The Parametrized Complexity of Quantum Verification

Srinivasan Arunachalam
IBM Quantum, Thomas J Watson Research Center, Yorktown Heights, NY, USA

Sergey Bravyi
IBM Quantum, Thomas J Watson Research Center, Yorktown Heights, NY, USA

Chinmay Nirkhe
IBM Quantum, Thomas J Watson Research Center, Yorktown Heights, NY, USA
Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA
Challenge Institute for Quantum Computation, University of California, Berkeley, CA, USA

Bryan O'Gorman
IBM Quantum, Thomas J Watson Research Center, Yorktown Heights, NY, USA

Abstract

We initiate the study of parameterized complexity of QMA problems in terms of the number of non-Clifford gates in the problem description. We show that for the problem of parameterized quantum circuit satisfiability, there exists a classical algorithm solving the problem with a runtime scaling exponentially in the number of non-Clifford gates but only polynomially with the system size. This result follows from our main result, that for any Clifford + t T-gate quantum circuit satisfiability problem, the search space of optimal witnesses can be reduced to a stabilizer subspace isomorphic to at most t qubits (independent of the system size). Furthermore, we derive new lower bounds on the T-count of circuit satisfiability instances and the T-count of the W-state assuming the classical exponential time hypothesis (ETH). Lastly, we explore the parameterized complexity of the quantum non-identity check problem.

2012 ACM Subject Classification Theory of computation → Quantum computation theory

Keywords and phrases parametrized complexity, quantum verification, QMA

Digital Object Identifier 10.4230/LIPIcs.TQC.2022.3

Related Version *Previous Version*: https://arxiv.org/abs/2202.08119

Funding *Sergey Bravyi*: Was supported in part by the IBM Research Frontiers Institute.
Chinmay Nirkhe: Acknowledges support from NSF Quantum Leap Challenges Institute Grant number OMA2016245 and an IBM Quantum PhD internship.

Acknowledgements Part of this work was completed while CN and BO were participants in the Simons Institute for the Theory of Computing *Summer Cluster on Quantum Computation*. Additionally, we thank Sam Gunn, Zeph Landau, Dimitri Maslov and Kristan Temme for insightful discussions.

1 Introduction

The solutions to many important computational problems require resources that seem to scale exponentially in the system size in the worst case. Parameterized complexity refines this phenomenon by identifying one or more parameters along with algorithms that scale exponentially only in the identified parameters (but polynomially in system size). In the worst case, these parameters typically scale with system size, but often interesting instances have an intermediate value of the parameter. In quantum computing, parameterized complexity has

been applied to the classical simulation of quantum systems [13, 9, 8, 22, 6, 23]. In particular, parameterizing circuits by the number of non-Clifford gates has yielded many state-of-the-art algorithms for classical simulation. Here, we initiate the study of the parameterized complexity of *non-deterministic* computation, i.e., *quantum verification*. Specifically, we consider QMA (Quantum Merlin-Arthur[1]) problems parameterized by the number of non-Clifford gates in their verification circuits and obtain non-trivial upper bounds on the classical complexity of finding an optimal witness and the number of qubits required for its representation.

1.1 The parameterized complexity of quantum circuit satisfiability

The first problem we consider is quantum circuit satisfiability (QCSAT), a canonical QMA-complete problem. In a QCSAT instance, the input is an s-gate quantum circuit U acting on $n+m$ qubits followed by the measurement in the standard basis of any $k > 0$ output qubits.[2] The goal in the QCSAT problem is to estimate the maximal probability that quantum circuit measurement outputs 1^k when run on input states (i.e., witnesses) of the form $|\psi\rangle \otimes |0^m\rangle$ for $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$:

$$\mathsf{Val} \stackrel{\text{def}}{=} \max_{|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}} \langle \psi, 0^m | \, U^\dagger \, |1^k\rangle\langle 1^k| \, U \, |\psi, 0^m\rangle \, . \tag{1}$$

The problem can also be phrased as a promise decision problem in which the goal is to decide if $\mathsf{Val} > a$ (yes instance) or $\mathsf{Val} < b$ (no instance) for $a > b$. The decision problem is known to be QMA-complete when $a = 2/3$ and $b = 1/3$.

