Lower Bounds on Stabilizer Rank

Shir Peleg[1], Amir Shpilka[1], and Ben Lee Volk[2]

[1]Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel
[2]Efi Arazi School of Computer Science, Reichman University, Herzliya, Israel

The *stabilizer rank* of a quantum state ψ is the minimal r such that $|\psi\rangle = \sum_{j=1}^{r} c_j |\varphi_j\rangle$ for $c_j \in \mathbb{C}$ and stabilizer states φ_j. The running time of several classical simulation methods for quantum circuits is determined by the stabilizer rank of the n-th tensor power of single-qubit magic states.

We prove a lower bound of $\Omega(n)$ on the stabilizer rank of such states, improving a previous lower bound of $\Omega(\sqrt{n})$ of Bravyi, Smith and Smolin [7]. Further, we prove that for a sufficiently small constant δ, the stabilizer rank of any state which is δ-close to those states is $\Omega(\sqrt{n}/\log n)$. This is the first non-trivial lower bound for approximate stabilizer rank.

Our techniques rely on the representation of stabilizer states as quadratic functions over affine subspaces of $\mathbb{F}_2^n$, and we use tools from analysis of boolean functions and complexity theory. The proof of the first result involves a careful analysis of directional derivatives of quadratic polynomials, whereas the proof of the second result uses Razborov-Smolensky low degree polynomial approximations and correlation bounds against the majority function.

1 Introduction

The conventional wisdom is that quantum computers are more powerful than classical computers. Among other reasons, this belief is supported by the fact that quantum computers are able to efficiently solve problems such as integer factorization [22], which are believed by some to be hard for classical computers; by provable black box separations [23, 11, 3, 20]; and by quantum computers' advantage in solving certain sampling problems that are deemed intractable for classical computers under well established complexity theoretic conjectures [1].

There is, however, very little that we can *unconditionally* prove with regard to the impossibility of efficiently simulating quantum computers using classical computers. Indeed, barring a computational complexity theoretic breakthrough, such as — at the very least — separating P from PSPACE, we can't hope to prove general and unconditional impossibility results.

Nevertheless, it remains an interesting and important problem to prove lower bounds on the running time of certain restricted types of simulation techniques for quantum circuits. One such result is a lower bound of Huang, Newman and Szegedy [12], who prove

Shir Peleg: shirpele@tauex.tau.ac.il, The research leading to these results has received funding from the Israel Science Foundation (grant number 514/20) and from the Len Blavatnik and the Blavatnik Family foundation.

Amir Shpilka: shpilka@tauex.tau.ac.il, The research leading to these results has received funding from the Israel Science Foundation (grant number 514/20) and from the Len Blavatnik and the Blavatnik Family foundation.

Ben Lee Volk: benleevolk@gmail.com, Part of this work was done while at the Department of Computer Science, University of Texas at Austin, USA, Supported by NSF Grant CCF-1705028.

unconditional exponential lower bounds for a subclass of simulators they call *monotone simulators*, which includes many, but not all, of the known simulation techniques.

Simulation algorithms based on *stabilizer rank decompositions* for quantum circuits dominated by Clifford gates [7, 5, 4] is a recent powerful class of algorithms for classically simulating quantum circuits (which is not covered by the lower bound of [12]). The computational cost of these algorithms is dominated by a certain natural algebraic and complexity-theoretic rank measure for quantum states, which we now define.

1.1 Clifford Circuits, Magic States and Stabilizer Rank

Clifford circuits are quantum circuits which only apply Clifford gates (for background on the Clifford group and the definitions of the type of gates considered in this paper, see Appendix A). Equivalently, such circuits only use CNOT, Hadamard, and phase gates. This is an important class of quantum circuits which, by the Gottesman-Knill theorem [10, 2], *can* be efficiently simulated (on, say, the input $|0^n\rangle$) by a classical algorithm. This highly non-obvious theorem follows from the fact that such circuits can only maintain certain states known as *stabilizer states*. These can be succinctly represented, and it is easy to track the state and update the succinct representation after any application of a Clifford gate.

Adding T gates (we refer again to Appendix A for the definition) on top of the Clifford gates results in a universal quantum gate set, that is, a set which can approximate every unitary operation. It is then possible, using a simple gadget-based transformation, to "push the T gates to the inputs" and obtain an equivalent circuit, of roughly the same size, which only uses Clifford operations, and is given, as additional auxiliary inputs, sufficiently many copies of qubits in a so-called *magic state* [6, 7]. This transformation only increases the circuit size by a polynomial factor. For classical circuit complexity theorists, a useful albeit imperfect analogy is the fact that any size-s circuit can be simulated by a *monotone* circuit of size polynomial in s, which is given as additional inputs the n negations of the input variables $x_1, \ldots, x_n$.

Two examples for such magic states are $|H\rangle = \cos(\pi/8)\,|0\rangle + \sin(\pi/8)\,|1\rangle$ and $|R\rangle = \cos(\beta)\,|0\rangle + e^{i\pi/4}\sin(\beta)\,|1\rangle$, where $\beta = \frac{1}{2}\arccos(1/\sqrt{3})$ [6].[1] This suggests the possibility of simulating a general quantum circuit by decomposing $|H^{\otimes n}\rangle$ or $|R^{\otimes n}\rangle$ as a linear combination of stabilizer states.

More formally, $|\varphi\rangle$ is a *stabilizer state* if $|\varphi\rangle = U\,|0^n\rangle$ where U is an n-bit Clifford unitary (see also Equation (1) and the following paragraph). The *stabilizer rank* of a state $|\psi\rangle$, denoted $\chi(\psi)$, is the minimal integer r such that

$$|\psi\rangle = \sum_{j=1}^{r} c_j\,|\varphi_j\rangle\,,$$

where for every $1 \le j \le r$, $|\varphi_j\rangle$ is a stabilizer state and $c_j \in \mathbb{C}$.

For any n-qubit state, the stabilizer rank is at most 2^n. Interestingly, much smaller upper bounds can be shown for the the stabilizer rank of $|H^{\otimes n}\rangle$: Bravyi, Smith and Smolin [7] proved that $\chi(H^{\otimes 6}) \le 7$ which implies that $\chi(H^{\otimes n}) \le 7^{n/6} \le 2^{0.468n}$. Bravyi, Smith and Smolin [7] then use this identity to obtain simulation algorithms for circuits with a small number of T gates, whose running time is much faster than the trivial brute-force simulation. A slightly faster algorithm was presented by Kocia who proved that

[1] The state $|R\rangle$ is often called $|T\rangle$. However, to avoid confusion, we follow the notation of [7], and reserve the notation $|T\rangle$ for a different state. For a handy reference to our notation, see Appendix A.

$\chi(H^{\otimes 12}) \leq 47$ [14], and upper bound was further improved by Qassim, Pashayan and Gosset [19] who proved that $\chi(H^{\otimes n}) = O(2^{\alpha n})$ for $\alpha = \frac{1}{4}\log_2 3 \leq 0.3963$.

When simulating quantum circuits, it is often enough, for all intents and purposes, to obtain an approximation of their output state. Thus, it's natural to define a similar approximation notion for stabilizer rank. The δ-approximate stabilizer rank of $|\psi\rangle$, denoted $\chi_\delta(\psi)$, is defined as the minimum of $\chi(\varphi)$ over all states $|\varphi\rangle$ such that $\||\psi\rangle - |\varphi\rangle\|_2 \leq \delta$ [4]. By considering approximate stabilizer decomposition of $|H^{\otimes n}\rangle$, improved simulation algorithms were obtained by Bravyi and Gosset [5].

A natural question is then what is the limit of such simulation methods. As the running time of the simulation scales with the stabilizer rank, an upper bound which is polynomial (in n) on $\chi(H^{\otimes n})$ or $\chi(R^{\otimes n})$ will imply that $\mathsf{BPP} = \mathsf{BQP}$ and even (by simulating quantum circuits with postselection) $\mathsf{P} = \mathsf{NP}$ [4], and thus seems highly improbable.[2] Much stronger hardness assumptions than $\mathsf{P} \neq \mathsf{NP}$, such as the exponential time hypothesis, imply that $\chi(H^{\otimes n}) = 2^{\Omega(n)}$ [17, 12].

However, the starting point of this discussion was our desire to obtain *unconditional* impossibility results, and thus we are interested in provable lower bounds on $\chi(H^{\otimes n})$ and $\chi_\delta(H^{\otimes n})$, and similarly for $R^{\otimes n}$.

While it's easy to see, using counting arguments, that the stabilizer rank of a random quantum state would be exponential, it is a challenging open problem to prove super-polynomial lower bounds on the rank of $|H^{\otimes n}\rangle$ or for other explicit states. Bravyi, Smith and Smolin proved that $\chi(H^{\otimes n}) = \Omega(\sqrt{n})$. In this paper, we improve this lower bound, and also prove the first non-trivial lower bounds for approximate stabilizer rank.

