Quantum Lower Bound for Recursive Fourier Sampling

Scott Aaronson*
Computer Science Department
University of California, Berkeley
aaronson@cs.berkeley.edu

Abstract

We revisit the oft-neglected 'recursive Fourier sampling' (RFS) problem, introduced by Bernstein and Vazirani to prove an oracle separation between BPP and BQP. We show that the known quantum algorithm for RFS is essentially optimal, despite its seemingly wasteful need to uncompute information. This implies that, to place BQP outside of PH [log] relative to an oracle, one would need to go outside the RFS framework. Our proof argues that, given any variant of RFS, either the adversary method of Ambainis yields a good quantum lower bound, or else there is an efficient classical algorithm. This technique may be of independent interest.

Subject classifications: quantum computing, lower bounds, query complexity.

1 Introduction

Quantum computing first gained notoriety with Shor's factoring algorithm [15], which built on earlier work of Simon [16]. It is sometimes claimed that, before Simon and Shor's breakthroughs, there was no credible evidence that quantum computers can yield a superpolynomial speedup over classical ones. For, although the Deutsch-Jozsa algorithm [10] was discovered earlier than Simon's algorithm, the former gives a speedup only in the exact case, not the bounded-error case.

However, there is an oft-neglected algorithm that appeared after Deutsch and Jozsa's and before Simon's. This is the recursive Fourier sampling (henceforth RFS) algorithm, which was used by Bernstein and Vazirani [6] to obtain the first oracle separation between BPP and BQP.[1] There are two likely reasons

*Supported by an NSF Graduate Fellowship and by the Defense Advanced Research Projects Agency (DARPA) and Air Force Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-01-2-0524.

[1]For the definitions of complexity classes used in this paper, see [2]. For background on quantum computing and on the quantum oracle model, see [5, 6].

for this neglect. First, the RFS problem seems artificial. It was introduced for the sole purpose of proving an oracle result, and is unlike all other problems for which a quantum speedup is known. (We will define RFS in Section 2; but for now, it involves a tree of depth $\log n$, where each vertex is labeled with a function to be evaluated via a Fourier transform.) Second, the speedup for RFS is only quasipolynomial (i.e. n versus $n^{\log n}$), rather than exponential as for the period-finding and hidden subgroup problems.

Nevertheless, we believe RFS is due for a comeback. Despite (or perhaps because of) its artificiality, the problem serves as an important link between quantum computing and the techniques of classical complexity theory. One reason is that, although other problems in BQP—such as the factoring, discrete logarithm, and 'shifted Legendre symbol' problems [9]—are thought to be classically intractable, these problems are quite low-level by complexity-theoretic standards. They, or their associated decision problems, are in $\mathsf{NP} \cap \mathsf{coNP}$.[2] By contrast, Bernstein and Vazirani [6] showed that, as an oracle problem, RFS lies outside NP and even MA (the latter result is unpublished, though not difficult). Subsequently Watrous [19] gave an oracle A, based on an unrelated problem, for which $\mathsf{BQP}^A \not\subseteq \mathsf{MA}^A$.[3] Also, Green and Pruim [12] gave an oracle B for which $\mathsf{BQP}^B \not\subseteq \mathsf{P}^{\mathsf{NP}^B}$. However, Watrous' problem was shown by Babai [4] to be in AM, while Green and Pruim's problem is in BPP. Thus, neither problem can be used to place BQP outside higher levels of the polynomial hierarchy.

On the other hand, Vazirani [18] and others have conjectured that RFS is not in PH, from which it would follow that there exists an oracle A relative to which $\mathsf{BQP}^A \not\subseteq \mathsf{PH}^A$. Proving this is, in our view, one of the central open problems of quantum complexity theory. Its solution is likely to require radically new techniques for circuit lower bounds.[4]

Here we address a different question. If BQP is indeed outside of PH relative to an oracle, how many alternations *are* needed to simulate it classically? The RFS problem is trivially in $\mathsf{PH}[\log]$ (that is, PH with logarithmic alternations). Yet several researchers independently expressed to us the hope that, if RFS were modified a bit, a quantum algorithm might be able to handle more than $\log n$ levels of recursion. One might thereby obtain an oracle problem that is in BQP, yet outside $\mathsf{PH}[\log]$ or even $\mathsf{PH}[poly \log]$. We were thus led to investigate whether this hope could be realized.

Our conclusion is negative. In this paper we examine a broad class of variations on RFS, and show that each of them has either (1) a good quantum

[2] For the shifted Legendre symbol problem, this is true assuming a number-theoretic conjecture of Boneh and Lipton [7].

[3] Actually, to place BQP outside MA relative to an oracle, it suffices to consider the complement of Simon's problem ("Does $f(x) = f(x \oplus s)$ only when $s = 0$?").