To the best of our knowledge, there are no previously known classical algorithms that exploit the structure of the circuit to solve QCSAT in less than $\mathsf{exp}(n)$ time; a simple exponential time algorithm for calculating Val can be achieved by searching over the entire Hilbert space of the witness $|\psi\rangle$. One of the main results of our work is that parameterized QCSAT instances with t T-gates for $t \ll n$ can be solved significantly faster than this naive algorithm. We show that there is a stabilizer subspace isomorphic to $(\mathbb{C}^2)^{\otimes t}$ of the the input Hilbert space which contains all optimal witnesses; here an optimal witness is any input state ψ that maximizes the probability of observing the output 1^k.

▶ **Theorem 1.** *For every* QCSAT *instance U with $t \leq n$ T-gates, there is an n-qubit Clifford unitary W and a t-qubit state $|\phi\rangle$ such that $W\left(|\phi\rangle \otimes |0^{n-t}\rangle\right)$ is an optimal input state. Furthermore, the Clifford unitary W only depends on the description of U and can be computed in (classical, deterministic) time* $\mathsf{poly}(n, s)$.

This insight can be used to construct a faster algorithm for parameterized QCSAT instances.

▶ **Theorem 2.** *There exists a classical randomized algorithm that takes as input a parameterized instance of* QCSAT *problem, a precision parameter $\delta > 0$, and outputs a real random variable ξ such that*

$$0 \leq \xi \leq \mathsf{Val} \quad \textit{and} \quad \Pr[\xi \geq (1-\delta)\mathsf{Val}] \geq \frac{99}{100}. \tag{2}$$

The algorithm has runtime $\mathsf{poly}(n, m, s, t) + O(\delta^{-1} t 2^t)$.

---

[1] The complexity class QMA is the quantum analog of the classical non-deterministic complexity classes, MA and NP [18].

[2] While one can map the k-qubit measurement to a single measurement, this requires the application of a coherent AND logical gate. Implementing this AND gate requires $\Omega(k)$ non-Clifford gates; in the case of parameterized complexity, this cost may be prohibitive. For this reason, we define the problem in terms of the measurement of multiple qubits in the standard basis.

Quantum circuit satisfiability (QCSAT) is the analog of quantum circuit simulation (QMA- vs. BQP-completeness) in the same way that classical circuit satisfiability is the analog[3] of circuit simulation (NP- vs. P-completeness). Recall that it is widely believed that any classical algorithm for parameterized quantum circuit simulation should have a runtime scaling exponentially in t; the current matching upper-bound scales as $2^{\alpha t}$, where $\alpha < 0.3963$ for exact simulators [23] and $\alpha < 0.23$ for approximate simulators [8]. Our result shows that the QCSAT verification problem is not much harder than its simulation counterpart; the resulting exponential scaling is worse, but there is no exponential dependence on the instance size or n, the size of the witness. This is surprising as *classical* circuit satisfiability is believed to require a runtime scaling exponentially with n (the size of the witness) to solve, while classical circuit simulation is trivially solvable in polynomial-time. Therefore, it would be reasonable to expect that a parameterized QCSAT instance would incur a slowdown scaling exponentially with the witness size n due to non-determinism of the problem and a slowdown scaling exponentially with t due to its quantumness. We instead show that the exponential time scaling can be brought to scale with only t when $t \ll n$. This is the primary power of Theorem 1: to efficiently reduce the search space of optimal inputs from n qubits to t qubits.

Our result may seem surprising in view of the earlier work by Morimae et al. [21] that studied a restricted version of the class QMA where the verifier can only perform Clifford gates. In [21], they found that QMA with a Clifford verifier coincides with the standard QMA. However, the computational model of [21] is different from ours since it allows *adaptive* Clifford gates that can be classically controlled by the outcomes of intermediate measurements. In contrast, here we consider unitary (non-adaptive) verification circuits with all measurements delayed until the end.

