

A New Lower Bound Technique for Quantum Circuits without Ancillæ

Debajyoti Bera

Abstract

We present a technique to derive depth lower bounds for quantum circuits. The technique is based on the observation that in circuits without ancillæ, only a few input states can set all the control qubits of a Toffoli gate to 1. This can be used to selectively remove large Toffoli gates from a quantum circuit while keeping the cumulative error low. We use the technique to give another proof that parity cannot be computed by constant depth quantum circuits without ancillæ.

1 Introduction

There are many models of computing used in computer science, e.g. Turing machines, circuits, decision trees to name a few. Each of the models have their own advantages. The circuit model is particularly useful for proving lower bounds. Unlike a Turing machine, a circuit computing a function is easy to write down and is closely linked to the algebraic and combinatorial properties of the function. This makes it tempting to try to bound the power of the circuit, when limited to a certain depth or size or both.

There is a considerable interest in low depth quantum circuits. The recent developments in physical realization of qubits and quantum gates suggest that the early quantum circuits will be of limited size and depth, if built using the classical like model of a network of gates. For example it has been shown that the quantum Fourier transform can be implemented in low depth and size [CW00]. Note as well that the quantum Fourier transform plays an important role in providing quantum speedups for a lot of problems.

Many of the currently known upper bounds in quantum circuit complexity require the *fanout gate* [HS05]. While fanout is a no-cost operation in classical circuits (by simply splitting a wire), duplication of qubits is prohibited by the *No-cloning theorem*. This makes it impossible to create identical copies of a qubit in an arbitrary state. However, limited fanout operation can be achieved by only duplicating the basis states¹; there is an unitary operation which achieves this and linearly extends to other states. A quantum gate performing this operation is known as the fanout gate.

Interestingly enough, the fanout gate was found to be depth-3 equivalent to the parity gate [M99]. On one hand this makes the upper bounds stronger since a fanout gate seems to be something that could be physically implementable [F03]. On the other hand, this begs the question whether something as simple as the fanout gate could be implemented in constant depth using the typical set of universal gates. Since fanout is equivalent to parity, this is the quantum analogue of the classical “PARITY $\in AC^0$ ” question.

It has actually been shown before that a quantum circuit requires logarithmic depth to compute parity without using any ancillæ [FFGHZ06]. Their approach is similar to the *random restriction* method used in classical circuit complexity; they selectively eliminate Toffoli gates by carefully setting input variables. Their result also shows that their lower bound extends to circuits with a sublinear number of extra qubits initialized to a fixed state.

¹ $|x, y, z\rangle \xrightarrow{F} |x, x \oplus y, x \oplus z\rangle$. If $y = z = 0$, then this copies x to the other qubits.

1.1 Importance of Ancillæ

Sometimes circuits use additional inputs gates which are initialized to a fixed value. For classical circuits, it is understood that circuits always have access to as many values of 0 and 1 as needed. It is easy to generate a 0 from any input variable (e.g. $0 = x \oplus x$) and the circuit can generate any number of 0 and 1 from a single 0 using not gates and fanout.

However, due to the restriction of arbitrary fanout in quantum circuits, the number of fixed input ancillary qubits (commonly referred to as *ancillæ*) become an important physical resource. So frequently we consider problems none, few or many ancillæ.

These ancillæ are frequently used as workspace qubits to store intermediate values. So a related question is what is the final state of the ancillæ. A circuit is said to *cleanly compute* a function if the ancillæ needs to be initialized to a particular state and the ancillæ should be returned to that state at the end of the computation. Clean computing circuits are useful for composing since the ancillæ can be reused. In contrast to this, a circuit is said to *robustly compute* a function if it works for any initial state of the ancillæ.

2 Motivation

Our technique is based on the usual observation about AND gates: even though an AND gate might act on a large number of inputs, but for most of the inputs the action of the gate has no effect. Different manifestations of this observation has led to a number of techniques in the past.

The *random restriction* approach [PB] notes that setting only one input of an AND gate to 0 is sufficient to replace the gate by a constant. It then uses this idea to randomly set several of the input variables and *kill* most of the gates of the circuit.

The algebraic techniques [B93] represent the inputs as variables and the gates as polynomial functions on these variables. The main idea here revolves around approximating AND or OR by a low-degree polynomial function. The idea is to compose the polynomials gate-by-gate and come up with a low-degree polynomial for the whole circuit.

