Adversary-Based Parity Lower Bounds with Small Probability Bias

Badih Ghazi

badih@mit.edu

1 Introduction

The quantum query complexity of the parity function on n variables is known to be exactly $n/2$, for any error probability $\epsilon < 1/2$. The upper bound follows from a modification of Deutsch's quantum algorithm that computes the parity using $n/2$ queries and with zero probability of error [CEMM98]. Also, the lower bound of $n/2$ can be obtained by a simple application of the polynomial method, and the proof holds for any error probability strictly less than $1/2$ [MNR11]. In particular, it holds for $\epsilon = 1/2 - q(n)$ where $q(n)$ is exponentially small. The question that initially motivated this project is: Can we recover this result using adversary-based arguments ?

The general additive adversary bound $Adv^{\pm}(f)$ was shown to asymptotically characterize the bounded-error quantum query complexity [Rei10]. However, this bound is very loose for the regime where $\epsilon = 1/2 - q(n)$ and $q(n) = o(1)$. In fact, the resulting lower bound is:

$$Q_{\epsilon}(Parity_n) \geq (\frac{1}{2} - \sqrt{\epsilon(1-\epsilon)})Q_2(Parity_n)$$

Setting $\epsilon = 1/2 - q(n)$ and noting that $Q_2(Parity_n) = \theta(n)$, we get that:

$$Q_{1/2-q(n)}(Parity_n) = \Omega((q(n))^2 n)$$

Thus, even for $q(n)$ as large as $1/\sqrt{n}$, this approach yields a trivial constant lower bound.

In this report, we will prove a lower bound of $\Omega(n)$ on the quantum query complexity of the parity function, that holds for any error probability $\epsilon = 1/2 - \Omega(e^{-d \cdot n})$ (where $d > 0$ is a fixed constant) and that is proven using adversary-based arguments. The proof further uses a "quantum reduction" to the t-fold search problem, which was itself studied by [Amb05] and [Spa08] using the earliest version of the multiplicative adversary method. This lower bound also holds for weaker parity algorithms, where a set of averages (over inputs) of error probabilities is required to be at most $1/2 - q(n)$, as opposed to the error probability on every input being at most $1/2 - q(n)$.

Shortly before the deadline, Robert Spalek pointed out to me a very recent result of Lee and Roland [LR12], where they prove a XOR lemma for quantum query complexity. This result, whose proof builds on recent work on quantum state generation and conversion algorithms, implies a lower bound on the quantum query complexity of the parity function even for error probabilities exponentially close to $1/2$. In the last section of this report, we show some consequences of this XOR lemma, related to parity algorithms which are required to succeed only on a subset of the Boolean hypercube and with error probability very close to $1/2$.

1.1 Outline

In Section 2, we briefly present the result of Ambainis and Spalek concerning the t-fold search problem that we will use. In Section 3, we prove the adversary-based lower bound for the parity function when the probability bias is exponentially small. In Section 4, we show some consequences of the XOR lemma that was proven in [LR12].

1.2 Notation

For every $n \in \mathbb{N}$, let $Q_n = \{0,1\}^n$ be the Boolean hypercube of dimension n and let $Par_n : Q_n \to \{0,1\}$ denote the parity function on n variables. For any $x \in Q_n$, let $wh(x)$ be the Hamming weight of x, i.e. the number of non-zero coordinates of x. We let $[n]$ denote the set $\{1,2,..,n\}$ and for any $b \in \{0,1\}$, b' denotes the complement of b. For any $x \in Q_n$ and any subset $S \subset [n]$, we let $x|_S$ be the restriction of x to the subset S of indices, i.e. $(x|_S)_i = x_i$ for all $i \in S$ and $(x|_S)_i = 0$ otherwise.

2 Lower bound for the t-fold search problem

First, we recall the t-fold search problem.