Let us briefly sketch the proof of Theorems 1 and 2; complete proofs are provided in Section 2. By definition, solving QCSAT is equivalent to estimating the largest eigenvalue of an operator $\rho \stackrel{\text{def}}{=} \langle 0^m| U^\dagger |1^k\rangle\langle 1^k| U |0^m\rangle$ acting on n qubits. Consider first a simple case when there are no T gates, i.e., $t = 0$ and U is a Clifford circuit. At a high level, ρ describes a state (unnormalized) generated by a sequence of Clifford operations: (1) initializing each qubit in a basis state or a maximally mixed state, (2) applying a unitary Clifford gate, and (3) post-selectively measuring a qubit in the standard basis. Such operations are known to have very limited computational power – they always produce a (mixed) stabilizer-type state [14]. Accordingly, the largest eigenvalue of ρ can be efficiently computed using the standard stabilizer formalism [10]. Suppose now that U contains t T-gates. It is well-known [5] that a T-gate can be implemented by a gadget that includes only Clifford operations and consumes one copy of a magic state $|A\rangle \propto |0\rangle + e^{i\pi/4}|1\rangle$. Replacing each T-gate in U by this gadget one gets $\rho = \mathrm{Tr}_{\{[t]\}}((\,|A^t\rangle\langle A^t| \otimes \mathbb{I}_n)\rho')$, where ρ' is a bipartite stabilizer state of $t+n$ qubits and we trace out the first t qubits. Our key technical tool is the characterization of bipartite stabilizer states [7]. This result implies that $\rho' = (C_1^\dagger \otimes C_2^\dagger)\rho''(C_1 \otimes C_2)$, where C_i are unitary Clifford operators and ρ'' is a tensor product of local single-qubit stabilizer states and at most t two-qubit stabilizer states shared between the two subsystems. Using this decomposition we are able to show that $\rho = C_2^\dagger(\rho_{\text{hard}} \otimes \rho_{\text{easy}})C_2$, where ρ_{hard} is some (non-stabilizer) state of t qubits and ρ_{easy} is a stabilizer state of $n-t$ qubits whose eigevalues are easy to compute. Thus, QCSAT reduces to estimating the largest eigenvalue of the t-qubit state ρ_{hard}. We remark that this theorem also holds if the T-gate is replaced by an arbitrary angle Z-rotation (since we use a post-selective magic state injection gadget, in which case

[3] Technically, they are the analogs of randomized circuit satisfiability and randomized circuit simulation (MA- vs. BPP-completeness).

it does not matter if the rotation angle is $\pi/8$ or not). To prove Theorem 2 we make use of the special structure of the state ρ_{hard} and reduce the problem of computing its largest eigenvalue to computing the largest Schmidt coefficient of a certain pure bipartite state of at most t qubits. The latter is computed using the power method with a random starting state [19].

Implications for QCMA vs. QMA

The description of a circuit generating a quantum state can constitute a classical witness for that state. It is an open question (QCMA vs. QMA) in complexity theory if all quantum witnesses have efficient classical descriptions [3, 2, 11]. Our work makes progress on the parameterized version of the question by proving that the witness to any QCSAT instance with t T-gates can be constructed with at most $\mathsf{exp}(t)$ T-gates as it only requires t qubits to describe. An interesting open question is if our techniques lend themselves to any sub-exponential in t upper-bound on the T-count of optimal witness states.

1.2 Implied lower bounds from the exponential time hypothesis

Theorem 2 provides an upper bound on the runtime of a classical algorithm for QCSAT with respect to the number of T-gates in the verifier. In conjunction with other complexity-theoretic assumptions, this also implies a *lower bound* on the T-count of the verification circuit. One such assumption is the Exponential-Time Hypothesis (ETH), introduced by Impagliazzo and Paturi [15]. Informally, ETH is the conjecture that classical k-SAT requires exponential classical time. By Theorem 2, a verifier circuit for QMA with $o(n)$ T-gates would imply a $2^{o(n)}$-time algorithm for k-SAT, because NP $\subset$ QMA. Thus we get the following lower bound.