Our approach is inspired by one of the algebraic approaches [SMO87] where classical AND gates are approximated as²:

$$\forall \delta, \exists \text{polynomial } \hat{A}(X), \Pr_{X \in_R \{0,1\}^n} [AND(X) = \hat{A}(X)] > 1 - \delta$$

where $\hat{A}(x)$ is a low-degree polynomial with degree dependent on δ . In this paper we observe that for quantum circuits without any ancillæ, the reversibility property of the circuit ensures that the output distribution after any gate is same as the input distribution. This allows us to replace the AND functions of the Toffoli gates by the degree-0 polynomial **1**. Since the states of a quantum circuit can be a superposition of basis states, instead of actually bounding the probability that the AND gate can be replaced by the degree-0 polynomial, we bound the distance between the states obtained by applying the Toffoli gate and by applying the approximation. This result gives us another proof that parity cannot be computed by constant depth circuits using unbounded Toffoli gates and single qubit gates.

3 Bounded fanin circuits

Our main result uses a communication argument to show that a circuit can ignore some of its inputs yet compute the original function approximately. A simpler variant of the argument shows that bounded fanin gates and single qubits gates cannot compute any function that is dependent on all input qubits. This was proved in [FFGHZ06]. We present a slightly different argument here that will be helpful in understanding the main result.

²It is stronger than what is given below, in fact the $x \in_R \{0,1\}^n$ needs to be replaced by x according to any distribution.

Consider a circuit of depth d , n inputs and a ancillæ consisting of single qubit and 2-qubit gates. Also assume at the end of the circuit, the measurement of the first qubit in some fixed basis is output as the circuit output (the fact is true as long as any constant number of output qubits are measured).

Since the gates are of fanin at most 2, at depth d , the state of the first qubit depends at most on a fixed set of 2^d input qubits. However, there might be gates between the other qubits. It is obvious that those gates will not affect the measurement output of the first qubit. Here is a formal proof of this fact.

For any quantum circuit with single and 2-qubit gates, there is an equivalent directed acyclic graph where single qubit gates correspond to vertices with fanin and fanout 1 and 2-qubit gates correspond to fanin and fanout 2. Let the vertex for the last gate on the first qubit be denoted by u (the measurement qubit). Then the vertices can be partitioned into two disjoint subsets U and V such that, U is the smallest subset containing u and there is no edge from V to U (perform a topological sorting of the vertices; U is the subset of vertices which have an edge to u and their closure). Note that $|U| \leq 2^d$.

Since there is no edge from V to U , this means in the circuit, the gates in U do not depend on any gates from V . So, the circuit of depth d can be transformed to an equivalent depth $2d$ circuit where the gates in V are *pushed* past layer d i.e. if V has a gate g in layer i , then the new circuit has g in layer $d + i$. The transformed circuit then has all gates from U in the first d layers working only on 2^d qubits, followed by d more layers consisting on gates from V working on all but the first qubit (since U contains u and u is the last gate on first qubit, there is no outgoing edge from u). Formally, $C = (I^{(1)} \otimes V^{(n+a-1)})(U^{(2^d)} \otimes I^{(n+a-2^d)})$ (the superscript denotes the size of the subspace the operator).

Since later gates cannot affect the measurement of the output of u , and the gates up to and including u depend only on 2^d gates, the measurement of the output of u can depend on at most 2^d inputs.

This argument also works with unbounded numbers of ancillæ.

4 Unbounded fanin circuits

For the main result we consider circuits with unbounded fanin gates. Specifically we consider circuits with unbounded Toffoli gates and any fixed set of single qubit gates. For details on quantum circuits, see the survey [SIGACT07].

4.1 Setup

We will consider a circuit on n qubits (no ancillæ) where the initial state is one of the 2^n standard basis states $\{|00\dots 0\rangle, |00\dots 1\rangle, \dots, |11\dots 1\rangle\}$. Each of the possible initial states $\{|x\rangle : x \in \{0, 1\}^n\}$ corresponds to the classical input x to the circuit. Hence the state of the circuit after any gate will be written as a function of x , e.g. $|\Psi(x)\rangle$ and should be understood as a family of states, one for each initial state.