[4] For the RFS function can be represented by a low-degree real polynomial—this follows from the existence of a polynomial-time quantum algorithm for RFS, together with the result of Beals et al. [5] relating quantum algorithms to low-degree polynomials. As a result, the circuit lower bound technique of Razborov [14] and Smolensky [17], which is based on the nonexistence of low-degree polynomials, seems unlikely to work. Even the random restriction method of Furst et al. [11] can be related to low-degree polynomials, as shown by Linial et al. [13].

lower bound or (2) an efficient classical algorithm—there are no 'in-between' cases. It follows that, to place BQP outside of PH [log] relative to an oracle, one would need to go outside the RFS framework as formulated in this paper. That rather arcane-sounding assertion has a broader context. Like a classical algorithm, a quantum algorithm can solve problems recursively by calling itself as a subroutine. When this is done, though, the algorithm typically needs to call itself *twice* for each subproblem to be solved. The second call's purpose is to *uncompute* garbage left over by the first call, and thereby enable interference between different branches of the computation. The need to uncompute is what causes the standard RFS algorithm to have query complexity 2^h, where h is the height of a tree to be evaluated. One might wonder, though, whether the uncomputing step is really necessary, or whether a cleverly designed algorithm might avoid it. Our result gives, to our knowledge, the first nontrivial example for which recursive uncomputation *is* provably necessary. We conjecture that uncomputation is needed as well for other recursive problems, such as game-tree evaluation.[5]

The plan is as follows. In Section 2 we define the RFS problem and show that it lies in BQP. In Section 3, we use the adversary method of Ambainis [3] to prove a lower bound on the quantum query complexity of any RFS variant. This bound, however, requires a parameter called the *nonparity coefficient* to be large. The crux of our argument is that the nonparity coefficient is always above a certain threshold, unless the RFS variant is trivial (i.e. admits an efficient classical algorithm). We prove the general case of this result in Section 3; then, in Section 4, we prove a much stronger version for the important special case of 'total' variants. We conclude in Section 5.

2 Preliminaries

In ordinary Fourier sampling, we are given oracle access to a Boolean function $A : \{0,1\}^n \rightarrow \{0,1\}$, and are promised that there exists a secret string $s \in \{0,1\}^n$ such that $A(x) = s \cdot x \pmod 2$ for all x. We are asked to return $g(s)$, where $g : \{0,1\}^n \rightarrow \{0,1\}$ is a known Boolean function. We may assume that either g is efficiently computable, or else we are given access to an oracle for g.

Then, to obtain a height-2 recursive Fourier sampling tree, we simply compose this problem. That is, we are no longer given direct access to $A(x)$, but instead are promised that $A(x) = g(s_x)$, where $s_x \in \{0,1\}^n$ is the secret string for another Fourier sampling problem. A query then takes the form (x, y), and produces as output $A_x(y) = s_x \cdot y \pmod 2$. As before, we are promised that there exists an s such that $A(x) = s \cdot x \pmod 2$ for all x, meaning that the s_x strings must be chosen consistent with this promise. Again we must return $g(s)$.

Continuing, we can define height-h recursive Fourier sampling, or RFS_h,

[5] Formally, we conjecture that the quantum query complexity of evaluating a game tree increases with depth as the number of leaves is held constant, even if there is at most one winning move per vertex (and hence no need to bound error probability).

inductively as follows. We are given oracle access to a function $A(x_1, \ldots, x_h)$ for all $x_1, \ldots, x_h \in \{0,1\}^n$, and are promised that

(1) for each fixed x_1^*, $A(x_1^*, x_2, \ldots, x_h)$ is an instance of RFS_{h-1} on $x_2, \ldots, x_h$, having answer $b(x_1^*) \in \{0,1\}$; and

(2) there exists a secret string $s \in \{0,1\}^n$ such that $A(x_1, \ldots, x_h) = s \cdot b(x_1) \pmod 2$ for all x_1.

Again we must return $g(s)$. Note that we take g to be the same everywhere in the tree. Allowing different g's at different vertices would, we believe, complicate the results without adding anything conceptually new. As an example that will be used later, we could take $g(s) = g_{\text{mod}\,3}(s)$, where $g_{\text{mod}\,3}(s) = 0$ if $|s| \equiv 0 \pmod 3$ and $g_{\text{mod}\,3}(s) = 1$ otherwise, and $|s|$ denotes the Hamming weight of s. We do *not* want to take g to be the parity of s, for if we did then $g(s)$ could be evaluated using a single query.

Bernstein and Vazirani [6] showed that $RFS_{\log n}$, or RFS with height $\log n$ (all logarithms are base 2), is solvable on a quantum computer in time polynomial in n. We include a proof for completeness. Let $A = (A_n)_{n \geq 0}$ be an oracle that, for each n, encodes an instance of $RFS_{\log n}$ whose answer is Ψ_n. Then let L_A be the unary language $\{0^n : \Psi_n = 1\}$.