Definition 2.1. *($Search_{t,n}$)*
Let $t \in \mathbb{N}$ and let $\mathcal{B}_t = \{x \in Q_n \mid wh(x) = t\}$. For every $x \in \mathcal{B}_t$, $Search_{t,n}(x)$ is the unique subset J of $[n]$ of cardinality t s.t. $x_i = 1$ if and only if $i \in J$.

Ambainis [Amb05] showed that the quantum query complexity of the $Search_{t,n}$ is $\Omega(\sqrt{tn})$ even if the error probability is allowed to be as large as $1 - \Omega(e^{-t/8})$. Ambainis's method, which was based on the analysis of the eigenspaces of the density matrix, was later generalized by Spalek [Spa08], into what he coined, the "multiplicative quantum adversary".

Theorem 2.2. *([Amb05], [Spa08], [AŠDW09])*
For every $t \leq \frac{n}{4e}$ and every $\epsilon = 1 - \Omega(e^{-t/8})$, $Q_\epsilon(Search_{t,n}) = \Omega(\sqrt{tn})$.

3 Lower bounds on the quantum query complexity of the parity function with small probability bias

In this section, we prove a lower bound of $\Omega(n)$ on the quantum query complexity of the parity function. The proof uses a "quantum reduction" to the t-fold search problem. The proof of the reduction is an adaptation to our setup of the proof, given by [CVDNT99], of the lower bound on the bounded-error quantum communication complexity of the inner product function. The lower bound that we obtain applies to a family of "weak" parity algorithms, where a set of averages (over inputs) of error probabilities is required to be at most $1/2 - q(n)$, as opposed to the error probability on every input being at most $1/2 - q(n)$.

Notation 3.1. *(Characteristic vector)*
Let $n \in \mathbb{N}$ and $S \subset [n]$. Then, the characteristic vector $\chi_S \in Q_n$ of S is given by:

$$(\chi_S)_i = \begin{cases} 1 & \text{if } i \in S. \\ 0 & \text{otherwise.} \end{cases}$$

for all $i \in [n]$.

The following lemma states an observation (due to Professor Scott Aaronson) which is central to the reduction. Note that this relation was used in the Bernstein-Vazirani problem [BV97] and in the lower bound on the bounded-error quantum communication complexity of the inner product function [CVDNT99].

Lemma 3.2. *For every* $n \in \mathbb{N}$*, let* H_n *denote the* $2^n \times 2^n$ *Hadamard matrix. Then, for all* $x \in Q_n$*, we have*

$$\frac{1}{\sqrt{2^n}} \sum_{S \subset [n]} (-1)^{Par_n(x|_S)} H_n |\chi_S\rangle = |x\rangle \tag{1}$$

Proof.

$$\begin{aligned}
\frac{1}{\sqrt{2^n}} \sum_{S \subset [n]} (-1)^{Par_n(x|_S)} H_n |\chi_S\rangle &= \frac{1}{\sqrt{2^n}} \sum_{S \subset [n]} (-1)^{Par_n(x|_S)} \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{Par_n(y|_S)} |y\rangle \\
&= \frac{1}{2^n} \sum_{y \in \{0,1\}^n} \sum_{S \subset [n]} (-1)^{Par_n(x|_S + y|_S)} |y\rangle \\
&= \sum_{y \in \{0,1\}^n} \big(\frac{1}{2^n} \sum_{S \subset [n]} (-1)^{Par_n((x+y)|_S)} \big) |y\rangle
\end{aligned}$$

But,

$$\frac{1}{2^n} \sum_{S \subset [n]} (-1)^{Par_n((x+y)|_S)} = \begin{cases} 1 & \text{if } x + y = 0 \mod 2, \text{ i.e. if } x = y. \\ 0 & \text{if } x + y \neq 0 \mod 2, \text{ i.e. if } x \neq y. \end{cases}$$

Thus, we conclude that

$$\frac{1}{\sqrt{2^n}} \sum_{S \subset [n]} (-1)^{Par_n(x|_S)} H_n |\chi_S\rangle = |x\rangle$$

□

The following definition is similar to the notion of "clean computation" used by [CVDNT99].