1.2 Our results: Improved Lower Bounds on Stabilizer Rank and Approximate Stabilizer Rank

Our first result is an improved lower bound on $\chi(H^{\otimes n})$ and $\chi(R^{\otimes n})$.

Theorem 1.1. *$\chi(H^{\otimes n}) = \Omega(n)$, and similarly, $\chi(R^{\otimes n}) = \Omega(n)$.*

As we remark in Section 1.3, proving super-linear lower bound on $\chi(H^{\otimes n})$ will solve a notable open problem in complexity theory. We discuss this challenge, as well us some barriers preventing our technique from proving super-linear lower bounds, in Section 1.5.

The result of Theorem 1.1 can be immediately adapted to prove the same lower bounds on the δ-approximate stabilizer rank for exponentially small δ. We are, however, interested in much coarser approximations, and we are able to prove a meaningful result even for δ being a small enough positive constant.

Theorem 1.2. *There exists an absolute constant $\delta > 0$ such that $\chi_\delta(H^{\otimes n}) = \Omega(\sqrt{n}/\log n)$, and similarly $\chi_\delta(R^{\otimes n}) = \Omega(\sqrt{n}/\log n)$.*

By definition, the stabilizer rank of any two states which are Clifford-equivalent is the same, and thus the lower bounds of Theorem 1.1 and Theorem 1.2, while stated as lower bounds on the ranks of $|H^{\otimes n}\rangle$ and $|R^{\otimes n}\rangle$ hold for any state which is Clifford-equivalent to them, even up to a phase.

[2] This implication holds up to uniformity issues having to do with finding the decomposition of $|H^{\otimes n}\rangle$ as a linear combination of stabilizer states. However, these complexity classes collapses are not believed to hold even in the non-uniform world, and further, by the Karp-Lipton theorem, a non-uniform collapse also implies a collapse of the polynomial hierarchy in the uniform world.

1.3 Technique: Stabilizer States as Quadratic Polynomials

The original proof of the Gottesman-Knill Theorem used the stabilizer formalism and tracked the current state of the circuit by storing the generators of the subgroup of the Pauli group which stabilizes the state, and updating them after each application of a Clifford operation. It turns out, however, that there is an alternative succinct representation of stabilizer states, using their amplitudes in the computational basis $\{|x\rangle\}_{x\in\mathbb{F}_2^n}$ [8, 26]. This representation also leads to an alternative proof of the theorem, as explained in [26].

If $|\varphi\rangle$ is a stabilizer state then (up to normalization)

$$|\varphi\rangle = \sum_{x\in A} i^{\ell(x)}(-1)^{q(x)}\,|x\rangle \tag{1}$$

where $A \subseteq \mathbb{F}_2^n$ is an affine subspace, $\ell(x)$ is an $\mathbb{F}_2$-linear function and $q(x)$ is a quadratic polynomial over $\mathbb{F}_2$. The amplitudes of $|H^{\otimes n}\rangle$ and $|R^{\otimes n}\rangle$ are also easy to compute. For example, recall that $|H\rangle = \cos(\pi/8)\,|0\rangle + \sin(\pi/8)\,|1\rangle$, and thus

$$|H^{\otimes n}\rangle = \sum_{x\in\mathbb{F}_2^n} \cos(\pi/8)^{n-|x|}\sin(\pi/8)^{|x|}\,|x\rangle\,,$$

where $|x|$ denotes the Hamming weight of x.

It is convenient to recast this problem as a problem about functions on the boolean cube in the following natural way. For an n-qubit state $|\psi\rangle$ we associate a function $F_\psi : \mathbb{F}_2^n \to \mathbb{C}$ such that $F_\psi(x)$ equals the amplitude of $|x\rangle$ when writing $|\psi\rangle$ in the computational basis. In this formulation, our "building blocks" are *stabilizer functions*, i.e., functions of the form

$$\varphi(x) = i^{\ell(x)}(-1)^{q(x)}\mathbb{1}_A$$

where A is an affine subspace, $\mathbb{1}_A$ is the indicator function of A (i.e., $\mathbb{1}_A(x) = 1$ if $x \in A$ and zero otherwise), ℓ is a linear function and q is a quadratic polynomial. Let H_n denote the function associated with $|H^{\otimes n}\rangle$. We would like to show that in any decomposition

$$H_n(x) = \sum_{j=1}^{r} c_j\varphi_j(x) = \sum_{j=1}^{r} c_j i^{\ell(x)}(-1)^{q_j(x)}\mathbb{1}_{A_j}(x)$$

where $c_j \in \mathbb{C}$ and $\varphi_j(x)$ are stabilizer functions, r must be large.

Our techniques for showing that use tools from the analysis of boolean functions and from complexity theory. In Section 1.4 we recall some similar questions that have arisen in complexity theory.

For the proof of Theorem 1.1, we show that if f is a function of stabilizer rank at most, say, $n/100$, then it is possible to find two vectors $x, y \in \mathbb{F}_2^n$ such that the Hamming weight of x is very small, the Hamming weight of y is very large, and $f(x) = f(x+y)$. Since $|x+y| \geq |y|-|x|$, for the correctly chosen parameters we get that $|x+y| > |x|$, which leads to a contradiction if $f = H_n$, since H_n takes different value on each layer of the Hamming cube.

To find such x and y, given a decomposition $\sum_{j=1}^{r} c_j i^{\ell(x)}(-1)^{q_j}\mathbb{1}_{A_j}$ with $r \leq n/100$, we find x, y such that $\ell_j(x) = \ell_j(x+y)$, $q_j(x) = q_j(x+y)$ and $\mathbb{1}_{A_j}(x) = \mathbb{1}_{A_j}(x+y)$ for all $j \in [r]$.

Observe that for a fixed $y \in \mathbb{F}_2^n$ and a quadratic polynomial $q(x)$, the equation $q(x) = q(x+y)$ is an affine linear equation in unknowns x. Thus, denoting $\Delta_y(q) = q(x)+q(x+y)$ (this is also called the directional derivative of q with respect to y), we get a system of affine

linear equations $\{\Delta_y(q_j) = 0\}_{j\in[r]}$ in x, which, assuming r is small, has many solutions (assuming it is solvable at all).

The additional requirements $\ell_j(x) = \ell_j(x+y)$ and $\mathbb{1}_{A_j}(x) = \mathbb{1}_{A_j}(x+y)$ make things more complicated. However, using an averaging argument and by again utilizing the fact that r is relatively small, we are able to find a large affine subspace U of vectors which satisfy those equations, and then we analyze the above system of linear equations over the affine subspace U.

In order to satisfy the conditions on the Hamming weights of x and y we use Kleitman's theorem [13] which gives an upper bound on the size of sets of the boolean cube with small diameter, as well as some elementary linear algebra. The full proof of Theorem 1.1 appears in Section 3.

The proof of Theorem 1.2 follows a different strategy. Starting from a state $|\psi\rangle$ of rank r which is δ-close to $|H^{\otimes n}\rangle$ for some small enough constant $\delta > 0$, we show how to use $|\psi\rangle$ in order to construct an $\mathbb{F}_2$-polynomial of degree $O(r \log r)$ which $(1-\varepsilon)$-approximates the majority function on $m = \Omega(n)$ bits. By a well known correlation bound of Razborov and Smolensky [21, 24, 25], this implies that $r = \Omega(\sqrt{n}/\log n)$.

We now explain how to obtain this polynomial approximating the majority function. Let $p = \sin^2(\pi/8) = 0.146\ldots$. Instead of majority, it is convenient to first consider the function THR_{pn} which is 1 on all inputs x whose Hamming weight is at least pn, and zero otherwise. Note that this function is trivial to approximate under the uniform distribution by the constant 1 polynomial, but the approximation question becomes meaningful when considering $B(n,p)$, the binomial distribution with parameter p on the n-dimensional cube. This is useful since the L_2 mass of the vector $|H^{\otimes n}\rangle$ is distributed according to this distribution. In particular it is heavily concentrated on coordinates x such that $|x| = pn \pm O(\sqrt{n})$, and a state $|\psi\rangle$ which is δ-close to $|H^{\otimes n}\rangle$ must contain in almost all of these coordinates values which are very close to those of $|H^{\otimes n}\rangle$. It is then possible to obtain from ψ a boolean function f which approximates the function THR_{pn}. We observe that a restriction g of f to a random set of $2pn$ coordinates will approximate the majority function, and further, assuming $|\psi\rangle$ has stabilizer rank r, and using standard techniques again borrowed from Razborov and Smolensky, g itself can be approximated by a polynomial $\tilde{g}$ of degree $O(r \log r)$. It follows that $\tilde{g}$ approximates the majority function over $2pn$ bits. The full proof of Theorem 1.2 appears in Section 4.