Lemma 1 *For any choice of A, $L_A \in \mathsf{EQP}^A \subseteq \mathsf{BQP}^A$.*

Proof. $RFS_1(n)$ can be solved exactly in four queries, with no garbage bits left over. The algorithm is as follows: first prepare the state

$$2^{-n/2} \sum_{x \in \{0,1\}^n} |x\rangle |A(x)\rangle,$$

using one query to A. Then apply a phase flip conditioned on $A(x) = 1$, and uncompute $A(x)$ using a second query, obtaining

$$2^{-n/2} \sum_{x \in \{0,1\}^n} (-1)^{A(x)} |x\rangle.$$

Then apply a Hadamard gate to each bit of the $|x\rangle$ register. It can be checked that the resulting state is simply $|s\rangle$. One can then compute $|s\rangle |g(s)\rangle$ and uncompute $|s\rangle$ using two more queries to A, to obtain $|g(s)\rangle$. To solve $RFS_{\log n}(n)$, we simply apply the above algorithm recursively at each level of the tree. The total number of queries used is $4^{\log n} = n^2$.

One can further reduce the number of queries to $2^{\log n} = n$ by using the "one-call kickback trick," described by Cleve et al. [8]. Here one prepares the state

$$2^{-n/2} \sum_{x \in \{0,1\}^n} |x\rangle \otimes \frac{|1\rangle - |0\rangle}{\sqrt{2}}$$

and then exclusive-OR's $A(x)$ into the second register. This induces the desired phase $(-1)^{A(x)}$ without the need to uncompute $A(x)$. (However, one still needs to uncompute $|s\rangle$ after computing $|g(s)\rangle$.) ■

A remark on notation: to avoid confusion with subscripts, we denote the i^{th} bit of string x by $x[i]$.

3 Quantum Lower Bound

In this section we prove a lower bound on the quantum query complexity of RFS. Crucially, the bound should hold for *any* nontrivial one-bit function of the secret strings, not just (say) the function $g_{\mathrm{mod}\,3}(s)$ defined in Section 2. We should even allow the function to be partial, meaning that every secret string must satisfy a promise, separate from the RFS promise of the previous section. Thus, let $g : \Omega_g \rightarrow \{0,1\}$ be a function of the secret strings, where $\Omega_g \subseteq \{0,1\}^n$ is the range of g. Then let RFS_h^g be height-h RFS in which the problem at each vertex is to return $g(s)$, under the promise that $s \in \Omega_g$. The following notion turns out to be crucial.

Definition 2 *The nonparity coefficient $\mu(g)$ of g is the maximum μ^* for which the following holds. There exist distributions D_0, D_1 over $g^{-1}(0), g^{-1}(1)$ respectively such that for all $z \in \{0,1\}^n \setminus \{0^n\}$, $\widehat{s}_0 \in g^{-1}(0)$ and $\widehat{s}_1 \in g^{-1}(1)$,*

$$\Pr_{s_0 \in D_0, s_1 \in D_1}\left[s_0 \cdot z \equiv \widehat{s}_1 \cdot z \pmod 2 \quad \vee \quad s_1 \cdot z \equiv \widehat{s}_0 \cdot z \pmod 2\right] \geq \mu^*.$$

One can verify that $\mu(g) \in [0, 3/4]$ for all g. Intuitively, the nonparity coefficient is high if the parity of some subset of bits of s is never strongly correlated with $g(s)$. For example, we show in Proposition 5 that $\mu(g_{\mathrm{mod}\,3}(s)) = 3/4 - O(1/n)$. At the other extreme, $\mu(g) = 0$ if and only if g is a parity function or a restriction thereof; note that RFS_h^g is solvable in a single classical query if $\mu(g) = 0$. In Theorem 4 we show that for all g,

$$Q_2\left(RFS_h^g\right) = \Omega\left(\left(\frac{1}{1-\mu(g)}\right)^{h/2}\right),$$

where Q_2 is bounded-error quantum query complexity as defined by Beals et al. [5]. In other words, any RFS problem with μ bounded away from 0 requires a number of queries exponential in the tree height h. To show this we use the adversary method of Ambainis [3].

However, there is an essential further part of the argument, which restricts the values of $\mu(g)$ itself. Suppose there existed a family $\{g_n\}$ of 'pseudoparity' functions: that is, $\mu(g_n) > 0$ for all n, yet $\mu(g_n) = O(1/\log n)$. Then the best lower bound we could obtain from Theorem 4 would be $\Omega\left((1 + 1/\log n)^{h/2}\right)$, suggesting that $RFS_{\log^2 n}^g$ might still be solvable in quantum polynomial time. On the other hand, it would be unclear a priori how to solve $RFS_{\log^2 n}^g$ classically with a logarithmic number of alternations. In Theorem 6 we rule out this scenario by showing that pseudoparity functions do not exist: if $\mu(g) < 0.146$ then g is a parity function, and hence $\mu(g) = 0$.

The theorem of Ambainis [3] that we need is his 'most general' lower bound, which he uses to show that the quantum query complexity of inverting a permutation is $\Omega(\sqrt{n})$.