Definition 3.3. *(Coherent computation of the parity function)*
A quantum query algorithm $\mathcal{A}$ *computing* Par_n *is said to be coherent if, on inputs* $x \in Q_n$ *and* $S \subset [n]$*,* $\mathcal{A}$ *takes the state* $|x\rangle|\chi_S\rangle|z\rangle$ *to the state* $|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle$*.*

Loosely, a coherent computation of the parity function determines the parity of the restriction of the input string to the input subset and stores it without leaving any "garbage".

Lemma 3.4. *If there is a coherent algorithm* $\mathcal{A}$ *that computes* Par_n *using* $r(n)$ *queries, then there is a quantum query algorithm* $\mathcal{B}$ *that, on input* $x \in Q_n$*, takes the state* $|x\rangle|0\rangle^{\otimes n}|1\rangle$ *to the state* $|x\rangle|x\rangle|1\rangle$ *using* $r(n)$ *queries.*

Proof. First, $\mathcal{B}$ Hadamards the last $n+1$ qubits of the start state $|x\rangle|0\rangle^{\otimes n}|1\rangle$, which gives the state

$$|\psi_1\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S) \in \{0,1\}^{n+1}} (-1)^z |x\rangle|\chi_S\rangle|z\rangle$$

Then, $\mathcal{B}$ runs the coherent algorithm $\mathcal{A}$ on $|\psi_1\rangle$ which performs $r(n)$ queries and gives the state

$$|\psi_2\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z |x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle$$

Changing variables, we get

$$|\psi_2\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^{z+Par_n(x|_S)} |x\rangle|\chi_S\rangle|z\rangle$$

Then, $\mathcal{B}$ Hadamards the last $n+1$ qubits of $|\psi_2\rangle$ which gives the state:

$$\begin{aligned}
|\psi_3\rangle &= \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^{z+Par_n(x|_S)} |x\rangle H_n|\chi_S\rangle H_1|z\rangle \\
&= \frac{1}{\sqrt{2^n}} \sum_{\chi_S\in\{0,1\}^n} (-1)^{Par_n(x|_S)} |x\rangle H_n|\chi_S\rangle \frac{1}{\sqrt{2}} \sum_{z\in\{0,1\}} (-1)^z H_1|z\rangle \\
&= |x\rangle|x\rangle|1\rangle \text{ (by Lemma 3.2)}
\end{aligned}$$

□

Note that if we follow the algorithm $\mathcal{B}$ in Lemma 3.4 by a measurement of the middle n qubits in the standard basis, we get the original input x with probability 1.

The following lemma shows how to deal with the general case and its proof is based on the treatment by [CVDNT99] of the bounded-error quantum communication complexity of the inner product function.

Lemma 3.5. *Let $f : \mathcal{D} \to \Sigma_O$ where $\mathcal{D} \subset Q_n$ and Σ_O is a finite set. If there exists a quantum query algorithm $\mathcal{A}$ that computes Par_n using $r(n)$ queries and with error probability $p_x(n)$ on input $x \in Q_n$ and s.t.*

$$\frac{1}{2^n} \sum_{S\subset\{0,1\}^n} p_{x|_S}(n) \leq 1/2 - q(n) \tag{2}$$

(for some function $q : \mathbb{N} \to \mathbb{R}$), then there exists a quantum query algorithm $\mathcal{C}$ that computes f using $2r(n)$ queries and with success probability at least $4q^2(n)$.