1.4 Related Work

As mentioned above, the previous best lower bound was an $\Omega(\sqrt{n})$ lower bound for exact stabilizer rank of $|H^{\otimes n}\rangle$ proved by Bravyi, Smith and Smolin [7]. Stronger lower bounds are known in restricted models. As mentioned by [7] (see also Lemma 2 in [5]), for every stabilizer state $|\varphi\rangle$ it holds that $|\langle\varphi|H^{\otimes n}\rangle| \leq 2^{-\Omega(n)}$ which immediately implies an exponential lower bound in the case that the coefficients c_j are bounded in magnitude (in particular, this holds if the states in the decomposition are orthogonal). It is worth noting that by Cramer's rule, in any rank r decomposition the coefficients c_j can be taken to be of magnitude at most exponential in n and r.

Bravyi et al. [4] present a different restricted model in which they prove an exponential lower bound.

Related questions have been considered before in complexity theory. The so called "quadratic uncertainty principle" [9, 27] is a conjecture which states that in any decompo-

sition of the AND function as a sum

$$\sum_{j=1}^{r} c_j (-1)^{q_j(x)}, \tag{2}$$

for quadratic functions $\{q_j\}_{j\in[r]}$ and $c_j \in \mathbb{C}$, $r = 2^{\Omega(n)}$. The best lower bound known is $r \geq n/2$ (see [27]). Note that since in the stabilizer rank case we allow functions of the form $(-1)^q \cdot \mathbb{1}_A$ for affine subspaces A, the model we consider in this paper is stronger: in particular the AND function itself is a stabilizer function and its stabilizer rank is 1.

Williams [27] has constructed, for every positive integer k, a function $f_k \in \mathsf{NP}$ which requires $r = \Omega(n^k)$ in any decomposition as in (2). It remains, however, an intriguing open problem to construct boolean function in P which requires a super-linear number of summands.

We remark that proving super linear lower bounds on the stabilizer rank of $|H^{\otimes n}\rangle$ will solve this problem. Indeed, as mentioned above, the stabilizer rank model is even stronger, and thus lower bounds carry over to weaker models. Furthermore, even though H_n itself is not a boolean function, $|H\rangle$ is Clifford-equivalent (up to an unimportant phase) to $|T\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle)$ (see [7]), which implies that the stabilizer rank of $|H^{\otimes n}\rangle$ equals the stabilizer rank of $|T^{\otimes n}\rangle$. Denoting T_n the function associated with $T^{\otimes n}$, it is now evident that $T_n(x)$ depends only on $|x|$ mod 8, and therefore

$$T_n = \sum_{j=0}^{7} b_j M_j(x),$$

where for $j \in \{0, ..., 7\}$, $b_j \in \mathbb{C}$ and $M_j : \mathbb{F}_2^n \to \{0, 1\}$ is a boolean function such that $M_j(x) = 1$ if and only if $|x| = j$ mod 8. Thus, a super-linear lower bound on the stabilizer rank of $|H^{\otimes n}\rangle$ will imply a super-linear lower bound on the rank of the (boolean) mod 8 function.

Following the initial publication of this work, our results were reproved using different techniques. Labib [15] used higher-order Fourier analysis in order to prove a result similar to Theorem 1.1, and extended it to qudits of any prime dimension. Lovitz and Steffan [16] proved nearly identical lower bounds for exact and approximate stabilizer rank using number-theoretic techniques.

1.5 Open Problems

While Theorem 1.1 improves upon the previous best lower bound known, we are unfortunately unable to prove super-polynomial or even super-linear lower bounds on $\chi(H^{\otimes n})$ or $\chi(R^{\otimes n})$. Further, our techniques seem incapable of proving super-linear lower bounds, as they extend to any representation of H_n as an arbitrary function of r stabilizer functions, and not necessarily a linear combination of them.

As mentioned in Section 1.4, it seems that a first step could be proving super-linear lower bounds for the quadratic uncertainty principle problem. A different approachable open problem is to improve our lower bound on the δ-approximate stabilizer rank to be closer to $\Omega(n)$. This could perhaps be easier assuming δ is polynomially small in n.

Acknowledgements The third author would like to thank Andru Gheorghiu for introducing him to the notion of stabilizer rank.

2 Preliminaries

2.1 General Notation

As mentioned in the introduction, it is often convenient to speak about functions on the boolean cube rather than quantum states. For an n-qubit state $|\psi\rangle = \sum_{x\in\mathbb{F}_2^n} c_x |x\rangle$, the associated function $F_\psi : \mathbb{F}_2^n \to \mathbb{C}$ is defined as $F_\psi(x) = c_x$.

The L_2 norm of the function $F : \mathbb{F}_2^n \to \mathbb{C}$ is then the same as the norm of the corresponding vector, i.e., $\|F\| = \left(\sum_{x\in\mathbb{F}_2^n} |F(x)|^2\right)^{1/2}$.

A function $\varphi : \mathbb{F}_2^n \to \mathbb{C}$ is called a *stabilizer function* if there exists an n-variate linear function $\ell(x)$, an n-variate quadratic polynomial $q(x) \in \mathbb{F}_2[x_1, \ldots, x_n]$ and an affine subspace $A \subseteq \mathbb{F}_2^n$ such that $\varphi(x) = i^{\ell(x)}(-1)^{q(x)}\mathbb{1}_A$, where $\mathbb{1}_A$ denotes the characteristic function of A. As shown in [8, 26], stabilizer functions indeed correspond to stabilizer states up to normalization (which has no effect on the stabilizer rank).

The *stabilizer rank* of a function $F : \mathbb{F}_2^n \to \mathbb{C}$, denoted $\chi(F)$, is the minimal r such that there exist $c_1, \ldots, c_r \in \mathbb{C}$ and stabilizer functions $\varphi_1, \ldots, \varphi_r$ such that $F(x) = \sum_{j=1}^r c_j\varphi_j(x)$.

For a vector $x \in \mathbb{F}_2^n$ we denote by $|x|$ its Hamming weight. We denote by $\mathsf{Maj}_m : \mathbb{F}_2^m \to \mathbb{F}_2$ the m-bit majority function, that is $\mathsf{Maj}_m(x) = 1$ if and only if $|x| \geq m/2$.

Definition 2.1. *Let $A \subset \mathbb{F}_2^n$. The* diameter *of A, denoted $\mathsf{diam}(A)$, is defined as*

$$\max_{u,v\in A} d(u, v) = \max_{u,v\in A} |u + v|.$$

Here $d(u, v)$ denotes the Hamming distance of u and v.

Kleitman [13] proved that sets of small diameter cannot be too large.

Theorem 2.2 ([13])**.** *Let $A \subset \mathbb{F}_2^n$ such that $\mathsf{diam}(A) \leq 2k$ for $k < n/2$. Then,*

$$|A| \leq \sum_{j=0}^{k} \binom{n}{j} \leq 2^{H_2(\frac{k}{n})n},$$

where $H_2(p) = -p\log_2 p - (1 - p)\log_2(1 - p)$ is the binary entropy function.

This result is obviously tight as shown by the example of the set of all vectors of Hamming weight at most k.

2.2 Linear Algebraic Facts

Recall that an affine subspace $U \in \mathbb{F}_2^n$ is a the set of solutions to a system of affine equations, i.e., a system of the form $Mx = b$ for some $M \in \mathbb{F}_2^{k\times n}$ and $b \in \mathbb{F}_2^k$. Every affine subspace can be written as $U = u + U_0$ for $u \in \mathbb{F}_2^n$ and a linear subspace $U_0 \subseteq \mathbb{F}_2^n$. In our terminology, linear subspaces are in particular affine subspaces (and similarly, linear functions are a special case of affine functions).

We record the following useful facts.

Fact 2.3. *Let $U \subsetneq \mathbb{F}_2^n$ be an affine subspace, and let $v \in \mathbb{F}_2^n \setminus U$. Then there is an affine function $a(x) : \mathbb{F}_2^n \to \mathbb{F}_2$ such that $a(v) = 1$ and for every $u \in U$, $a(u) = 0$.*

Fact 2.4. *Let $U_1, U_2 \subseteq \mathbb{F}_2^n$ be affine subspaces such that $U_1 \cap U_2 \neq \emptyset$. Then*

$$\dim(U_1 + U_2) = \dim(U_1) + \dim(U_2) - \dim(U_1 \cap U_2).$$

Claim 2.5. *Let $U \subseteq \mathbb{F}_2^n$ be an affine subspace, with $\dim(U) = n - k > 0$. There exists a subset $S \subset [n]$, of size $|S| = n - k$ such that for every $v \in \mathbb{F}_2^{n-k}$ there is $u \in U$ with $u|_S = v$ (where $u|_S$ denotes the restriction of u to the coordinates indexed by S).*

Proof. U is the set of solutions for an equation $Mx = b$ for a matrix $M \in \mathbb{F}_2^{k \times n}$ and $b \in \mathbb{F}_2^k$. The fact that $\dim(U) = n-k$ implies that M has rank k, and there is a $k \times k$ non-singular submatrix M' of M. Denote by S the columns of M that do not appear in M'. For every $v \in \mathbb{F}_2^{n-k}$, fixing $x|_S = v$ in the equation $Mx = b$ gives a system of equations $M'x' = b'$ in the set of remaining k unknowns x', which has a solution since M' is non-singular. □

Corollary 2.6. *Let $U \subseteq \mathbb{F}_2^n$ be an affine subspace with $\dim(U) = n - k > 0$. Then, there exists $u \in U$ with $|u| \leq k$.*

Proof. Follows immediately from applying Claim 2.5 with $v = 0$. □

Finally, we define the directional derivative of a quadratic function over $\mathbb{F}_2$.