Theorem 3 (Ambainis) *Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a Boolean function and X, Y two sets of inputs such that $f(x) \neq f(y)$ if $x \in X$ and $y \in Y$. Let $R \subset X \times Y$ be such that*

1. *For every $x \in X$, there are at least m different $y \in Y$ such that $(x, y) \in R$.*
2. *For every $y \in Y$, there are at least m' different $x \in X$ such that $(x, y) \in R$.*

For all $x \in X$, let $l_{x,i}$ be the number of $y \in Y$ such that $(x, y) \in R$ and $x[i] \neq y[i]$; also, for all $y \in Y$, let $l_{y,i}$ be the number of $x \in X$ such that $(x, y) \in R$ and $x[i] \neq y[i]$. Let $l_{\max}$ be the maximum of $l_{x,i} l_{y,i}$ over all $(x, y) \in R$ and $i \in \{1, \ldots, n\}$ such that $x[i] \neq y[i]$. Then $Q_2(f) = \Omega\left(\sqrt{mm'/l_{\max}}\right)$.

We can trivially extend Theorem 3 to allow weighted inputs by, for example, letting X and Y contain multiple copies of each input. We can then obtain the lower bound for RFS.

Theorem 4 *For all g, $Q_2(RFS_h^g) = \Omega\left((1 - \mu(g))^{-h/2}\right)$.*

Proof. Let X be the set of all inputs to RFS_h with output 0, and Y the set of all inputs with output 1. Say that $x \in X$ and $y \in Y$ *differ minimally* if, at every vertex v of the RFS tree, if the answer bit $g(s)$ at v is the same for x and y, then the subtrees rooted at v are identical in x and y. We will take $(x, y) \in R$ if and only if x and y differ minimally.

We weight the inputs according to the distributions D_0, D_1 from the definition of the nonparity coefficient. In particular, input x is weighted by the product, over all vertices v in the tree, of the probability of the secret string s at v, if s is drawn from $D_{g(s)}$ (where we condition on v's output bit, $g(s)$). Fix $i \in \{1, \ldots, n\}$ and a pair $(\widehat{x}, \widehat{y}) \in R$. Then

$$\Pr_{x \in D_0 : (x, \widehat{y}) \in R}\left[x[i] \neq \widehat{y}[i]\right] \Pr_{y \in D_1 : (\widehat{x}, y) \in R}\left[\widehat{x}[i] \neq y[i]\right] \leq (1 - \mu)^h.$$

This says the following. Suppose we choose $x \in X$ at random, conditioned on it differing minimally from $\widehat{y}$, and $y \in Y$ at random, conditioned on it differing minimally from $\widehat{x}$. Then the probability that, at any particular bit i, x differs from $\widehat{y}$ and $\widehat{x}$ from y, decreases exponentially in h. To see this, observe that $x[i] = \widehat{y}[i]$ if $s_x \cdot z \equiv s_{\widehat{y}} \cdot z \pmod 2$ at *any* vertex v along the path from the root to i, where s_x and $s_{\widehat{y}}$ are the secret strings of x and $\widehat{y}$ respectively at v. Similarly, $\widehat{x}[i] = y[i]$ if $s_{\widehat{x}} \cdot z \equiv s_y \cdot z \pmod 2$ at any v from the root to i. The events

$$s_x \cdot z \equiv s_{\widehat{y}} \cdot z \pmod 2, \quad s_{\widehat{x}} \cdot z \equiv s_y \cdot z \pmod 2$$

occur independently, both with probability at least μ.

So applying Theorem 3, with inputs weighted so that $m = m' = 1$ and $l_{\max} \leq (1-\mu)^h$, we conclude that

$$Q_2\left(RFS_h^g\right) = \Omega\left(\sqrt{mm'/l_{\max}}\right) = \Omega\left((1-\mu)^{-h/2}\right).$$

■

We now show that there is a natural choice of g—the function $g_{\operatorname{mod} 3}(s)$ defined in Section 2—for which the nonparity coefficient is almost $3/4$. Thus, for $g = g_{\operatorname{mod} 3}$, the algorithm of Lemma 1 is essentially optimal.

Proposition 5 $\mu\left(g_{\operatorname{mod} 3}\right) = 3/4 - O(1/n)$.

Proof. Let $n \geq 6$. Let D_0 be the uniform distribution over all s with $|s| = 3\lfloor n/6 \rfloor$ (so $g_{\operatorname{mod} 3}(s) = 0$); likewise let D_1 be the uniform distribution over s with $|s| = 3\lfloor n/6 \rfloor + 2$ ($g_{\operatorname{mod} 3}(s) = 1$). We consider only the case of s drawn from D_0; the D_1 case is analogous. We will show that for any z,

$$\left|\Pr_{s \in D_0}[s \cdot z \equiv 0] - 1/2\right| = O(1/n)$$

(all congruences are mod 2). The theorem then follows, since by the definition of the nonparity coefficient, given any z the choices of $s_0 \in D_0$ and $s_1 \in D_1$ are independent.