Proof. Let $\mathcal{A}$ be an algorithm that computes Par_n using $r(n)$ queries and that satisfies Equation 2. Then, $\mathcal{A}$ can be transformed into an algorithm $\mathcal{B}$ that takes as inputs $x \in Q_n$ (to which it has oracle access) and $S \subset [n]$ and outputs the parity of $x|_S$ with error probability equal to $p_{x|_S}(n)$. Moreover, $\mathcal{A}$ can be assumed to measure one qubit and output the obtained value as the (Boolean) answer. Those two assumptions are due to the fact that $\mathcal{B}$ can use additional qubits (not used by $\mathcal{A}$) to store any new garbage. Thus, on inputs $x \in Q_n$ and $S \subset [n]$, $\mathcal{B}$ takes the state $|x\rangle|\chi_S\rangle|z\rangle|0\rangle|0\rangle^{\otimes w}$ to the state

$$|\psi_1\rangle = a_{x,S}|x\rangle|\chi_S\rangle|z\rangle|Par_n(x|_S)\rangle|J_{x,S}\rangle + b_{x,S}|x\rangle|\chi_S\rangle|z\rangle|Par'_n(x|_S)\rangle|K_{x,S}\rangle$$

where $|b_{x,S}|^2 = p_{x|_S}(n)$, $|a_{x,S}|^2 + |b_{x,S}|^2 = 1$ and $|J_{x,S}\rangle$ and $|K_{x,S}\rangle$ are arbitrary unit vectors. Applying a $CNOT$ operation (controlled by the answer register) on $|\psi_1\rangle$ gives:

$$\begin{aligned}
|\psi_2\rangle =& a_{x,S}|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|Par_n(x|_S)\rangle|J_{x,S}\rangle + b_{x,S}|x\rangle|\chi_S\rangle|z + Par'_n(x|_S)\rangle|Par'_n(x|_S)\rangle|K_{x,S}\rangle \\
&= a_{x,S}|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|Par_n(x|_S)\rangle|J_{x,S}\rangle + b_{x,S}|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|Par'_n(x|_S)\rangle|K_{x,S}\rangle \\
&\quad - b_{x,S}|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|Par'_n(x|_S)\rangle|K_{x,S}\rangle + b_{x,S}|x\rangle|\chi_S\rangle|z + Par'_n(x|_S)\rangle|Par'_n(x|_S)\rangle|K_{x,S}\rangle
\end{aligned}$$

Running the algorithm $\mathcal{B}$ in reverse on $|\psi_2\rangle$ yields the state

$$|\psi_3\rangle = |x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|0\rangle^{\otimes(w+1)} + \sqrt{2}b_{x,S}|M_{x,S,z}\rangle$$

where $|M_{x,S,z}\rangle = \mathcal{B}^{-1}|x\rangle|\chi_S\rangle(\frac{1}{\sqrt{2}}|z + Par'_n(x|_S)\rangle - \frac{1}{\sqrt{2}}|z + Par_n(x|_S)\rangle)|Par'_n(x|_S)\rangle|K_{x,S}\rangle$.
Moreover, the set $\{|M_{x,S,z}\rangle\}_{x,S,z}$ satisfies the following 2 properties:

1. For every $x \in Q_n$ and every $S \subset [n]$, $|M_{x,S,z}\rangle$ is odd in z i.e. $|M_{x,S,0}\rangle = -|M_{x,S,1}\rangle$.
2. The set $\{|M_{x,S,0}\rangle\}_{x,S}$ is an orthonormal set because the algorithm $\mathcal{B}$ does not modify the contents of the registers $|x\rangle$ and $|\chi_S\rangle$ (and thus the same is true for $\mathcal{B}^{-1}$).

By Lemma 3.4, we can see that $|\psi_3\rangle$ is the sum of the output of a coherent parity algorithm (when run on x and S) and a residual "garbage" term (namely, $\sqrt{2}b_{x,S}|M_{x,S,z}\rangle$) and which is not necessarily orthogonal to the first term. Having transformed any parity algorithm to the form above with only a factor of 2 increase in the number of queries, we now show how the reduction works for general parity algorithms. More precisely, given an input $x \in Q_n$, we will show how we can, using 1 call to the algorithm $\mathcal{B}$ above, recover the input x with probabilty at least $4q^2(n)$. The algorithm $\mathcal{C}$ starts with the state $|x\rangle|0\rangle^{\otimes n}|1\rangle|0\rangle^{\otimes w}$ and first Hadamards the middle $n+1$ qubits, which gives the state

$$|\varphi_1\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z|x\rangle|\chi_S\rangle|z\rangle|0\rangle^{\otimes w}$$