Definition 2.7. *Let $q \in \mathbb{F}_2[x_1, \ldots, x_n]$ be a polynomial of degree 2. Let $0 \neq y \in \mathbb{F}_2^n$. The* directional derivative *of q in direction y is defined to be the function*

$$\Delta_y(q)(x) := q(x) + q(x + y) \in \mathbb{F}_2[x_1, \ldots, x_n].$$

Observe that for every y, $\Delta_y(q)$ is an affine function in x.

3 A Lower Bound for Exact Stabilizer Rank

In this section we prove Theorem 1.1. We first present the main lemma of this section.

Lemma 3.1. *Let $F : \mathbb{F}_2^n \to \mathbb{C}$ be a function of stabilizer rank r such that $r \leq n/100$. Then, there exist $y, z \in \mathbb{F}_2^n$ such that $|y| \neq |z|$ and $F(y) = F(z)$.*

Theorem 1.1, which we now restate, is an immediate corollary of Lemma 3.1.

Theorem 3.2. *Let $|B\rangle$ be either $|H\rangle$ or $|R\rangle$. Then $\chi(B^{\otimes n}) = \Omega(n)$.*

Proof. In the case where $|B\rangle = |H\rangle$, the associated function $F_H : \mathbb{F}_2^n \to \mathbb{C}$ is defined by $F_H(x) = \cos(\pi/8)^{n-|x|} \sin(\pi/8)^{|x|}$. If $|B\rangle = |R\rangle$, the associated function $F_R : \mathbb{F}_2^n \to \mathbb{C}$ is defined by $F_R(x) = \cos(\beta)^{n-|x|}(e^{i\pi/4} \sin(\beta))^{|x|}$ where $\beta = \arccos(1/\sqrt{3})/2$.

It is immediate to verify that for every $y, z \in \mathbb{F}_2^n$ of different Hamming weight those functions attain different values. Thus, by Lemma 3.1, their stabilizer rank is at least $n/100$. □

We turn to the proof of Lemma 3.1.

Proof of Lemma 3.1. Let $F : \mathbb{F}_2^n \to \mathbb{C}$ be a function of stabilizer rank at most $r \leq n/100$, i.e.,

$$F(x) = \sum_{j=1}^{r} c_j i^{\ell_j(x)} (-1)^{q_j(x)} \mathbb{1}_{A_j}(x),$$

where for every $j \in [r]$, ℓ_j is a linear function, q_j is a quadratic function, and $A_j \subseteq \mathbb{F}_2^n$ is an affine subspace.

To prove the statement of the lemma, we will show that there exist $y, z \in \mathbb{F}_2^n$ such that $|y| < |z|$ and for every $j \in [r]$ all of the following hold:

1. $\ell_j(y) = \ell_j(z)$
2. $\mathbb{1}_{A_j}(y) = \mathbb{1}_{A_j}(z)$
3. $q_j(y) = q_j(z)$.

The first two items are handled by the following claim, which shows that there is a large affine subspace satisfying both conditions.

Claim 3.3. *There's an affine subspace $U \subseteq \mathbb{F}_2^n$ of dimension at least $n - 3r$ such that for every $j \in [r]$ and for every $u_1, u_2 \in U$, $\ell_j(u_1) = \ell_j(u_2)$ and $\mathbb{1}_{A_j}(u_1) = \mathbb{1}_{A_j}(u_2)$.*

We defer the proof of Claim 3.3 to the end of this proof. Write $U = u + U_0$ where $u \in \mathbb{F}_2^n$ and $U_0 \subseteq \mathbb{F}_2^n$ is a *linear* subspace. The next claim handles the third item above.

Claim 3.4. *There exists $v \in U_0$ with $|v| \geq 2n/3$ such that the system of equations*

$$\{q_j(x) = q_j(x+v)\}_{j\in[r]}$$

(in unknowns x) has a solution in U.

We postpone the proof of this claim as well, and now explain how it implies the result. Let $v \in U_0$ as promised in Claim 3.4. The set of solutions in U to the system of affine equations

$$\{q_j(x) = q_j(x+v)\}_{j\in[r]} = \{\Delta_v(q)(x) = 0\}_{j\in[r]} \tag{3}$$

is non-empty (by Claim 3.4), and thus by Fact 2.4, the set of solutions in U to (3) is an affine subspace $V \subseteq U$ of dimension at least $n - 4r$.

By Corollary 2.6, there is $y \in V$ with $|y| \leq 4r$. Set $z = y+v$, so that $q_j(y) = q_j(y+v) = q_j(z)$ for all $j \in [r]$. Observe that $z \in U$, since Claim 3.4 promises that $v \in U_0$. Thus y and z attain the same values on ℓ_j and $\mathbb{1}_{A_j}$ for all $j \in [r]$ as well. Finally note that $|y| \leq 4r$ whereas

$$|z| = |y + v| \geq |v| - |y| \geq \frac{2n}{3} - 4r > 4r. \qquad \square$$

It remains to prove Claim 3.3 and Claim 3.4.

Proof of Claim 3.3. Let $V_1 \subset \mathbb{F}_2^n$ be the linear subspace defined by the system of equations $\{\ell_j = 0\}$ for all $j \in [r]$. It holds that $\dim(V_1) \geq n - r > 0$.

Consider now the map $E : V_1 \to \{0,1\}^r$, defined by

$$E(x) = (\mathbb{1}_{A_1}(x), \ldots, \mathbb{1}_{A_r}(x)).$$

By the pigeonhole principle, there is $\alpha \in \{0,1\}^r$ with $|E^{-1}(\alpha)| \geq 2^{\dim V_1 - r} \geq 2^{n-2r}$. Let S be the support of α, that is, the set of indices $j \in [r]$ such that $\alpha_j = 1$. We have that

$$E^{-1}(\alpha) = \left(\left(\bigcap_{j\in S} A_j\right) \setminus \left(\bigcup_{j\notin S} A_j\right)\right) \cap V_1 \subseteq \left(\bigcap_{j\in S} A_j\right) \cap V_1$$

(for notational convenience, if $S = \emptyset$, then $\bigcap_{j\in S} A_j = \mathbb{F}_2^n$).

Let $V_2 = \left(\bigcap_{j\in S} A_j\right) \cap V_1$. Then V_2 is an affine subspace, and $|V_2| \geq |E^{-1}(\alpha)| \geq 2^{n-2r}$, so $\dim(V_2) \geq n - 2r > 0$.

Pick now an arbitrary $x_0 \in E^{-1}(\alpha)$. Thus, $x_0 \in V_2$, and for every $j \notin S$, $x_0 \notin A_j$. By Fact 2.3, for every $j \notin S$ there is an affine equation a_j such that $a_j(x_0) = 1$ and for all $x \in A_j$, $a_j(x) = 0$. Let

$$U = \{x \in V_2 : \text{for all } j \notin S, a_j(x) = 1\}.$$

Then U is an affine subspace (as it is defined by at most r additional affine constraints on V_2), and it is non-empty (since $x_0 \in U$). By Fact 2.4, it follows that $\dim(U) \geq n - 2r - r = n - 3r$. Further, for every $x \in U$ and $j \in [r]$, it holds that $\ell_j(x) = 0$ and

$$\mathbb{1}_{A_j}(x) = \begin{cases} 1 & j \in S \\ 0 & j \notin S \end{cases}$$

which completes the proof. □

We finish the section by proving Claim 3.4.

Proof of Claim 3.4. Consider the map $\Gamma : U \to \{0,1\}^r$ defined by

$$\Gamma(x) = (q_1(x), \ldots, q_r(x)).$$

For every $\alpha \in \{0,1\}^r$, let $\Gamma_\alpha = \{x_1 + x_2 : x_1, x_2 \in \Gamma^{-1}(\alpha)\}$. Observe that for every α, $\Gamma_\alpha \subseteq U_0$. Furthermore, for every $v \in \Gamma_\alpha$, the set of affine equations

$$\{\Delta_v(q_j)(x) = 0\}_{j \in [r]},$$

in unknowns x, has a solution in U. Indeed, $v = x_1 + x_2$ where $x_1, x_2 \in \Gamma^{-1}(\alpha)$, and thus $q_j(x_1) = q_j(x_2) = q_j(x_1 + v)$ for every $j \in [r]$, which implies that x_1 is a solution.