4.2 Approximating Toffoli gates

Consider an orthonormal basis $\mathcal{B} = \{|y_1\rangle, \dots, |y_{2^n}\rangle\}$ for a 2^n dimensional vector space over n qubits. We will treat \mathcal{B} as a $2^n \times 2^n$ matrix where $\mathcal{B}_{i,j} = \langle i|y_j\rangle$ (amplitude of $|i\rangle$ in $|y_j\rangle$). Consider the unitary operator $C|i\rangle = |y_i\rangle$, $\forall i = \{0, 1\}^n$ i.e. $C = \sum_k |y_k\rangle\langle k|$. Notice that $\langle i|C|j\rangle = \langle i|y_j\rangle = \mathcal{B}_{i,j}$. So, \mathcal{B} is just the representation of C when written in the standard basis. Henceforth we will always use \mathcal{B} and think of \mathcal{B} as the unitary operator producing \mathcal{B} .

Now consider a state of the circuit $|\Psi(x)\rangle$. Note that $\{|\Psi(x)\rangle : x \in \{0, 1\}^n\}$ forms an orthonormal basis for a 2^n dimensional vector space on n qubits. If we write the matrix $[|\Psi(0)^n\rangle \dots |\Psi(1^n)\rangle]$ as $\{|\Psi(x)\rangle\}_x$, then by the previous observation we get,

Claim 4.1 *For any state $|\Psi(x)\rangle$, $\{|\Psi(x)\rangle\}_x$ is the unitary operator written in the standard basis which produces this state.*

Consider a group of t qubits T when the circuit is in $|\Psi(x)\rangle$. Wlog assume $T = \{1, \dots, t\}$. Let $f_x(j) = \langle j|\Psi(x)\rangle$ denote the (x, j) th entry of the matrix $\{|\Psi(x)\rangle\}_x$ and $q = n - (t + 1)$ denote the number of qubits not in T . Then we can write the states $\Psi(x)$, $\forall x$ as the following.

$$\begin{aligned}\Psi(x) &= \sum_{k \in \{0,1\}^n} f_x(k)|k\rangle \\ &= \sum_{j \in \{0,1\}^q} \sum_{i \in \{0,1\}^{t+1}} f_x(i \cdot j)|i\rangle|j\rangle \\ &= \sum_{j \in \{0,1\}^q} \left[f_x(1^t \cdot 0 \cdot j)|1^t \cdot 0 \cdot j\rangle + f_x(1^t \cdot 1 \cdot j)|1^t \cdot 1 \cdot j\rangle + \sum_{\substack{i \neq 1^t \cdot 0 \\ i \neq 1^t \cdot 1}} f_x(i \cdot j)|i \cdot j\rangle \right]\end{aligned}$$

Now, assume a $t + 1$ -qubit Toffoli gate \mathcal{T} is applied whose control qubits belong to T and qubit- $(t + 1)$ is the target qubit. Then for each x , all but two of the basis states will remain unchanged and two basis states will be swapped. We will get the following set of states for all x .

$$\begin{aligned}\mathcal{T}|\Psi(x)\rangle &= \sum_{j \in \{0,1\}^q} \left[f_x(1^t \cdot 0 \cdot j)\mathcal{T}|1^t \cdot 0 \cdot j\rangle + f_x(1^t \cdot 1 \cdot j)\mathcal{T}|1^t \cdot 1 \cdot j\rangle + \sum_{\substack{i \neq 1^t \cdot 0 \\ i \neq 1^t \cdot 1}} f_x(i \cdot j)\mathcal{T}|i \cdot j\rangle \right] \\ &= \sum_{j \in \{0,1\}^q} \left[f_x(1^t \cdot 0 \cdot j)|1^t \cdot 1 \cdot j\rangle + f_x(1^t \cdot 1 \cdot j)|1^t \cdot 0 \cdot j\rangle + \sum_{\substack{i \neq 1^t \cdot 0 \\ i \neq 1^t \cdot 1}} f_x(i \cdot j)|i \cdot j\rangle \right]\end{aligned}$$

Definition 4.2 For a state $|\Psi(x)\rangle$ and a gate \mathcal{T} , define the error³ from \mathcal{T} at $|\Psi(x)\rangle$ denoted by $\Delta_{\mathcal{T}}(|\Psi(x)\rangle)$, as

$$\Delta_{\mathcal{T}}(|\Psi(x)\rangle) = \|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\|^2$$