Then, $\mathcal{C}$ runs the algorithm $\mathcal{B}$ described above on $|\varphi_1\rangle$ which gives the state:

$$\begin{aligned}
|\varphi_2\rangle &= \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z\mathcal{B}|x\rangle|\chi_S\rangle|z\rangle|0\rangle^{\otimes w} \\
&= \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z(|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|0\rangle^{\otimes(w+1)} + \sqrt{2}b_{x,S}|M_{x,S,z}\rangle) \\
&= \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z|x\rangle|\chi_S\rangle|z + Par_n(x|_S)\rangle|0\rangle^{\otimes(w+1)} + \frac{1}{\sqrt{2^{n+1}}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z\sqrt{2}b_{x,S}|M_{x,S,z}\rangle
\end{aligned}$$

Applying a Hadamard operation to the middle $n+1$ qubits of $|\varphi_2\rangle$, we get the state

$$\begin{aligned}
|\varphi_3\rangle &= |x\rangle|x\rangle|1\rangle|0\rangle^{\otimes(w+1)} + U\frac{1}{\sqrt{2^n}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z b_{x,S}|M_{x,S,z}\rangle \quad \text{(By Lemma 3.4)} \\
&= |x\rangle|x\rangle|1\rangle|0\rangle^{\otimes(w+1)} + U|\Lambda\rangle
\end{aligned}$$

where $|\Lambda\rangle = \frac{1}{\sqrt{2^n}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z b_{x,S}|M_{x,S,z}\rangle$ and where $U = I^{\otimes n} \otimes H_{n+1} \otimes I^{\otimes(w+1)}$ is a unitary operation. The next step of $\mathcal{C}$ is to measure the second set of n qubits in the standard basis. Since $|x\rangle|x\rangle|1\rangle|0\rangle^{\otimes(w+1)}$ and $U|\Lambda\rangle$ are not necessarily orthogonal, the outcome of the measurement might be different from x with some non-zero probability. In order to find the cosine of the angle θ between $|\varphi_3\rangle$ and

$|x\rangle|x\rangle|1\rangle|0\rangle^{\otimes(w+1)}$, we first upper bound the squared magnitude of $U|\Lambda\rangle$:

$$\begin{aligned}
\|U|\Lambda\rangle\|_2^2 &= \||\Lambda\rangle\|_2^2 \quad \text{(since } U \text{ is a unitary matrix)} \\
&= \|\frac{1}{\sqrt{2^n}} \sum_{(z,\chi_S)\in\{0,1\}^{n+1}} (-1)^z b_{x,S}|M_{x,S,z}\rangle\|_2^2 \\
&= \|\frac{2}{\sqrt{2^n}} \sum_{\chi_S\in\{0,1\}^n} b_{x,S}|M_{x,S,0}\rangle\|_2^2 \quad \text{(since } |M_{x,S,0}\rangle = -|M_{x,S,1}\rangle \text{ for all } S\subset[n] \text{ and all } x\in Q_n\text{)} \\
&= \frac{4}{2^n} \sum_{\chi_S\in\{0,1\}^n} |b_{x,S}|^2 \quad \text{(since } \{|M_{x,S,0}\rangle\}_{x,S} \text{ is an orthonormal set)} \\
&= \frac{4}{2^n} \sum_{\chi_S\in\{0,1\}^n} p_{x|_S}(n) \leq 4(\frac{1}{2} - q(n))
\end{aligned}$$

But, $\||\varphi_3\rangle\|_2^2 + \||x\rangle|x\rangle|1\rangle|0\rangle^{\otimes(w+1)}\|_2^2 - 2\cos\theta = \|U|\Lambda\rangle\|_2^2$ which implies that $\cos\theta \geq 2q(n)$. Therefore, $Pr[\text{Obtaining } x] = (\cos\theta)^2 \geq 4q^2(n)$. Thus, with probability at least $4q^2(n)$ and using at most $2r(n)$ queries, we can compute the value $f(x)$ of any function $f : \mathcal{D} \to \Sigma_O$ where $\mathcal{D} \subset Q_n$ and Σ_O is a finite set. □