In order to finish the proof we need to show that there is $\alpha \in \{0,1\}^r$ and $v \in \Gamma_\alpha$ such that $|v| \geq \frac{2n}{3}$. By the pigeonhole principle there is $\alpha_0 \in \{0,1\}^r$ such that $|\Gamma^{-1}(\alpha_0)| \geq |U|/2^r = 2^{n-4r}$. Observe that the maximal Hamming weight of an element in Γ_{α_0} equals the diameter of the set $\Gamma^{-1}(\alpha_0)$.

By Theorem 2.2 (for $k = n/3$), the size of every set of diameter $2n/3$ is at most $2^{H_2(1/3)n} \leq 2^{0.92n}$. Since $r \leq n/100$, $|\Gamma^{-1}(\alpha_0)| > 2^{0.95n}$, so $\mathsf{diam}(\Gamma^{-1}(\alpha_0)) \geq 2n/3$, and there is $v \in \Gamma_{\alpha_0}$ of weight at least $2n/3$. □

4 A Lower Bound for Approximate Stabilizer Rank

In this section we prove Theorem 1.2. In Section 4.1, we show how to obtain, given a function f that approximates the function THR_{pn} (with respect to the binomial distribution on the n-dimensional cube with parameter p, $B(n,p)$), a random restriction of f which approximates the majority function over $m = 2\lfloor pn \rfloor$, bits with respect to the uniform distribution.[3] In Section 4.2, we construct, given a state $|\psi\rangle$ that is δ close to either $|H^{\otimes n}\rangle$ or $|R^{\otimes n}\rangle$, a boolean function f_ψ that approximates THR_{pn}. In Section 4.3 we then show how to get low-degree polynomial approximations to restrictions of f_ψ, which, as we specify in Section 4.4, completes the proof.

[3] In what follows, in order to help with the readability of the argument we often omit the floor and ceiling signs. For example, we'll use "pn" to refer to $\lfloor pn \rfloor$. We reintroduce floor and ceiling signs in cases where there is a chance of confusion.

4.1 A Reduction from Threshold Functions to Majority

Let $0 < p < 1/2$. Recall that $\mathsf{THR}_{pn}(x)$ equals 1 if $|x| \geq pn$ and 0 otherwise. In this section we prove that given any function $f : \mathbb{F}_2^n \to \mathbb{F}_2$ that approximates THR_{pn} with respect to $B(n,p)$, we can find a function g, which is a restriction of f to $2pn$ random coordinates, which approximates the majority function on those bits with respect to the uniform distribution.

In anticipation of the next section, when considering approximations for THR_{pn} we will work with a slightly different notion of approximation than approximation with respect to $B(n,p)$, which we now explain.

Let $L_k = \{x \in \mathbb{F}_2^n \mid |x| = k\}$ denote the k-th layer of the boolean cube. We say that a function $f : \mathbb{F}_2^n \to \mathbb{F}_2$ is ε-*wrong* on L_k (with respect to THR_{pn}) if the fraction of elements $x \in L_k$ such that $f(x) \neq \mathsf{THR}_{pn}(x)$ is at least ε.

We say that f (ε,γ)-*approximates* THR_{pn} if f is ε-wrong on at most a γ fraction of the layers L_k for $k \in [pn - \lceil 5\sqrt{2pn}\rceil, pn + \lceil 5\sqrt{2pn}\rceil]$.

For the rest of the proof we will always set $\varepsilon = \gamma = 0.01$.

Since $B(n,p)$ is heavily concentrated on layers L_k with $k \in [pn - O(\sqrt{n}), pn + O(\sqrt{n})]$, and for every k in that range, $\Pr_{x\sim B(n,p)}[x \in L_k] = \Theta(1/\sqrt{n})$, this notion and the notion of approximation with respect to $B(n,p)$ are in fact very similar, up to the precise choice of constants.

Lemma 4.1. *Let $0 < p < 1/2$ be an absolute constant, and let $f : \mathbb{F}_2^n \to \mathbb{F}_2$ be a boolean function that $(0.01, 0.01)$-approximates THR_{pn}. For every $D \subseteq [n]$ of size $m := 2pn$, let $g_D : \mathbb{F}_2^m \to \mathbb{F}_2$ be the function obtained from f by fixing all input bits outside of D to 0. Then there exists D_0 such that for $g := g_{D_0}$, $\Pr_{x\in\mathbb{F}_2^m}[g(x) = \mathsf{Maj}_m(x)] \geq 3/4$, where x is chosen according to the uniform distribution.*

Proof. Let $m = 2pn$. For every $D \subseteq [n]$ of size m, let g_D be the function obtained from f by fixing all input bits outside of D to 0. It will be convenient to consider g_D as a function whose domain is $\mathbb{F}_2^m$ using some bijection between D and $[m]$. Every $x \in \mathbb{F}_2^n$ which is zero on coordinates outside of D then corresponds to a unique $\bar{x} \in \mathbb{F}_2^m$, and vice versa.

We will now pick D uniformly at random among all subsets of $[n]$ of size m, so that g_D is a random restriction of f.

We say $x \in \mathbb{F}_2^n$ *survives* D if the set of indices $j \in [n]$ such that $x_j = 1$ is contained in D. The probability that $x \in L_k$ survives D is $\binom{m}{k}/\binom{n}{k}$.

For an input $x \in \mathbb{F}_2^n$, we say x is *correct* if $f(x) = \mathsf{THR}_{pn}(x)$, and incorrect otherwise. If x is correct and survives, then $\mathsf{Maj}_m(\bar{x}) = \mathsf{THR}_{pn}(x) = f(x) = g_D(\bar{x})$.

Let X_k be a random variable, which denotes the number of incorrect inputs $x \in L_k$ that survive D, and

$$X = \sum_{k=pn-\lceil 5\sqrt{2pn}\rceil}^{pn+\lceil 5\sqrt{2pn}\rceil} X_k.$$

By the assumption, for at least 0.99 fraction of the layers L_k, the number of incorrect x's is at most $0.01\binom{n}{k}$, and thus for each such layer L_k for $k \in [pn - \lceil 5\sqrt{2pn}\rceil, pn + \lceil 5\sqrt{2pn}\rceil]$, $\mathbb{E}_D[X_k] \leq 0.01\binom{m}{k}$. We call such layers *good*. For the rest of the layers, which we call *bad*, obviously $\mathbb{E}_D[X_k] \leq \binom{m}{k}$. The total number of layers in the interval $[pn - \lceil 5\sqrt{2pn}\rceil, pn + \lceil 5\sqrt{2pn}\rceil]$ is at most

$$2 \cdot (5\sqrt{2pn} + 1) + 1 \leq 11\sqrt{2pn},$$

and thus the number of bad layers is at most $0.01 \cdot 11 \cdot \sqrt{2pn}$. Further, for every k, $\binom{m}{k} \leq \frac{1}{\sqrt{m}} \cdot 2^m$.

Therefore,

$$\begin{aligned}\mathbb{E}_D[X] &= \sum_{L_k \text{ is good}} \mathbb{E}_D[X_k] + \sum_{L_k \text{ is bad}} \mathbb{E}_D[X_k] \\ &\leq \sum_{L_k \text{ is good}} 0.01\binom{m}{k} + 0.01 \cdot 11\sqrt{m} \cdot \frac{1}{\sqrt{m}} \cdot 2^m \\ &\leq 0.01 \cdot \left(\sum_k \binom{m}{k}\right) + \frac{11}{100} \cdot 2^m \leq \frac{12}{100} \cdot 2^m.\end{aligned}$$

In particular, there is some D_0 such that the number if incorrect x's in layers $[pn - \lceil 5\sqrt{2pn}\rceil, pn + \lceil 5\sqrt{2pn}\rceil]$ that survive D_0 is at most $\frac{12}{100}2^m$. Let $g := g_{D_0}$. We now claim that g and Maj_m agree on more than $3/4$ of the inputs in $\mathbb{F}_2^m$.

First, By the Chernoff bound, the number of vectors $\bar{x} \in \mathbb{F}_2^m$ whose Hamming weight is *not* in the range

$$[pn - \lceil 5\sqrt{2pn}\lceil, pn + \lceil 5\sqrt{2pn}\rceil] = [m/2 - \lceil 5\sqrt{m}\rceil, m/2 + \lceil 5\sqrt{m}\rceil],$$

is at most $2e^{-((5/\sqrt{pn})^2 \cdot pn/6)} \cdot 2^m \leq \frac{1}{15}2^m$. On these inputs we have no guarantee. By the choice of D_0, the number of $\bar{x}$'s such that $|\bar{x}| \in [m/2 - \lceil 5\sqrt{m}\rceil, m/2 + \lceil 5\sqrt{m}\rceil]$ and $g(x) \neq \mathsf{Maj}(x)$ is at most $\frac{12}{100}2^m$. It follows that $g(\bar{x}) \neq \mathsf{Maj}_m(\bar{x})$ on less than $\frac{1}{4} \cdot 2^m$ inputs. □

4.2 From Stabilizer Decompositions to Threshold Functions

Let $|B\rangle = \alpha|0\rangle + \beta|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$. Let $p = |\beta|^2$ and suppose that $0 < p < 1/2$. Let $F_B : \mathbb{F}_2^n \to \mathbb{C}$ be the function associated with $|B^{\otimes n}\rangle$, i.e., $F_B(x) = \alpha^{n-|x|}\beta^{|x|}$.

