \in [-1, 1]$ for $x \in [-1, 1]$.
- $T_d(1 + M/d^2) \geq 10$ for a constant M independent of d.
- The extreme points of T_d in $[-1, 1]$ are the degree-d *Chebyshev nodes*, which take the form $\cos(i\pi/d)$ for $0 \leq i \leq d$.

Truncating the Taylor expansion of $\cos(x) = 1 - x^2/2 + \ldots$ after the quadratic term, one sees that the Chebyshev nodes are well-approximated via the expression $\cos(i\pi/d) \approx 1 - (ci^2/d^2)$ for some constant c.

Applying an appropriate affine transformation to T_d, we obtain a polynomial A_d with the following properties:

- $A_d(0) = 1$.
- $A_d(i) \in [-1, -3/4]$ for all real numbers $i \in [1, d^2/M]$.
- $A_d(i) \in [-1, 1]$ for all real numbers $i \in [0, d^2/M]$.
- The extreme points of A_d are well approximated by the points $c \cdot i^2$ for $i \in \{0, 1, \ldots, \lfloor d \cdot M^{-1/2} \rfloor\}$.

Let $p_S(x) = \mathbb{I}_{\mathrm{ED},S}(x) \cdot A_d(\mathsf{cross}_S(x_1, \ldots, x_N)/r)$ for $d = 100 \cdot M \cdot N^{1/3}$, and let

$$p(x) = \mathbb{E}_{|S|=r}[p_S(x)] = \frac{1}{\binom{N}{r}} \sum_{|S|=r} p_S(x).$$

Then p is a polynomial of degree $|S| \log R + 2 \cdot d \cdot \log R = O(N^{1/3} \log R)$. We argue that p approximates $\mathrm{Col}_{N,R}$ to error ε for some $\varepsilon \leq 1/3$. The analysis falls into three cases.

Case 1: For $x \in T_1$, $p_S(x) = A_d(0) = -1$ for all S, where the first equality follows from Property 1 above. So $p(x) = \mathbb{E}_{|S|=r}[p_S(x)] = 1$.

Case 2: For $x \in T_2$, $\mathbb{I}_{\mathrm{ED},S}(x) = 1 \implies p_S(x) = A_d(1) \in [-1, -3/4]$, where the equality follows from Property 2 above. Meanwhile, $\mathbb{I}_{\mathrm{ED},S}(x) \neq 1 \implies p_S(x) = 0$. Combining these two facts with Property 3 above establishes that $p(x) = \mathbb{E}_{|S|=r}[p_S(x)] \in [-1, -2/3]$.

Case 3: For $x \in \{-1, 1\}^n$, $p_S(x) \in [-1, 1]$. This follows from Property 4 above.

Identifying maximum-error points of p. For any fixed S, the maximum error points of p_S are well-approximated by the $x \in \{1,1\}^n$ for which the following two equations hold:

$$\mathsf{cross}_S(x) = c \cdot i^2 \cdot r \text{ for some } i \in \{0, 1, \ldots, \lfloor d \cdot M^{-1/2} \rfloor\} \tag{8}$$

and

$$\mathbb{I}_{\mathrm{ED},S}(x) = 1. \tag{9}$$

(This follows from the fact that the extreme points of A_d are roughly of the form $c \cdot i^2$ for $0 \leq i \leq d \cdot M^{-1/2}$).

However, the maximum-error points for the averaged polynomial $p(x) = \mathbb{E}_{|S|=r}[p_S(x)]$ are the points x that satisfy Eq. (8) and Eq. (9) *with high probability* over the choice of S. Indeed, for these points x, the error of $p(x)$ is at least $\varepsilon \cdot (1 - o(1)) \approx \varepsilon$.

Consider any k of the form $k = c \cdot i^2 + 1$ for some $i \in \{0, 1, \ldots, \lfloor d \cdot M^{-1/2} \rfloor\}$, such that $k = o(N^{1/3})$. Consider any $x \in T_k$; we claim that x satisfies Eq. (8) and Eq. (9) with probability $1 - o(1)$ over choice of S. To see this, observe that the probability that $\mathbb{I}_{\mathrm{ED},S}(x_S) = 0$ is at most $(N/k) \cdot k^2 \cdot (r/N)^2 = \frac{k \cdot r^2}{N} = o(1)$. And if $\mathbb{I}_{\mathrm{ED},S}(x_S) \neq 0$, then the number of cross-collisions is exactly

$$\mathsf{cross}_S(x_1, \ldots, x_N) = r \cdot (k - 1).$$

When k takes the form $k = c \cdot i^2 + 1$, this means that x satisfies Eq. (8). Hence, x has nearly maximal error even for the averaged polynomial p.

Acknowledgements. We are grateful to Guy Kindler, Yaoyun Shi, and Mario Szegedy for several illuminating discussions during the early stages of this work. We also thank Scott Aaronson and Emanuele Viola for helpful comments on an earlier version of this manuscript.