Theorem 3.6. *Let $t = c \cdot n$ for some constant $0 < c \leq 1/(4e)$ and let $q(n) = \Omega(e^{-t/16})$. If $\mathcal{A}$ is a quantum query algorithm that computes Par_n with error probability $p_x(n)$ on input $x \in Q_n$ and s.t.*

$$\frac{1}{2^n} \sum_{S\subset\{0,1\}^n} p_{x|_S}(n) \leq 1/2 - q(n) \tag{3}$$

for all x of Hamming weight t, then $\mathcal{A}$ should make $\Omega(n)$ queries.

Proof. We apply Lemma 3.5 with the function f set to the t-fold search function $Search_{t,n}$ and we get that, given $x \in \mathcal{B}_t$ (where $\mathcal{B}_t$ is the set of all strings in Q_n of Hamming weight t), we can compute $Search_{t,n}(x)$ with probability at least $4q^2(n)$ and by making twice as many queries as $\mathcal{A}$ does. Since $q(n) = \Omega(e^{-t/16})$, Theorem 2.2 gives a lower bound of $\Omega(n)$ on the $(1 - \Omega(q^2(n)))$-error quantum query complexity of the t-fold search problem. Thus, the parity algorithm $\mathcal{A}$ cannot have query complexity $o(n)$. □

Note that the reason why the required conditions on the error probabilities in Equation 3 treat strings of low weights differently than those of large weights, is that our reduction in Lemma 3.5 is not "symmetric". However, knowing that the parity function is symmetric, one can "symmetrize" the conditions in Theorem 3.6.

As a special case of Theorem 3.6, the next corollary shows a lower bound on the quantum query complexity of the parity function.

Corollary 3.7. *For any constant $d \leq 1/(64e)$, $Q_{\frac{1}{2}-e^{-d\cdot n}}(Par_n) = \Omega(n)$.*

Proof. First, note that if for all $x \in Q_n$, $p_x(n) \leq 1/2 - q(n)$, then for every $x \in Q_n$:

$$\frac{1}{2^n} \sum_{S\subset\{0,1\}^n} p_{x|_S}(n) \leq 1/2 - q(n)$$

Applying Theorem 3.6 with $c = 1/(4e)$, we get the desired claim as long as $d \leq c/16 = 1/(64e)$. □

4 Some consequences of the XOR lemma

Recently, [LR12] proved a XOR lemma for the quantum query complexity. Loosely, the XOR lemma says that if one is to compute the parity of k values of a Boolean-valued function f, then at least k times the number of queries required to compute f on 1 instance are needed, even if the error probability is allowed to be exponentially close to $1/2$. The proof of [LR12] uses the recent result of [AMRR11] who proved that the multiplicative adversary is stronger than the general additive adversary, as well as recent work on quantum state generation and conversion by [AMRR11] and [LMR$^+$11].

Lemma 4.1. *(The XOR lemma [LR12])*
Let f be a Boolean-valued function, $0 \leq \delta \leq 1$ and $k \in \mathbb{N}$. Then $Q_{(1-\delta^{k/2})/2}(\oplus \circ f^{(k)}) \geq \frac{k\delta}{8} Adv^{\pm}(f)$.