In this section we prove that if $\psi : \mathbb{F}_2^n \to \mathbb{C}$ is such that $\chi(\psi) \leq r$ and $\|\psi - F_B\| \leq \delta$, then it is possible to construct a boolean function f_ψ that $(0.01, 0.01)$-approximates THR_{pn}. In Section 4.3, we will prove that if $\chi(\psi) \leq r$, f_ψ has low degree polynomial approximations.

From here on, δ will denote a sufficiently small constant, which may depend on $|B\rangle$ and its parameters (i.e., δ is some function of p), but does *not* depend on n. Since we are interested in the case $|B\rangle = |H\rangle$ or $|B\rangle = |R\rangle$, δ can be taken to be some small universal constant.

For $k \in [n]$, let $m_k := |\alpha^{n-k}\beta^k|$ denote the absolute value of F_B on the k-th layer. Let $w_k = m_k^2 = p^k(1-p)^{n-k}$ and $W_k = \binom{n}{k}w_k$ the total mass on the k-th layer, with respect to $B(n,p)$. Let $\eta = \frac{|\beta|}{|\alpha|}$. Observe that by assumption, $0 < \eta < 1$.

Suppose $\psi : \mathbb{F}_2^n \to \mathbb{C}$ is such that $\chi(\psi) \leq r$ and $\|\psi - F_B\| \leq \delta$. We define a boolean function $f_\psi : \mathbb{F}_2^n \to \mathbb{F}_2$ as follows:[4]

$$f_\psi(x) = \begin{cases} 1 & \text{if } |\psi(x)| \leq \left(\frac{1+\eta}{2}\right) m_{pn-1} \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

The intuition for the definition is that, since $\|\psi - F_B\| \leq \delta$, we expect $\psi(x)$ to be very close to $F_B(x)$ for most inputs x. For every such x, f_ψ will correctly compute THR_{pn}. Further, inputs x such that $f_\psi(x) \neq \mathsf{THR}_{pn}(x)$ correspond to inputs x such that $|\psi(x) - F_B(x)|$ is large. Assuming there are many such x's will lead to a contradiction to the assumption that $\|\psi - F_B\| \leq \delta$.

[4] Observe that if $|x| < |x'|$ then $|\psi(x)| > |\psi(x')|$.

Lemma 4.2. *Let $\psi : \mathbb{F}_2^n \to \mathbb{C}$ be a function such that $\|\psi - F_B\| \leq \delta$ for a sufficiently small δ. Let f_ψ the boolean function defined as in (4). Then f_ψ $(0.01, 0.01)$-approximates THR_{pn}.*

We begin with the following calculation.

Claim 4.3. *Suppose $x \in \mathbb{F}_2^n$ is such that $f_\psi(x) \neq \mathsf{THR}_{pn}(x)$ and $|x| = k$. Then $|\psi(x) - F_B(x)|^2 \geq w_k \cdot \left(\frac{1-\eta}{2}\right)^2$.*

Proof. Since $|x| = k$, $|F_B(x)| = m_k$. Suppose first that $k \leq pn - 1$ so that $\mathsf{THR}_{pn}(x) = 0$. By assumption, $f_\psi(x) = 1$, which implies that

$$|\psi(x)| \leq \left(\frac{1+\eta}{2}\right) m_{pn-1}.$$

Observe that $m_k = (\eta^{-1})^{pn-1-k} m_{pn-1} \geq m_{pn-1}$ for $k \leq pn - 1$, and therefore by the triangle inequality

$$\begin{aligned} |\psi(x) - F_B(x)| \geq |F_B(x)| - |\psi(x)| &\geq m_k - \left(\frac{1+\eta}{2}\right) m_{pn-1} \\ &\geq m_k - \left(\frac{1+\eta}{2}\right) m_k = \left(\frac{1-\eta}{2}\right) m_k, \end{aligned}$$

which implies the statement of the claim (for $k \leq pn - 1$) by squaring both sides.

If $k \geq pn$, then $\mathsf{THR}_{pn}(x) = 1$ which implies $f_\psi(x) = 0$, i.e.,

$$|\psi(x)| \geq \left(\frac{1+\eta}{2}\right) m_{pn-1}.$$

Note that $m_k = \eta^{k-pn+1} m_{pn-1}$ and in particular $m_k \leq \eta m_{pn-1}$ for all $k \geq pn$. Thus,

$$\begin{aligned} |\psi(x) - F_B(x)| \geq |\psi(x)| - |F_B(x)| &\geq \left(\frac{1+\eta}{2}\right) m_{pn-1} - m_k \\ &\geq \left(\frac{1+\eta}{2}\right) m_{pn-1} - \eta m_{pn-1} = \left(\frac{1-\eta}{2}\right) m_{pn-1} \\ &\geq \left(\frac{1-\eta}{2}\right) m_k, \end{aligned}$$

which proves the lemma for this case as well. □

We use the following standard estimates on the concentration of the binomial distribution. Recall that $W_k = \binom{n}{k} p^k (1-p)^{n-k}$.

Claim 4.4. *Let $C \in \mathbb{R}$. Then $W_{pn+C\sqrt{n}} = \Omega(1/\sqrt{n})$, where the constant hidden under the Ω notation depends on C and p, but not on n.*

Observe that C in the above claim may be negative. The proof is a direct application of Stirling's approximation. For completeness, we provide a crude estimate which suffices for us in Appendix B.

We are now ready to prove the main lemma of the section.

Proof of Lemma 4.2. Let k be a layer such that f_ψ is 0.01-wrong on L_k. By Claim 4.3,

$$\sum_{x \in L_k} |\psi(x) - F_B(x)|^2 \geq 0.01 \cdot \binom{n}{k} \cdot \left(\frac{1-\eta}{2}\right)^2 w_k = 0.01 \left(\frac{1-\eta}{2}\right)^2 W_k.$$

Suppose, towards a contradiction, f_ψ is 0.01-wrong on more than 0.01 fraction of the layers $k \in [pn - \lceil 5\sqrt{2pn} \rceil, pn + \lceil 5\sqrt{2pn} \rceil]$, i.e., on more than $0.1\sqrt{2pn}$ layers. By Claim 4.4, for every such k, $W_k \geq c/\sqrt{n}$ for some constant c which does not depend on n. It follows that

$$\|\psi - F_B\| \geq 0.1\sqrt{2pn} \cdot 0.01 \left(\frac{1-\eta}{2}\right)^2 \cdot \frac{c}{\sqrt{n}},$$

which is a contradiction for $\delta < 0.001\sqrt{2p}c \left(\frac{1-\eta}{2}\right)^2$. □

4.3 A Low Degree Polynomial Approximation

In this section we show that for the function f_ψ defined as in (4), and for any restriction g_D of f_ψ as in Lemma 4.1, the function g_D has a polynomial approximating it, whose degree is at most $O(r \log r)$. To prove this we apply standard approximation techniques used for proving lower bounds for bounded depth circuits with modular gates, although in our case the details are somewhat simpler.

We begin with the following lemma that shows how to approximate indicator functions of affine subspaces with low degree polynomials.

Claim 4.5 ([21, 25])**.** *Let $A \subseteq \mathbb{F}_2^m$ be an affine subspace. For every $t \in \mathbb{N}$, there exists a polynomial $P \in \mathbb{F}_2[x_1, \ldots, x_m]$ of degree at most t such that $\Pr_{x \in \mathbb{F}_2^m}[P(x) \neq \mathbb{1}_A(x)] \leq 2^{-t}$.*

Proof. Since A is an affine subspace, there exist $k \leq m$ affine functions $a_1, \ldots, a_k$ such that $x \in A$ if and only if $a_j(x) = 0$ for every $j \in [k]$, or equivalently, $\mathbb{1}_A(x) = \prod_{j=1}^k (a_j(x) + 1)$.

Let D be a uniformly random subset of $[k]$ and $a_D = \sum_{j \in D} a_j$. Observe that for $x \in A$, $a_D(x) = 0$ with probability 1, whereas for $x \notin A$, there is some $j \in [k]$ such that $a_j(x) = 1$ and hence $\Pr_D[a_D(x) = 0] = 1/2$ (as j is included in D with probability 1/2).