Assume without loss of generality that $1 \leq |z| \leq n/2$ (if $|z| > n/2$, then replace z by its complement). We apply induction on $|z|$. If $|z| = 1$, then clearly

$$\Pr[s \cdot z \equiv 0] = 3\lfloor n/6 \rfloor / n = 1/2 \pm O(1/n).$$

For $|z| \geq 2$, let $z = z_1 \oplus z_2$, where z_2 contains only the rightmost 1 of z and z_1 contains all the other 1's. Suppose the proposition holds for $|z| - 1$. Then

$$\begin{aligned}\Pr[s \cdot z \equiv 0] = {} & \Pr[s \cdot z_1 \equiv 0] \Pr[s \cdot z_2 \equiv 0 | s \cdot z_1 \equiv 0] + \\ & \Pr[s \cdot z_1 \equiv 1] \Pr[s \cdot z_2 \equiv 1 | s \cdot z_1 \equiv 1],\end{aligned}$$

where

$$\Pr[s \cdot z_1 \equiv 0] = 1/2 + \alpha, \quad \Pr[s \cdot z_1 \equiv 1] = 1/2 - \alpha$$

for some $|\alpha| = O(1/n)$. Furthermore, even conditioned on $s \cdot z_1$, the expected number of 1's in s outside of z_1 is $(n - |z_1|)/2 \pm O(1)$ and they are uniformly distributed. Therefore

$$\Pr[s \cdot z_2 \equiv b | s \cdot z_1 \equiv b] = 1/2 + \beta_b$$

for some $|\beta_0|, |\beta_1| = O(1/n)$. So

$$\begin{aligned}\Pr[s \cdot z \equiv 0] &= 1/2 + \beta_0/2 + \alpha\beta_0 - \beta_1/2 - \alpha\beta_1 \\ &= 1/2 \pm O(1/n).\end{aligned}$$

■

Finally we need to show that pseudoparity functions do not exist. That is, if g is too close to a parity function for the bound of Theorem 4 to apply, then g actually *is* a parity function, from which it follows that RFS_h^g admits an efficient classical algorithm.

Theorem 6 *Suppose $\mu(g) < 0.146$. Then g is a parity function (equivalently, $\mu(g) = 0$).*

Proof. By linear programming duality, there exists a distribution D over $z \in \{0,1\}^n \setminus \{0^n\}$, $\widehat{s}_0 \in g^{-1}(0)$ and $\widehat{s}_1 \in g^{-1}(1)$ such that for all $s_0 \in g^{-1}(0)$ and $s_1 \in g^{-1}(1)$,

$$\Pr_D\left[s_0 \cdot z \equiv \widehat{s}_1 \cdot z \pmod 2 \quad \vee \quad s_1 \cdot z \equiv \widehat{s}_0 \cdot z \pmod 2\right] < \mu.$$

Furthermore $\widehat{s}_0 \cdot z \not\equiv \widehat{s}_1 \cdot z \pmod 2$, since otherwise we could violate the hypothesis by taking $s_0 = \widehat{s}_0$ or $s_1 = \widehat{s}_1$. It follows that there exists a joint distribution $\widehat{D}$ over $z \in \{0,1\}^n$ and $b \in \{0,1\}$ such that for all (s_0, s_1),

$$\begin{aligned} \Pr_{\widehat{D}}\left[s_0 \cdot z \equiv b \pmod 2\right] &> 1-\mu, \\ \Pr_{\widehat{D}}\left[s_1 \cdot z \not\equiv b \pmod 2\right] &> 1-\mu. \end{aligned}$$

We have established that g is a bounded-error threshold function of parity functions. More precisely, there exist probabilities p_z, summing to 1, as well as $b_z \in \{0,1\}$ such that for all s,

$$\Psi(s) = \sum_{z \in \{0,1\}^n \setminus \{0^n\}} p_z \left(s \cdot z \oplus b_z\right) \text{ is } \begin{cases} > 1-\mu & \text{if } g(s) = 1 \\ < \mu & \text{if } g(s) = 0. \end{cases}.$$

We will draw s uniformly at random from $\{0,1\}^n$, and consider $var(\Psi)$, the variance of $\Psi(s)$ under this distribution. First, if $p_z \geq 1/2$ for any z, then $g(s) = s \cdot z \oplus b_z$ and $\mu(g) = 0$. Second, suppose $p_z < 1/2$ for all z. Then since s is uniform, for each $z_1 \neq z_2$ we know that $s \cdot z_1 \oplus b_{z_1}$ and $s \cdot z_2 \oplus b_{z_2}$ are pairwise independent 0-1 random variables, both with expectation $1/2$. So

$$var(\Psi) = (1/4) \sum_z p_z^2 < 1/8.$$

On the other hand, since $\Psi(s)$ is always less than μ or greater than $1 - \mu$,

$$var(\Psi) > (1/2 - \mu)^2.$$

Combining,

$$\mu > \left(2 - \sqrt{2}\right)/4 > 0.146.$$

■

4 Parity Coefficients for Total Functions

To make our results as general as possible, in the previous section we allowed g to be partial; that is, an arbitrary promise could be placed on g's inputs. However, if we require g to be total (defined on every $s \in \{0,1\}^n$), then we can obtain a sharper and more surprising variant of Theorem 6. The variant allows the RFS lower bound to be shown using only Ambainis' standard lower bound result, rather than his 'most general' result. Apart from that, the variant may be of independent combinatorial interest.