Note that in the statement of the XOR lemma, f is a general partial function and need not be total. Given the XOR lemma, we can deduce the following corollary:

Corollary 4.2. *Let $\delta > 0$. Any algorithm computing the parity function on n variables with error probability at most $(1 - \delta^{k/2})/2$ should make $\Omega(n)$ queries to the oracle.*

Proof. Follows from Lemma 4.1 by letting $k = n$ and the "base function" f be the identity function on 1 bit, which gives that $Adv^{\pm}(f) \geq 1$. □

Lemma 4.3. *Let $\mathcal{A} \subset Q_n$. If $\mathcal{A}$ contains a subcube of dimension $d(n)$, then the query complexity of any algorithm that computes the parity of every string in $\mathcal{A}$ with error probability at most $(1-\delta^{d(n)/2})/2$ (where $0 < \delta \leq 1$ is any constant) must perform $\Omega(d(n))$ queries.*

Proof. Assume that $\mathcal{A} \subset Q_n$ contains a subcube of dimension $d(n)$. Then, there exist indices $i_1, i_2, .., i_{d(n)} \in [n]$ s.t. if $y \in Q_n$ satisfies $y_j = x_{i_j}$ for all $j \in [d(n)]$, then $y \in \mathcal{A}$. Letting the "base function" in Lemma 4.1 be the parity function on 1 bit (i.e. the identity) and letting $k = d(n)$, we see that every algorithm that computes the parity of every string in $\mathcal{A}$ with error probability at most $(1-\delta^{d(n)/2})/2$ can compute the value of the function $\oplus \circ f^{(k)} : \{0,1\}^{d(n)} \to \{0,1\}$ on every input with error probability at most $(1 - \delta^{d(n)/2})/2$. By Lemma 4.1, such an algorithm should make $\Omega(d(n))$ queries. □

The previous lemma raises the following combinatorial question, which does not seem to have been addressed:

Question 4.4. *Let $n \in \mathbb{N}$ and let $1 \leq d(n) \leq n$. What is the number of subsets of Q_n that contain at least one subcube of dimension $d(n)$?*

One would be tempted to try to prove that if the cardinality of $\mathcal{A}$ is sufficiently large, then $\mathcal{A}$ should contain a subcube of non-negligible dimension. However, such an approach would be limited by the following result of Alon et. al.

Theorem 4.5. *[AKS07] For every constant $d \in \mathbb{N}$, there exists a subset $\mathcal{A}$ of Q_n of size $\geq 2^n - 2^n/d$ and that does not contain a subcube of dimension d.*

For instance, for $d = 2$, Theorem 4.5 says that there exists a subset of Q_n of cardinality at least $\frac{2}{3}2^n$ and that does not contain any hypercube of dimension 2!
We now further use the XOR lemma to prove a lower bound on the $(1/2-q(n))$- quantum query complexity of algorithms computing Par_n on certain types of subsets of Q_n and with $q(n)$ exponentially small.

Definition 4.6. *(Product subsets of Q_n)*
Let $t \in \mathbb{N}$. A subset $\mathcal{A}$ of Q_n is said to be a product subset of order t if $\mathcal{A} = \mathcal{D}^t$ for some $\mathcal{D} \subset Q_{\frac{n}{t}}$ s.t. there exist $x_0, x_1 \in \mathcal{D}$ with $Par_{\frac{n}{t}}(x_0) = 0$ and $Par_{\frac{n}{t}}(x_1) = 1$.

Loosely, product subsets are cartesian products of "partial non-trivial subsets", where "partial" means that they are not necessarily subcubes and "non-trivial" means that not all their elements have the same parity.
We now use the XOR lemma to show a lower bound on the quantum query complexity of algorithms that are only required to succeed on a product subset of Q_n and with error probability at most $(1/2 - q(n))$, where $q(n)$ is exponentially small. Note that this implies a lower bound for algorithms that are only required to succeed on a set that contains such a product subset.

Lemma 4.7. *Let $g : \mathbb{N} \to \mathbb{N}$ s.t. $1 \leq g(n) \leq n$ for all $n \in \mathbb{N}$. For every $\mathcal{A} \subset Q_n$ containing a product subset of order $g(n)$, any algorithm succeeding on all the binary strings in $\mathcal{A}$ with error probability at most $(1 - \delta^{g(n)/2})/2$ (where $0 < \delta \leq 1$ is any constant) must perform $\Omega(g(n))$ queries.*