Hence, for $t \in \mathbb{N}$, define $P_{\mathbf{D}}(x) = \prod_{k=1}^t (a_{D_k}(x)+1)$, where $\mathbf{D} = (D_1, \ldots, D_t)$ are chosen uniformly and independently. Then, $P_{\mathbf{D}}$ is a degree t (random) polynomial, $P_{\mathbf{D}}(x) = 1$ for all $x \in A$, and for $x \notin A$, $\Pr_{\mathbf{D}}[P_{\mathbf{D}}(x) = 1] \leq 2^{-t}$. In particular, in expectation $P_{\mathbf{D}}$ and $\mathbb{1}_A$ disagree on at most 2^{m-t} of the inputs, which implies that there exists a choice of $\mathbf{D}' = (D_1, ..., D_t)$ such that $P := P_{\mathbf{D}'}$ satisfies the properties required in the lemma. □

We now show how to approximate restrictions of the boolean function f_ψ.

Lemma 4.6. *Let F_B and ψ be functions as in Section 4.2 and let f_ψ defined as in (4). Let $D \subseteq [n]$ and denote $g := g_D$ the restriction of f_ψ obtained by setting variables outside of D to 0, as in Section 4.1. Then, there is a polynomial $\tilde{g}$ of degree $O(r \log r)$ such that $\Pr_{\bar{x} \in \mathbb{F}_2^m}[g(\bar{x}) \neq \tilde{g}(\bar{x})] \leq \frac{1}{20}$.*

Proof. Write

$$\psi(x) = \sum_{j=1}^r c_j \varphi_j(x) = \sum_{j=1}^r c_j i^{\ell_j(x)} (-1)^{q_j(x)} \mathbb{1}_{A_j}(x), \tag{5}$$

where for every $j \in [r]$, ℓ_j is a linear function, q_j a quadratic function, and A_j an affine subspace.

For every $j \in [r]$, let A'_j, ℓ'_j, q'_j denote the projection of A_j, ℓ_j, q_j respectively, obtained by setting the coordinates outside of D to zero. Observe that $A'_j \subseteq \mathbb{F}_2^m$ is an affine subspace, ℓ'_j an m-variate linear function over $\mathbb{F}_2$, and q'_j an m-variate quadratic function over $\mathbb{F}_2$, and that

$$g(\bar{x}) = \begin{cases} 1 & \text{if } |\sum_{j=1}^r c_j \cdot i^{\ell'_j(\bar{x})} \cdot (-1)^{q'_j(\bar{x})} \cdot \mathbb{1}_{A'_j}(\bar{x})| \leq \left(\frac{1+\eta}{2}\right) m_{pn-1} \\ 0 & \text{otherwise.} \end{cases}$$

Let $h : \mathbb{F}_2^{3r} \to \mathbb{F}_2$ denote the following function:

$$h(y_1, \ldots, y_r, z_1, \ldots, z_r, v_1, \ldots, v_r) = \begin{cases} 1 & \text{if } |\sum_{j=1}^{r} c_j \cdot i^{y_j} \cdot (-1)^{z_j} \cdot v_j| \leq \left(\frac{1+\eta}{2}\right) m_{pn-1} \\ 0 & \text{otherwise.} \end{cases}$$

(note that here $v_j \in \{0, 1\}$ is considered as a real number). Then

$$g(\bar{x}) = h(\ell'_1(\bar{x}), \ldots, \ell'_r(\bar{x}), q'_1(\bar{x}), \ldots, q'_r(\bar{x}), \mathbb{1}_{A'_1}(\bar{x}), \ldots, \mathbb{1}_{A'_r}(\bar{x})).$$

For every $j \in [r]$, let P_j be a polynomial of degree $O(\log(r))$ such that $\Pr_{\bar{x}\in\mathbb{F}_2^m}[P_j(\bar{x}) \neq \mathbb{1}_{A'_j}(\bar{x})] \leq \frac{1}{20r}$, as guaranteed by Claim 4.5. Note that h is a function on $3r$ boolean variables, and hence can be represented exactly by a polynomial of degree at most $3r$. As the ℓ'_j's have degree 1 and q'_j's degree 2, it follows that

$$\tilde{g}(\bar{x}) = h(\ell'_1(\bar{x}), \ldots, \ell'_r(\bar{x}), q'_1(\bar{x}), \ldots, q'_r(\bar{x}), P_1(\bar{x}), \ldots, P_r(\bar{x}))$$

is a polynomial of degree $O(r \log r)$, and by the union bound

$$\Pr_{\bar{x}\in\mathbb{F}_2^m}[\tilde{g}(\bar{x}) \neq g(\bar{x})] \leq \Pr_{\bar{x}\in\mathbb{F}_2^m}\left[\exists j \in [r] \text{ such that } P_j(x) \neq \mathbb{1}_{A'_j(x)}\right] \leq \frac{1}{20}. \qquad \square$$

4.4 A Lower Bound for Approximate Stabilizer Rank via Correlation Bounds

We now observe that the results of Section 4.1, Section 4.2 and Section 4.3 imply our lower bounds. The final ingredient we require is the following correlation lower bound.

Lemma 4.7 ([21, 25])**.** *Let $f : \mathbb{F}_2^m \to \mathbb{F}_2$ be a boolean function such that*

$$\Pr_{x\in\mathbb{F}_2^m}[f(x) = \mathsf{Maj}_m(x)] \geq \frac{2}{3}.$$

Then $\deg(f) = \Omega(\sqrt{m})$.

We recall Theorem 1.2

Theorem 4.8 (Restatement of Theorem 1.2)**.** *Let $|B\rangle$ be either $|H\rangle$ or $|R\rangle$. Then, for a sufficiently small constant δ, it holds that $\chi_\delta(B^{\otimes n}) = \Omega(\sqrt{n}/\log n)$.*

Proof. Let ψ be a state such that $\|\psi - B^{\otimes n}\| \leq \delta$. By Lemma 4.2, this implies that the boolean function $f := f_\psi$, as defined in (4), $(0.01, 0.01)$-approximates THR_{pn}. By Lemma 4.1 this implies that there exists a restriction of f, g, such that

$$\Pr_{\bar{x}\in\mathbb{F}_2^m}[g(\bar{x}) \neq \mathsf{Maj}_m(\bar{x})] \leq \frac{1}{4}$$

for $m = 2pn$. Further, by Lemma 4.6, there is a polynomial $\tilde{g}$, of degree $O(r \log r)$, such that

$$\Pr_{\bar{x}\in\mathbb{F}_2^m}[g(\bar{x}) \neq \tilde{g}(\bar{x})] \leq \frac{1}{20}.$$

It follows that

$$\Pr_{\bar{x}\in\mathbb{F}_2^m}[\tilde{g}(\bar{x}) \neq \mathsf{Maj}_m(\bar{x})] \leq \frac{1}{3}$$

and thus, by Lemma 4.7, $r \log r = \Omega(\sqrt{2pn})$, as the theorem states. $\square$

References

[1] Scott Aaronson and Alex Arkhipov. The Computational Complexity of Linear Optics. *Theory Comput.*, 9:143–252, 2013. doi:10.4086/toc.2013.v009a004.

[2] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, Nov 2004. doi:10.1103/PhysRevA.70.052328.

[3] Ethan Bernstein and Umesh V. Vazirani. Quantum Complexity Theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997. doi:10.1137/S0097539796300921.

[4] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, September 2019. doi:10.22331/q-2019-09-02-181.

[5] Sergey Bravyi and David Gosset. Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates. *Phys. Rev. Lett.*, 116:250501, Jun 2016. doi:10.1103/PhysRevLett.116.250501.

[6] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A*, 71:022316, Feb 2005. doi:10.1103/PhysRevA.71.022316.

[7] Sergey Bravyi, Graeme Smith, and John A. Smolin. Trading Classical and Quantum Computational Resources. *Phys. Rev. X*, 6:021043, Jun 2016. doi:10.1103/PhysRevX.6.021043.

[8] Jeroen Dehaene and Bart De Moor. Clifford group, stabilizer states, and linear and quadratic operations over GF(2). *Phys. Rev. A*, 68:042318, Oct 2003. doi:10.1103/PhysRevA.68.042318.

[9] Yuval Filmus, Hamed Hatami, Steven Heilman, Elchanan Mossel, Ryan O'Donnell, Sushant Sachdeva, Andrew Wan, and Karl Wimmer. Real Analysis in Computer Science: A collection of Open Problems. Simons Institute, 2014. URL: https://simons.berkeley.edu/sites/default/files/openprobsmerged.pdf.

[10] Daniel Gottesman. Stabilizer Codes and Quantum Error Correction. *Dissertation (Ph.D.), California Institute of Technology*, 1997. doi:10.7907/rzr7-dt72.

[11] Lov K. Grover. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, pages 212–219. ACM, 1996. doi:10.1145/237814.237866.