Now notice that,

$$\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle = \sum_{j \in \{0,1\}^q} (f_x(1^t \cdot 0 \cdot j) - f_x(1^t \cdot 1 \cdot j)) |1^t \cdot 1 \cdot j\rangle + (f_x(1^t \cdot 1 \cdot j) - f_x(1^t \cdot 0 \cdot j)) |1^t \cdot 0 \cdot j\rangle \quad (1)$$

So, taking the difference and summing over all x gives us,

$$\begin{aligned}\|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\|^2 &= \sum_{j \in \{0,1\}^q} 2|f_x(1^t \cdot 0 \cdot j) - f_x(1^t \cdot 1 \cdot j)|^2 \\ \frac{1}{2}\|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\|^2 &= \sum_{j \in \{0,1\}^q} |f_x(1^t \cdot 0 \cdot j) - f_x(1^t \cdot 1 \cdot j)|^2 \\ \frac{1}{2}\sum_x \|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\|^2 &= \sum_{j \in \{0,1\}^q} \sum_x |f_x(1^t \cdot 0 \cdot j) - f_x(1^t \cdot 1 \cdot j)|^2\end{aligned}$$

Since $\{|\Psi(x)\rangle\}_x$ is unitary, so is $\{|\Psi(x)\rangle\}_x^\dagger$. Form a vector from the rows of $\{|\Psi(x)\rangle\}_x^\dagger$. Let $|\phi(r)\rangle$ denote these vectors: $\forall r = 0^n \dots 1^n$, $|\phi(r)\rangle = \sum_x f_x(r)^*|x\rangle = \sum_x \langle \Psi(x)|r\rangle|x\rangle$.

Note that these vectors form an orthonormal basis for the 2^n dimensional state over n qubits. So for $r_1 \neq r_2$, $\| |\phi(r_1)\rangle - |\phi(r_2)\rangle \| = \sqrt{2}$. Again, $\| |\phi(r_1)\rangle - |\phi(r_2)\rangle \| = \| \sum_x (f_x(r_1)^* - f_x(r_2)^*)|x\rangle \| = \sqrt{\sum_x |f_x(r_1) - f_x(r_2)|^2}$. This gives us,

$$\frac{1}{2}\sum_x \|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\|^2 = \sum_{j \in \{0,1\}^q} 2 = 2 \cdot 2^q$$

This immediately gives us an estimate of the total error,

³We actually mean the square of the error!

Claim 4.3 $\sum_x \Delta_{\mathcal{T}}(|\Psi(x)\rangle) = 2 \cdot 2^n / 2^t$

Definition 4.4 For a gate \mathcal{T} applied to a state $|\Psi(x)\rangle$, the average error from \mathcal{T} at $|\Psi(x)\rangle$ is defined as

$$\mathcal{E}_{\mathcal{T}}(|\Psi(x)\rangle) = E_x(\Delta_{\mathcal{T}}(|\Psi(x)\rangle))$$

where E_x denotes the expectation taken over the random variable x ⁴.

Assuming the input to the circuit is chosen uniformly at random from $\{0,1\}^n$, it follows that $\mathcal{E}_{\mathcal{T}} = 2/2^t$. Applying the Chebyshev inequality to the above, we get the main lemma.

Lemma 4.5 For a Toffoli gate \mathcal{T} with t control qubits, acting on a state $|\Psi(x)\rangle$,

$$\Pr_x [\|\mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle\| < \epsilon] > 1 - \frac{4}{2^{t\epsilon}}$$

This gives us the average error when a Toffoli gate is removed from the circuit. We can get a similar result if we have multiple gates composed together. We can also allow gates that are not removed to be present in the circuit. The crucial fact to note here is that the states after each gate $\{|\Psi(x)\rangle\}_x$ have a one-to-one correspondence with the initial states $\{|x\rangle\}_x$ i.e. the random variables $\{|x\rangle\}_x$ and $\{|\Psi(x)\rangle\}_x$ after each gate are identically distributed. The rest is a simple application of the union bound and using the fact that errors from quantum operations are additive.