References

[1] Scott Aaronson. Quantum lower bound for the collision problem. In John H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 635–642. ACM, 2002.

[2] Scott Aaronson. The polynomial method in quantum and classical computing. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, page 3. IEEE Computer Society, 2008.

[3] Scott Aaronson. The collision lower bound after 12 years. In *Qstart Conference*, 2013.

[4] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004.

[5] Andris Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. Syst. Sci.*, 64(4):750–767, 2002.

[6] Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1(1):37–46, 2005.

[7] Andris Ambainis. Polynomial degree vs. quantum query complexity. *Journal of Computer and System Sciences*, 72(2):220 – 238, 2006. {JCSS} {FOCS} 2003 Special Issue.

[8] Andris Ambainis. Quantum walk algorithm for element distinctness. *SIAM J. Comput.*, 37(1):210–239, 2007.

[9] Howard Barnum, Michael E. Saks, and Mario Szegedy. Quantum query complexity and semi-definite programming. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*, pages 179–193. IEEE Computer Society, 2003.

[10] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001.

[11] Richard Beigel. The polynomial method in circuit complexity. In *Proceedings of the Eigth Annual Structure in Complexity Theory Conference, San Diego, CA, USA, May 18-21, 1993*, pages 82–95. IEEE Computer Society, 1993.

[12] Richard Beigel. Perceptrons, pp, and the polynomial hierarchy. *Computational Complexity*, 4:339–349, 1994.

[13] A. Belovs and A. Rosmanis. Adversary Lower Bounds for the Collision and the Set Equality Problems. *ArXiv e-prints*, October 2013.

[14] Aleksandrs Belovs and Robert Spalek. Adversary lower bound for the k-sum problem. In Robert D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 323–328. ACM, 2013.

[15] G. Brassard, P. Høyer, and A. Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News (Cryptology Column)*, 28:14–19, 1997. quant-ph/9705002.

[16] Harry Buhrman, Nikolai K. Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *22nd Annual IEEE Conference on Computational Complexity (CCC 2007), 13-16 June 2007, San Diego, California, USA*, pages 24–32. IEEE Computer Society, 2007.

[17] Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and markov-bernstein inequalities. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *ICALP (1)*, volume 7965 of *Lecture Notes in Computer Science*, pages 303–314. Springer, 2013.

[18] Mark Bun and Justin Thaler. Hardness amplification and the approximate degree of constant-depth circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:151, 2013.

[19] Karthekeyan Chandrasekaran, Justin Thaler, Jonathan Ullman, and Andrew Wan. Faster private release of marginals on small databases. *CoRR*, abs/1304.3754, 2013.

[20] Peter Hoyer, Troy Lee, and Robert Spalek. Negative weights make adversaries stronger. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 526–535, New York, NY, USA, 2007. ACM.

[21] Peter Høyer, Michele Mosca, and Ronald de Wolf. Quantum search on bounded-error inputs. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, *Automata, Languages and Programming, 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30 - July 4, 2003. Proceedings*, volume 2719 of *Lecture Notes in Computer Science*, pages 291–299. Springer, 2003.

[22] Adam Tauman Kalai, Adam R. Klivans, Yishay Mansour, and Rocco A. Servedio. Agnostically learning halfspaces. *SIAM J. Comput.*, 37(6):1777–1805, 2008.

[23] Varun Kanade and Justin Thaler. Distribution-independent reliable learning. In Maria-Florina Balcan and Csaba Szepesvári, editors, *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, volume 35 of *JMLR Proceedings*, pages 3–24. JMLR.org, 2014.

[24] Adam R. Klivans and Rocco A. Servedio. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. *J. Comput. Syst. Sci.*, 68(2):303–318, 2004.

[25] Adam R. Klivans and Rocco A. Servedio. Toward attribute efficient learning of decision lists and parities. *Journal of Machine Learning Research*, 7:587–602, 2006.

[26] Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(1):29–36, 2005.

[27] Scott Aaronson lecturer. The collision problem: Notes for lecture 13 of mit course 6.845: Quantum complexity theory, 2010.

[28] Noam Nisan and Mario Szegedy. On the degree of boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994.

[29] Ryan O'Donnell and Rocco A. Servedio. New degree bounds for polynomial threshold functions. *Combinatorica*, 30(3):327–358, 2010.

[30] Ramamohan Paturi. On the degree of polynomials that approximate symmetric boolean functions (preliminary version). In S. Rao Kosaraju, Mike Fellows, Avi Wigderson, and John A. Ellis, editors, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada*, pages 468–474. ACM, 1992.

[31] Rocco A. Servedio, Li-Yang Tan, and Justin Thaler. Attribute-efficient learning and weight-degree tradeoffs for polynomial threshold functions. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT*, volume 23 of *JMLR Proceedings*, pages 14.1–14.19. JMLR.org, 2012.