[12] Cupjin Huang, Michael Newman, and Mario Szegedy. Explicit Lower Bounds on Strong Quantum Simulation. *IEEE Trans. Inf. Theory*, 66(9):5585–5600, 2020. doi:10.1109/TIT.2020.3004427.

[13] Daniel J. Kleitman. On a combinatorial conjecture of Erdös. *Journal of Combinatorial Theory*, 1(2):209–214, 1966. doi:10.1016/S0021-9800(66)80027-3.

[14] Lucas Kocia. Improved Strong Simulation of Universal Quantum Circuits. *arXiv preprint arXiv:2012.11739*, 2020. URL: https://arxiv.org/abs/2012.11739.

[15] Farrokh Labib. Stabilizer rank and higher-order Fourier analysis. *arXiv preprint arXiv:2107.10551*, 2021. URL: https://arxiv.org/abs/2107.10551.

[16] Benjamin Lovitz and Vincent Steffan. New techniques for bounding stabilizer rank. *arXiv preprint arXiv:2110.07781*, 2021. URL: https://arxiv.org/abs/2110.07781.

[17] Tomoyuki Morimae and Suguru Tamaki. Fine-grained quantum computational supremacy. *Quant. Inf. Comput.*, 19:1089, 2019. doi:10.26421/QIC19.13-14-2.

[18] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016. doi:10.1017/CBO9780511976667.

[19] Hammam Qassim, Hakop Pashayan, and David Gosset. Improved upper bounds on the stabilizer rank of magic states. *Quantum*, 5:606, December 2021. doi:10.22331/q-2021-12-20-606.

[20] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 13–23. ACM, 2019. doi:10.1145/3313276.3316315.

[21] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mat. Zametki*, 41:598–607, 1987. English translation in Mathematical Notes of the Academy of Sci. of the USSR, 41(4):333-338, 1987. doi:10.1007/BF01137685.

[22] Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. doi:10.1137/S0097539795293172.

[23] Daniel R. Simon. On the Power of Quantum Computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997. doi:10.1137/S0097539796298637.

[24] Roman Smolensky. Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 77–82. ACM, 1987. doi:10.1145/28395.28404.

[25] Roman Smolensky. On Representations by Low-Degree Polynomials. In *34th Annual Symposium on Foundations of Computer Science*, pages 130–138. IEEE Computer Society, 1993. doi:10.1109/SFCS.1993.366874.

[26] Maarten Van den Nest. Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond. *Quant. Inf. Comput.*, 10:258, 2010. doi:10.26421/QIC10.3-4-6.

[27] R. Ryan Williams. Limits on Representing Boolean Functions by Linear Combinations of Simple Functions: Thresholds, ReLUs, and Low-Degree Polynomials. In *33rd Computational Complexity Conference, CCC 2018*, volume 102 of *LIPIcs*, pages 6:1–6:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.CCC.2018.6.

A The Clifford Group and Magic States

The purpose of this section is to provide a brief introduction to the Clifford group for readers who are unfamiliar with it. We shall not cover the entire background, motivation and various applications of this group in quantum computing and quantum information, but rather only provide the bare minimum of definitions needed to understand this work and its motivation. The book [18] is good extensive reference on these topics, and in particular Sections 10.5.1 and 10.5.2 which deal with the stabilizer formalism. We also provide a notational reference to the various gates and magic states we consider in this paper.

A.1 Pauli and Clifford Group

The *Pauli matrices* are three 2×2 complex unitary matrices defined as follows:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

These matrices generate a subgroup of $2x2$ matrices of order 16, denoted by P_1 and called the single qubit Pauli group, that contains the elements

$$\{\pm I, \pm iI, \pm X, \pm iX, \pm Y, \pm iY, \pm Z, \pm iZ\}.$$

The n-qubit Pauli group, denoted P_n is defined as

$$P_n = \{\sigma_1 \otimes \sigma_2 \otimes \cdots \otimes \sigma_n : \text{for all } j \in [n], \sigma_j \in P_1\}.$$

The Clifford group $\mathcal{C}_n$ can now be defined as the normalizer of P_n in the group $U(n)$ of n-qubit unitary matrices. It is convenient, however, to consider $\mathcal{C}_n$ as a finite group, which is why it is usually defined modulo $U(1)$, i.e., we identify two matrices U and V if $U = cV$ for some $c \in \mathbb{C}$ with $|c| = 1$ (c is called a *global phase*):

$$\mathcal{C}_n := \left\{U \in U(n) : UP_nU^\dagger = P_n\right\} / U(1).$$

It turns out that $\mathcal{C}_n$ has a set of generators which is very easy to describe. Every $U \in \mathcal{C}_n$ can be generated using the following simple set of gates:

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}.$$

H is called the Hadamard gate and S is called the phase gate. The set of *stabilizer states* is the set of states φ such that $|\varphi\rangle = U|0^n\rangle$.

Evidently, $\{\text{CNOT}, H, S\}$ is thus not a universal quantum gate set. However, the set $\{\text{CNOT}, H, S, T\}$, where

$$T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}$$

is the so-called $\pi/8$ gate, *is* universal.

A.2 Magic States

As explained in Section 1.1, any circuit over the (universal) gate set $\{\text{CNOT}, H, S, T\}$ can be converted to a circuit of roughly the same size with only Clifford gates, which is given as additional inputs an ample supply of qubits in a *magic state*. The two types of magic states defined by Bravyi and Kitaev [6] are

$$|H\rangle = \cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle, \quad \text{and } |R\rangle = \cos(\beta)|0\rangle + e^{i\pi/4}\sin(\beta)|1\rangle,$$

where $\beta = \arccos(1/\sqrt{3})/2$.

We say two n-qubit states ψ and φ are *Clifford-equivalent* if $|\psi\rangle = U|\varphi\rangle$ for $U \in \mathcal{C}_n$. Up to a phase, state $|H\rangle$ is Clifford-equivalent to the state $|T\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle)$ (see [7]), and thus Clifford circuits provided with $|H^{\otimes n}\rangle$ as auxiliary inputs have the same computational power as Clifford circuits provided with $|T^{\otimes n}\rangle$.

B Proof of Claim 4.4

Proof of Claim 4.4. Recall that by Stirling's approximation, $m! \sim \sqrt{2\pi m}\,(\frac{m}{e})^m$. In particular, for large enough n,

$$\begin{aligned}\binom{n}{pn} &= \frac{n!}{(pn)!((1-p)n)!} \\ &\geq \frac{1}{2}\frac{\sqrt{2\pi n}\cdot(n/e)^n}{\sqrt{2\pi(pn)}(pn/e)^{pn}\cdot\sqrt{2\pi(1-p)n}\cdot((1-p)n/e)^{(1-p)n}}.\end{aligned}$$

Thus,

$$W_{pn} = \binom{n}{pn}p^{pn}(1-p)^{(1-p)n} = \Omega(1/\sqrt{n}),$$

where the constant hidden under the Ω notation depends on p. Now, for $C > 0$, we will show that $W_{pn}/W_{pn+C\sqrt{n}} = O(1)$ (where again, the constant depends on C and p).

$$\begin{aligned}\frac{W_{pn}}{W_{pn+C\sqrt{n}}} &= \frac{\binom{n}{pn}p^{pn}(1-p)^{(1-p)n}}{\binom{n}{pn+C\sqrt{n}}p^{pn+C\sqrt{n}}(1-p)^{(1-p)n-C\sqrt{n}}} \\ &= \frac{(pn+C\sqrt{n})\cdots(pn+1)}{((1-p)n)\cdots((1-p)n-C\sqrt{n}+1)}\cdot\left(\frac{1-p}{p}\right)^{C\sqrt{n}} \\ &\leq \left(\frac{pn+C\sqrt{n}}{(1-p)n-C\sqrt{n}}\right)^{C\sqrt{n}}\cdot\left(\frac{1-p}{p}\right)^{C\sqrt{n}} \\ &= \frac{pn}{(1-p)n}^{C\sqrt{n}}\cdot\frac{\left(1+\frac{C}{p\sqrt{n}}\right)^{C\sqrt{n}}}{\left(1-\frac{C}{(1-p)\sqrt{n}}\right)^{C\sqrt{n}}}\cdot\left(\frac{1-p}{p}\right)^{C\sqrt{n}} \\ &= \frac{\left(1+\frac{C}{p\sqrt{n}}\right)^{C\sqrt{n}}}{\left(1-\frac{C}{(1-p)\sqrt{n}}\right)^{C\sqrt{n}}}.\end{aligned}$$

The last term is bounded by a constant, as

$$\lim_{n\to\infty}\left(1+\frac{C}{p\sqrt{n}}\right)^{C\sqrt{n}} = e^{C^2/p},$$

and similarly

$$\lim_{n\to\infty}\left(1-\frac{C}{(1-p)\sqrt{n}}\right)^{C\sqrt{n}} = e^{-C^2/(1-p)}.$$

A similar calculation works when $C < 0$. □