Proof. Any such algorithm $\mathcal{B}$ should succeed on the product subset $\mathcal{D}^{g(n)}$ of order $g(n)$. Thus, if we let the "base function" in Lemma 4.1 be the function $f : \mathcal{D} \to \{0,1\}$ given by $f(x) = Par_{\frac{n}{g(n)}}(x)$ for all $x \in \mathcal{D}$, then $Adv^{\pm}(f) \geq 1$ since $Par_{\frac{n}{g(n)}}$ is not constant on $\mathcal{D}$. Letting $k = g(n)$, we see that every algorithm that computes the parity of every string in $\mathcal{D}$ with error probability at most $(1 - \delta^{g(n)/2})/2$ must be computing the value of the function $\oplus \circ f^{(k)} : \{0,1\}^n \to \{0,1\}$ on every input with error probability at most $(1 - \delta^{g(n)/2})/2$. By Lemma 4.1, such an algorithm should make $\Omega(g(n))$ queries. □

As a sidenote, we note the following lemma:

Lemma 4.8. *(Number of product subsets)*
For every $n, t \in \mathbb{N}$ with $1 \leq t \leq n$, the number of product subsets of Q_n of order t is exactly $2^{2^t} - 2\sqrt{2^{2^t}} + 1$.

Proof. First, note that the number of strings in Q_t of even parity is exactly $\frac{2^t}{2}$. Thus, the number of subsets of Q_t consisting only of strings having even parity (or only of strings having odd parity) is exactly $2^{2^{t-1}} - 1$, where the "-1" is to exclude the empty set. Thus, the claim follows. □

Finally, we conclude with the following question:

Question 4.9. *If for some subset $A \subset [n]$, the bounded error quantum query complexity is $Q_2(Parity|_A) = \Omega(r(n))$ for some function r. Is it true that the quantum query complexity will still be $\Omega(r(n))$, even if the error probability is allowed to be as large as $1/2 - q(n)$ for some $q(n) = o(1)$?*

Acknowledgments

I would like to thank Robert Spalek for pointing out to me the work of [LR12]. Also, thanks to Professor Scott Aaronson and to Adam Bouland for very helpful discussions.

References

[AKS07] N. Alon, A. Krech, and T. Szabó. Turan's theorem in the hypercube. *SIAM Journal on Discrete Mathematics*, 21(1):66–72, 2007.

[Amb05] A. Ambainis. A new quantum lower bound method, with an application to strong direct product theorem for quantum search. *arXiv preprint quant-ph/0508200*, 2005.

[AMRR11] A. Ambainis, L. Magnin, M. Roetteler, and J. Roland. Symmetry-assisted adversaries for quantum state generation. In *Computational Complexity (CCC), 2011 IEEE 26th Annual Conference on*, pages 167–177. IEEE, 2011.

[AŠDW09] A. Ambainis, R. Špalek, and R. De Wolf. A new quantum lower bound method, with applications to direct product theorems and time-space tradeoffs. *Algorithmica*, 55(3):422–461, 2009.

[BV97] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.

[CEMM98] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998.

[CVDNT99] R. Cleve, W. Van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. *Quantum Computing and Quantum Communications*, pages 61–74, 1999.

[LMR$^+$11] T. Lee, R. Mittal, B.W. Reichardt, R. Spalek, and M. Szegedy. Quantum query complexity of state conversion. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 344–353. IEEE, 2011.

[LR12] T. Lee and J. Roland. A strong direct product theorem for quantum query complexity. In *Computational Complexity (CCC), 2012 IEEE 27th Annual Conference on*, pages 236–246. IEEE, 2012.

[MNR11] A. Montanaro, H. Nishimura, and R. Raymond. Unbounded-error quantum query complexity. *Theoretical Computer Science*, 412(35):4619–4628, 2011.

[Rei10] B. Reichardt. Span programs and quantum query algorithms. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 17, page 110, 2010.

[Spa08] R. Spalek. The multiplicative quantum adversary. In *Computational Complexity, 2008. CCC'08. 23rd Annual IEEE Conference on*, pages 237–248. IEEE, 2008.