For any g (partial or total), we can modify our notion of a 'nonparity coefficient' by allowing each z to test the distributions D_0 and D_1 *separately*:

Definition 7 *The one-sided nonparity coefficient $\mu_1(g)$ of g is the maximum μ_1^* for which the following holds. There exist distributions D_0, D_1 over $g^{-1}(0)$ and $g^{-1}(1)$ respectively such that for all $z \in \{0,1\}^n \setminus \{0^n\}$, $\widehat{s}_0 \in g^{-1}(0)$ and $\widehat{s}_1 \in g^{-1}(1)$,*

$$\begin{aligned} \Pr_{s_0 \in D_0}[s_0 \cdot z \equiv \widehat{s}_1 \cdot z \pmod 2] &\geq \mu_1^* \text{ and} \\ \Pr_{s_1 \in D_1}[s_1 \cdot z \equiv \widehat{s}_0 \cdot z \pmod 2] &\geq \mu_1^*. \end{aligned}$$

One can verify that $\mu_1 \in [0, 1/2]$ for all g, and (the analogue of Proposition 5) that $\mu_1(g_{\text{mod } 3}) = 1/2 - O(1/n)$. We also have $\mu_1(g) \leq \mu(g)$, since if two events both occur with probability at least μ_1 then their disjunction does as well. On the other hand, $\mu_1(g)$ can be much smaller than $\mu(g)$, and indeed if g is partial then $\mu_1(g) = 0$ does not imply that g is a parity function. To illustrate, for even n let $g_{LR}(s) = 0$ if $s[1], \ldots, s[n/2]$ are all 0, and $g_{LR}(s) = 1$ if $s[n/2+1], \ldots, s[n]$ are all 0, under the promise that exactly one of these is the case. Then:

Proposition 8 $\mu(g_{LR}) = 1/2 + \Theta(2^{-n/2})$*, whereas* $\mu_1(g_{LR}) = 0$.

Proof. Let D_0 and D_1 be the uniform distributions over 0 and 1 inputs respectively. Choose any $z \neq 0^n$ for which (say) $z[n/2+1], \ldots, z[n]$ are all 0. Then $s_0 \cdot z \equiv 0 \pmod 2$ for every $s_0 \in g^{-1}(0)$, but

$$\Pr_{s_1 \in D_1}[s_1 \cdot z \equiv 1 \pmod 2] = \frac{1}{2 - 2^{-(n/2-1)}}$$

since $s[1], \ldots, s[n/2]$ are chosen uniformly at random conditioned on not being all 0. Hence $\mu(g_{LR}) = 1/2 + O(2^{-n/2})$. This is the best bound achievable, since if z has 1 bits on both the left and right halves then $s \cdot z$ is close to uniformly random for both $s \in D_0$ and $s \in D_1$.

In the one-sided case, for any s we can choose a $z \neq 0^n$ such that $|z| \equiv 1 \pmod 2$ and $s[i] = 1$ whenever $z[i] = 1$. Then $s \cdot z$ disagrees with $\widehat{s} \cdot z$ for every $\widehat{s}$ such that $g_{LR}(\widehat{s}) \neq g_{LR}(s)$, and hence $\mu_1(g_{LR}) = 0$. ■

Given the existence of such large gaps for partial g, it is remarkable that in the total case, even μ_1 must be bounded away from 0 if g is not a parity function. To show this, we will use the following lemma about finite groups.

Lemma 9 *Partition a finite group G into (A,B). Then either for all $a \in A$, there exist $b_1, b_2 \in B$ such that $a = b_1 \cdot b_2$, or for all $b \in B$, there exist $a_1, a_2 \in A$ such that $b = a_1 \cdot a_2$.*

Proof. Suppose without loss of generality that $1 \in A$. Then either $|A|-1 \leq |B|$ or $|A|-1 \geq |B|$. Suppose the former case, and choose an $a \in A$. Let the elements of G be vertices of a directed graph, with an edge from g_1 to g_2 if and only if $g_2 = g_1^{-1}a$. Then, since each vertex has indegree and outdegree 1, the graph is a union of disjoint cycles. Furthermore, after we remove the cycle $(1,a)$, there are more vertices in B than in A. Hence there exists an edge from some $b_1 \in B$ to $b_2 \in B$, and $b_1 b_2 = a$. Similarly, if $|A|-1 \geq |B|$, then for each $b \in B$, after we remove the cycle $(1,b)$ there are more vertices in A than in B. ■

We can now prove the main result; note that the constant we obtain is half that of Theorem 6.