Lemma 4.6 Consider a circuit with $C = C_k \mathcal{T}_k \cdots C_1 \mathcal{T}_1 C_0$ with k Toffoli gates that are removed and C_i s denote the intermediate gates that are not removed. Let \mathcal{T}_i have t_i control qubits. Then the total average error when all the k Toffoli gates are removed is the sum of the average error for each gate, i.e.

$$\Pr_x \left[\|C_k \mathcal{T}_k C_{k-1} \mathcal{T}_{k-1} \cdots C_1 \mathcal{T}_1 C_0 |\Psi(x)\rangle - C_k \cdots C_0 |\Psi(x)\rangle\| < \epsilon \right] > 1 - \frac{4k}{\epsilon} (1/2^{t_1} + \dots + 1/2^{t_k})$$

Proof. We will give a sketch of the proof for two consecutive gates. The general case where k Toffoli gates are removed and there are other intermediate gates is similar.

Say we have two Toffoli gates \mathcal{T}_1 and \mathcal{T}_2 with t_1 and t_2 control qubits respectively. And we have a circuit $C = \mathcal{T}_2 \mathcal{T}_1$. Let $|\Psi(x)\rangle = \mathcal{T}_1 |x\rangle$ denote the state of the circuit after applying \mathcal{T}_1 . So we know,

$$\Pr_x [\|\mathcal{T}_1 |x\rangle - |x\rangle\| \geq \epsilon] < \frac{4}{2^{t_1 \epsilon}}$$

$$\Pr_x [\|\mathcal{T}_2 |\Psi(x)\rangle - |\Psi(x)\rangle\| \geq \epsilon] < \frac{4}{2^{t_2 \epsilon}}$$

By union bound,

$$\Pr_x [\|\mathcal{T}_1 |\Psi(x)\rangle - |\Psi(x)\rangle\| \geq \epsilon \text{ or } \|\mathcal{T}_2 |\Psi(x)\rangle - |\Psi(x)\rangle\| \geq \epsilon] < \frac{4}{\epsilon} (1/2^{t_1} + 1/2^{t_2})$$

$$\Pr_x [\|\mathcal{T}_1 |\Psi(x)\rangle - |\Psi(x)\rangle\| < \epsilon \text{ and } \|\mathcal{T}_2 |\Psi(x)\rangle - |\Psi(x)\rangle\| < \epsilon] > 1 - \frac{4}{\epsilon} (1/2^{t_1} + 1/2^{t_2})$$

Due to additive nature of errors from quantum gates, $\|\mathcal{T}_2 \mathcal{T}_1 |x\rangle - |x\rangle\| \leq \|\mathcal{T}_2 |\Psi(x)\rangle - |\Psi(x)\rangle\| + \|\mathcal{T}_1 |\Psi(x)\rangle - |x\rangle\|$.

This proves our two gate case, $\Pr_x [\|\mathcal{T}_2 \mathcal{T}_1 |x\rangle - |x\rangle\| < 2\epsilon] > 1 - \frac{4}{\epsilon} (1/2^{t_1} + 1/2^{t_2})$.

⁴The left hand side does not depend on x anymore but is a function of the level of the circuit. $|\Psi(x)\rangle$ merely denotes the state at that level.

4.3 Approximating the circuit

To approximate a circuit of depth d and on n inputs without any ancillæ, we start from the input level of the circuit, remove *large* gates, move to the next higher level and then repeat the same. At each level we remove Toffoli gates with on t or more number of control qubits; $t = t(n, d)$ is a fixed value per circuit family that will be fixed later. Thus the error from each removed Toffoli gate is upper bounded by

$$\Pr_x [\| \mathcal{T}|\Psi(x)\rangle - |\Psi(x)\rangle \| < \epsilon] > 1 - \frac{4}{2^{|T|-1}\epsilon} > 1 - \frac{4}{2^t\epsilon}$$

There are at most n/t large Toffoli gates in each layer. Thus at most nd/t Toffoli gates are removed. The earlier lemma about composing error when multiple Toffoli gates are removed also work when there are other unitary gates between the replaced Toffoli gates as long as they are not removed. This gives us:

Claim 4.7 *For the whole circuit C on n qubits and with depth d , if we replace all Toffoli gates of width $t+1$ or more then we get a modified circuit C' where,*

$$\Pr_x [\| C|\Psi(x)\rangle - C'|\Psi(x)\rangle \| < \epsilon] > 1 - \frac{4nd}{t\epsilon} \left(\frac{nd}{t} \frac{1}{2^t} \right)$$

4.4 Lower bound for parity

If $t+1$, the maximum width of allowed Toffoli gates is less than $n^{1/d}$, then no qubit at the final layer, specially the measurement qubit, can be connected to all n input qubits. So we choose a suitable $t < n^{1/d} - 1$.