[32] A. A. Sherstov. Breaking the Minsky-Papert barrier for constant-depth circuits. In *STOC*, 2014.

[33] Alexander A. Sherstov. Communication lower bounds using dual polynomials. *Bulletin of the EATCS*, 95:59–93, 2008.

[34] Alexander A. Sherstov. Separating ac^0 from depth-2 majority circuits. *SIAM J. Comput.*, 38(6):2113–2129, 2009.

[35] Alexander A. Sherstov. The pattern matrix method. *SIAM J. Comput.*, 40(6):1969–2000, 2011.

[36] Alexander A. Sherstov. Strong direct product theorems for quantum communication and query complexity. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 41–50. ACM, 2011.

[37] Alexander A. Sherstov. Approximating the and-or tree. *Theory of Computing*, 9(20):653–663, 2013.

[38] Alexander A. Sherstov. The intersection of two halfspaces has high threshold degree. *SIAM J. Comput.*, 42(6):2329–2374, 2013.

[39] Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 513–519. IEEE Computer Society, 2002.

[40] Yaoyun Shi and Yufan Zhu. Quantum communication complexity of block-composed functions. *Quantum Information & Computation*, 9(5):444–460, 2009.

[41] Robert Spalek. A dual polynomial for OR. *CoRR*, abs/0803.4516, 2008.

[42] Robert Spalek and Mario Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2(1):1–18, 2006.

[43] Justin Thaler. Lower bounds for the approximate degree of block-composed functions. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:150, 2014.

[44] Justin Thaler, Jonathan Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, volume 7391 of *Lecture Notes in Computer Science*, pages 810–821. Springer, 2012.

[45] Henry Yuen. A quantum lower bound for distinguishing random functions from random permutations. *Quantum Information & Computation*, 14(13-14):1089–1097, 2014.

[46] Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Information & Computation*, 15(7&8):557–567, 2015.

[47] Shengyu Zhang. On the power of ambainis lower bounds. *Theoretical Computer Science*, 339(23):241 – 256, 2005.

A On the Tightness of Theorem 13 and Corollary 17

To complement Theorem 13, we construct an approximating polynomial that gives a nearly matching upper bound on the approximate degree of $\mathrm{Col}_{N,R}$. The construction is a refinement of the approximating polynomial given in Section 6.

Proposition 19. *For $0 \leq \delta \leq 1/N$, there exists a polynomial p of degree $O(\delta^{1/3}N^{1/3}\log R)$ that $(1-\delta)$-approximates $\mathrm{Col}_{N,R}$.*

Proof sketch. See Section 6 for the construction of an approximating polynomial of degree $O(N^{1/3}\log R)$ in the case where δ is constant. In order to obtain an improved upper bound for vanishing δ, we make the following changes to that construction:

1. We instead choose $r = \delta^{1/3}N^{1/3}$. Now if x is a 2-1 input, the probability over the random choice of the set S of obtaining a collision inside S, i.e. the probability that $\mathbb{I}_{\mathrm{ED},S} = 0$, is at most $(N/2)\cdot(r/N)^2 \leq \delta/2$.

2. We instead let A_d be an affine transformation of a Chebyshev polynomial with the following properties for some constant M:

   - $A_d(0) \geq \frac{\delta}{2}$
   - $A_d(i) \in [-1, -\frac{\delta}{2}]$ for $i \in [1, d^2/M\delta]$
   - $A_d(i) \in [-1, 1]$ for $x \in [0, d^2/M\delta]$.

3. Setting $d = 100 \cdot M \cdot r$ ensures that the polynomial p has degree $O(\delta^{1/3}N^{1/3}\log R)$ and is a $(1-\delta)$-approximation of $\mathrm{Col}_{N,R}$.

□

We now show that Corollary 17 is tight up to a factor of $\log R$, when $\delta \leq 1/M^2$. This gives mild evidence that the lower bound has the right dependence on both parameters M, δ for vanishing δ.

Proposition 20. *Let $\delta \leq 1/M^2$. Then there exists a $(1-\delta)$-approximating polynomial for $\mathrm{ED}_{M,R}$ with degree $O(\log R)$.*

Proof. We write

$$\mathrm{ED}_{M,R}(x_1, \ldots, x_M) = \bigwedge_{i \neq j} \mathrm{NEQ}(x_i, x_j),$$

where $\mathrm{NEQ}(x_i, x_j) = 1$ if i and inputs j are distinct, and is zero otherwise. The function NEQ can be computed exactly by a polynomial of degree $O(\log R)$. Therefore, the polynomial

$$\frac{1}{\binom{M}{2}} \left(\frac{1}{2} - \sum_{i \neq j} \mathrm{NEQ}(x_i, x_j) \right)$$

has degree $O(\log R)$ and approximates $\mathrm{ED}_{M,R}$ to within error $1 - 1/M^2$. □

ECCC ISSN 1433-8092

http://eccc.hpi-web.de