Theorem 10 *Suppose g is total and $\mu_1(g) < 0.0732$. Then $\mu_1(g) = 0$ (equivalently, g is a parity function).*

Proof. As in Theorem 6, by duality there exists a distribution D over $z, \widehat{s} \in \{0,1\}^n$, $z \neq 0^n$ such that for all $s \in \{0,1\}^n$,

$$\Pr_{(z,\widehat{s}) \in D : g(s) \neq g(\widehat{s})} [s \cdot z \equiv \widehat{s} \cdot z \,(\operatorname{mod} 2)] < \mu.$$

It follows that there exist distributions $\widehat{D}_0$ and $\widehat{D}_1$ over $z \in \{0,1\}^n$, as well as functions $b_0(z), b_1(z) \in \{0,1\}$ (which we can take to be deterministic without loss of generality), such that for all $s_0 \in g^{-1}(0)$ and $s_1 \in g^{-1}(1)$,

$$\begin{aligned} \Pr_{z \in \widehat{D}_0} [s_0 \cdot z \equiv b_0(z) \,(\operatorname{mod} 2)] &> 1-\mu, \\ \Pr_{z \in \widehat{D}_1} [s_1 \cdot z \equiv b_1(z) \,(\operatorname{mod} 2)] &> 1-\mu. \end{aligned}$$

Observe that if s is drawn uniformly at random from $\{0,1\}^n$, then for any $b_i(z)$,

$$\Pr_s \left[\Pr_{z \in \widehat{D}_1} [s \cdot z \equiv b_i(z) \,(\operatorname{mod} 2)] > \frac{1}{2} \right] \leq \frac{1}{2}$$

by symmetry. Hence $\left|g^{-1}(i)\right| \leq 2^{n-1}$ for $i \in \{0,1\}$, and therefore $\left|g^{-1}(0)\right| = \left|g^{-1}(1)\right| = 2^{n-1}$.

Now identify $\{0,1\}^n$ with the group $\mathbb{Z}_2^n$. By Lemma 9, either for all $s_0 \in g^{-1}(0)$ there exist $s_\alpha, s_\beta \in g^{-1}(1)$ such that $s_0 = s_\alpha \oplus s_\beta$, or for all $s_1 \in g^{-1}(1)$ there exist $s_\alpha, s_\beta \in g^{-1}(0)$ such that $s_1 = s_\alpha \oplus s_\beta$. Suppose the former without loss of generality; then for all $s_0 \in g^{-1}(0)$,

$$s_0 \cdot z \equiv (s_\alpha \cdot z) \oplus (s_\beta \cdot z) \,(\operatorname{mod} 2).$$

Here we have the parity of two binary random variables, $s_\alpha \cdot z$ and $s_\beta \cdot z$, which are not necessarily independent, but both of which take the value $b_1(z)$ with probability greater than $1-\mu$. Hence, by the union bound,

$$\Pr_{z \in \widehat{D}_1} [s_0 \cdot z \equiv 1 \,(\operatorname{mod} 2)] < 2\mu.$$

for all $s_0 \in g^{-1}(0)$.

Furthermore, if s is drawn uniformly from $\{0,1\}^n$ then the expectation of $s \cdot z$ is $1/2$ under any distribution over $z \neq 0$. Thus, if s_1 is drawn uniformly from $g^{-1}(1)$ then

$$\Pr_{z \in \widehat{D}_1,\ s_1}[s_1 \cdot z \equiv 1 \pmod 2] > 1 - 2\mu.$$

It follows by the union bound that

$$\Pr_{z \in \widehat{D}_1}[b_1(z) \equiv 1 \pmod 2] > 1 - 3\mu.$$

and hence, for all $s_1 \in g^{-1}(1)$,

$$\Pr_{z \in \widehat{D}_1}[s_1 \cdot z \equiv 1 \pmod 2] > 1 - 4\mu.$$

We have established that g is a bounded-error threshold function of parity functions: there exist p_z summing to 1 such that for all s,

$$\Psi(s) = \sum_{z \in \{0,1\}^n \setminus \{0^n\}} p_z(s \cdot z) \text{ is } \begin{cases} > 1 - 4\mu & \text{if } g(s) = 1 \\ < 2\mu & \text{if } g(s) = 0. \end{cases}.$$

As in Theorem 6, we will draw s uniformly from $\{0,1\}^n$ and consider the variance $var(\Psi)$. Suppose $p_z < 1/2$ for all z, since otherwise g is a parity function. Then by pairwise independence,

$$var(\Psi) = (1/4)\sum_z p_z^2 < 1/8.$$

Also, we derived previously that if s is drawn uniformly from $g^{-1}(1)$, then the expectation of $\Psi(s)$ is at least $1 - 2\mu$. It follows that

$$var(\Psi) > (1/2 - 2\mu)^2.$$

Combining,

$$\mu > \left(2 - \sqrt{2}\right)/8 > 0.0732.$$

■

5 Concluding Remarks

An intriguing open problem is whether Theorem 4 can be proved using the polynomial method of Beals et al. [5], rather than the adversary method of Ambainis [3]. It is known that one can lower-bound polynomial degree in terms of *block sensitivity*, which is (roughly) the maximum number of disjoint changes to an input that change the output value. The trouble is that the RFS function has block sensitivity 1—the 'sensitive blocks' of each input tend to have small intersection, but are not disjoint. See [1] for more about the quantum query complexity of such functions.