Then the measurement value of C' (in any fixed basis) can be used as an algorithm to compute the output of C (assuming C is a circuit with a classical output i.e. the measurement gives 0 with probability either 0 or 1). C' will output the same value as C with probability $> (1 - \epsilon)(1 - \frac{4n^2d^2}{t^2\epsilon 2^t})$ but that output value will not depend on all inputs.

If C is a circuit to compute parity in constant depth (d is constant), then we can choose $t = n^{1/d} - 2$ and $\epsilon = 1/4$. Then,

$$(1 - \epsilon) \left(1 - \frac{4n^2d^2}{t^2\epsilon 2^t} \right) > \frac{3}{4} \left(1 - 64 \frac{n^2d^2}{2^{n/d}} \right) > 3/4$$

for most n . Since parity cannot be computed with probability greater than $1/2$ without looking at all the inputs, C could not computed parity exactly on all its inputs.

Theorem 4.8 *There is no constant depth quantum circuit which computes exactly the parity of its inputs on all possible inputs using only single qubit gates and Toffoli gates and using no ancillæ qubits.*

5 Conclusion

The obvious open question is how to use this technique for circuits with ancillæ. Fang et al. [FFGHZ06] showed how to get a similar lower bound for quantum circuits with a sublinear number of ancillæ. However we believe the lower bound is true for unbounded number of ancillæ. While it is tempting to use the technique used in the paper for circuits with ancillæ, it is not immediately clear what constitutes a *large gate*. This is because, with ancillæ, the states of the qubits are copied and so it could happen that all the inputs in a seemingly large Toffoli gate are just the copies of a single actual input.

There are two ways an unbounded fanin quantum circuit is different than an unbounded fanin classical circuit. One is the usage of the single qubit gates which are the primary sources of *entanglement*. Observe that the Toffoli gate is basically a classical gate. The other is the restriction of fanout.

Here we would like to point out an interesting resemblance between quantum and classical circuits. Quantum circuits without any ancillæ behave very much like classical bounded fanin circuits. Without unbounded fanout, the unbounded fanin gates are not capable of quickly mixing up information to generate

a complex output. The number of gates in a circuit with no fanout gate is bounded by the product of the depth and the total number of qubits (inputs and ancillæ), so unlike the classical circuits, the number of ancillæ qubits in a quantum circuit is an important resource. It would be interesting to come up with explicit functions which cannot be computed without or with a limited number ancillæ but can be computed with lots of ancillæ.

We would like to thank Jin-Yi Cai, Fred Green and Steve Homer for interesting discussions.

References

- [B93] R. Beigel. “The Polynomial Method in Circuit Complexity”. In *Proc. Structure in Complexity Theory Conference*, 82–95, 1993.
- [CW00] R. Cleve and J. Watrous. “Fast parallel circuits for the quantum Fourier transform”. In *Proc. of 41st Symp. on Foundations of Computer Science*, 2000
- [F03] S.A. Fenner. “Implementing the fanout gate by a Hamiltonian”. (Unpublished) *arxiv:quant-ph/0309163*, 2003.
- [FFGHZ06] M. Fang, S. Fenner, F. Green, S. Homer and Y. Zhang. “Quantum lower bounds for fanout”. *Quantum Information and Computation*, 6(1):046–057, 2006.
- [HS05] P. Hyer and R. palek. “Quantum circuits with unbounded fan-out”. *Theory of Computing*, 1:81-103, 2005.
- [M99] Cristopher Moore. “Quantum Circuits: Fanout, Parity, and Counting”. In *Los Alamos Preprint archives* quant-ph/9903046, 1999.
- [PB] P. Beame. “A Switching Lemma Primer”. manuscript <http://www.cs.washington.edu/homes/beame/papers.html>.
- [SIGACT07] D. Bera, F. Green and S. Homer. “Small depth quantum circuits”. *SIGACT News*, 38(2):35–50, 2007.
- [SMO87] R. Smolensky. “Algebraic methods in the theory of lower bounds for Boolean circuit complexity”. In *Proc. 19th ACM Symposium on Theory of Computing (STOC’87)*, ACM, 77-82, 1987.