We believe the constants of Theorems 6 and 10 can be improved. The smallest nonzero $\mu(g)$ and $\mu_1(g)$ values we know of are both attained by $g = OR(s[1], s[2])$:

Proposition 11 $\mu(OR) = \mu_1(OR) = 1/3$.

Proof. First, $\mu_1(OR) \geq 1/3$, since D_1 can choose $s[1]\,s[2]$ to be 01, 10, or 11 each with probability 1/3; then for any $z \neq 0$ and $\widehat{s}_0 \in g^{-1}(0)$, $s \cdot z \not\equiv \widehat{s}_0 \cdot z$ with probability at most 2/3. Second, $\mu(OR) \leq 1/3$, since applying linear programming duality, we can let the first two bits of z and $\widehat{s}_1$, $(z[1]\,z[2], \widehat{s}_1[1]\,\widehat{s}_1[2])$, be $(01, 01)$, $(10, 10)$, or $(11, 10)$ each with probability 1/3, while $z[3], \ldots, z[n]$ are all 0. Then $s_0 \cdot z \not\equiv \widehat{s}_1 \cdot z$ always, and for any $s_1 \in g^{-1}(1)$, $s_1 \cdot z \not\equiv \widehat{s}_0 \cdot z$ with probability 2/3. ■

A final note. Define *affine recursive Fourier sampling*, or $ARFS_h$, similarly to RFS_h, except that at each level, we are guaranteed there exists a secret string s and additive constant $c \in \{0,1\}$ such that $A_n(x) = (s \cdot x + c) \pmod 2$ for all $x \in \{0,1\}^n$. As before, the goal at each level is to compute $g(s)$; we do not care about c. Clearly the algorithm of Lemma 1 works also for $ARFS_h$, since c is simply encoded into the phase when we apply the Hadamard gates. On the other hand, we can show a quantum lower bound for $ARFS_h^g$ (for any g) much more easily than for RFS_h^g. For observe that we can encode the parity of 2^h bits into $ARFS_h^g$; thus, by a result of Beals et al. [5], a quantum algorithm for $ARFS_h^g$ requires $\Omega(2^h)$ queries.

6 Acknowledgments

I thank Lisa Hales, Umesh Vazirani, and Ronald de Wolf for helpful comments.

References

[1] S. Aaronson. Quantum certificate complexity. Submitted, 2002. quant-ph/0210020.[6]

[2] S. Aaronson. *The Complexity Zoo.*
http://www.cs.berkeley.edu/~aaronson/zoo.html.

[3] A. Ambainis. Quantum lower bounds by quantum arguments. In *Proceedings of ACM STOC'2000*, pp. 636–643, 2000. Journal version to appear in *Journal of Computer and System Sciences.* quant-ph/0002066.

[4] L. Babai. Bounded round interactive proofs in finite groups. *SIAM Journal on Discrete Math* 5(1):88–111, 1992.

[5] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. In *Proceedings of IEEE FOCS'98*, pp. 352–361, 1998. quant-ph/9802049.

[6] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing* 26(5):1411–1473, 1997.

[6] Available at `www.arxiv.org`.

[7] D. Boneh and R. Lipton. Algorithms for black box fields and their application to cryptography. In *Proceedings of CRYPTO'96*, Lecture Notes in Computer Science Vol. 1109, Springer-Verlag, pp. 283–297, 1996.

[8] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London* A454:339–354, 1998. quant-ph/9708016.

[9] W. van Dam, S. Hallgren, and L. Ip. Algorithms for hidden coset problems. To appear in *Proceedings of SODA'2003*, 2003.

[10] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London* A439:553–558, 1992.

[11] M. Furst, J. Saxe, and M. Sipser. Parity, circuits, and the polynomial time hierarchy. In *Proceedings of IEEE FOCS'81*, pages 260–270, 1981.

[12] F. Green and R. Pruim. Relativized separation of EQP from P^{NP}. *Information Processing Letters* 80(5):257–260, 2001.

[13] N. Linial, Y. Mansour, and N. Nisan. Constant depth circuits, Fourier transform, and learnability. *Journal of the ACM* 40(3):607–620, 1993.

[14] A. A. Razborov. Lower bounds for the size of circuits of bounded depth with basis $\{\&, \oplus\}$. *Mathematicheskie Zametki* 41:4, pp. 598–607, 1987. English translation in *Math. Notes. Acad. Sci. USSR* 41:4, pp. 333–338, 1987.

[15] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* 26(5):1484–1509, 1997. quant-ph/9508027.

[16] D. Simon. On the power of quantum computation. In *Proceedings of IEEE FOCS'94*, pp. 116–123, 1994.

[17] R. Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proceedings of ACM STOC'87*, pp. 77–82, 1987.

[18] U. Vazirani. Personal communication, 2001.

[19] J. Watrous. Succinct quantum proofs for properties of finite groups. In *Proceedings of IEEE FOCS'2000*, pp. 537–546, 2000.