▶ **Corollary 3.** ETH *implies that any* QMA*-complete family of Clifford+T verifier circuits must include circuits with $\Omega(n)$ T-gates, where n is the size of the witness.*

Interestingly, we can also use ETH and Theorem 2 to get a lower bound on the T-count of a quantum circuit that prepares the m-qubit W-state

$$|W_m\rangle = \frac{1}{\sqrt{m}}\big(\,|100\cdots0\rangle + |010\cdots0\rangle + \cdots + |000\cdots1\rangle\,\big). \tag{3}$$

▶ **Corollary 4.** ETH *implies that any Clifford+T circuit V that, when applied to the all-zero state, outputs $|W_m\rangle \otimes |\mathsf{junk}\rangle$ must include $\Omega(m)$ T gates.*

Proofs of both corollaries are provided in Section 3. To our knowledge, this is the first lower bound for the T-count of state preparation based on complexity-theoretic assumptions. The proof uses the NP-hardness [4] of Hamiltonians of the form $H = \sum_{\{i,j\}\in E} Z_iZ_j + \sum_{i\in V} Z_i$. We show that a verifier circuit can be built from only a W state and Clifford operations.

1.3 The complexity of the non-identity check problem

The second QMA-complete problem we consider in this work is the Non-Identity Check (NIC) problem: given a classical description of an n-qubit quantum circuit U, decide if U is close to the identity operation, i.e., is $\min_\phi \left\|U - e^{i\phi}\cdot\mathbb{I}\right\| \geq c$ or $\leq s$ for $c - s \geq 1/\mathsf{poly}(n)$, promised

one of them is the case.[4] NIC was first considered by Janzing et al. [16] who showed that this problem is QMA-complete, by reducing it to the QCSAT problem. Subsequently, Ji and Wu [17] showed that the NIC problem remains QMA-complete for even depth-2 circuits with *arbitrary* gates. This motivates a natural question: what is the parameterized complexity of NIC? In particular, how does the complexity of NIC scale with the number of non-Clifford gates? Below we show that NIC can be solved in polynomial-time for Clifford circuits.

▶ **Theorem 5.** NIC *for Clifford circuits is contained in* P*.*

We sketch the proof of this theorem here while a complete proof is provided in Section 4. By definition, solving NIC for a Clifford circuit U on n qubits is equivalent to estimating the maximum eigenvalue of a Hamiltonian $H = (\mathbb{I} - U)^\dagger(\mathbb{I} - U)$, that is, checking if $\|H\| \geq c^2$ or $\leq s^2$. In the former case, observe that $\mathrm{Tr}(H^p) \geq c^{2p}$ and in the latter case $\mathrm{Tr}(H^p) \leq 2^n \cdot s^{2p}$. Furthermore, since $c - s \geq 1/\mathsf{poly}(n)$, we can pick $p = \mathsf{poly}(n)$ to make $c^{2p} \gg 2^n \cdot s^{2p}$. Thus it suffices to estimate $\mathrm{Tr}(H^p)$ with $p = \mathsf{poly}(n)$. To this end, observe that H^p can be written as a weighted sum of Clifford powers U^i with $i \in \{-p, \ldots, p\}$. Thus $\mathrm{Tr}(H^p) = \sum_{i=-p}^{p} \alpha_i \mathrm{Tr}\big(U^i\big)$ and the coefficients α_i can be efficiently computed using a simple recursive formula. Furthermore, we can efficiently compute $\mathrm{Tr}\big(U^i\big)$ for every i using the identity $\mathrm{Tr}\big(U^i\big) = 2^n \langle \Phi^{\otimes n} | U^i \otimes \mathbb{I} | \Phi^{\otimes n} \rangle$, where $|\Phi\rangle$ is the EPR state. The right-hand side is the inner product between two stabilizer states of $2n$ qubits. Such inner products can be computed exactly in time $O(n^3)$, see [12, 8]. Putting all this together, we can compute $\mathrm{Tr}(H^p)$ exactly in time $\mathsf{poly}(n)$, which determines if this is a "yes" or "no" instance of the NIC problem.

In addition, we also prove that the NIC problem for constant-depth circuits using gates from a constant-sized gate set is solvable for vanishing completeness parameters. This is in contrast to [17] which shows constant-depth circuits using gates from a general gate set is QMA-complete.

▶ **Theorem 6.** *Let $\mathcal{G}$ be any constant-sized gate set of 1 and 2 qudit gates and U a quantum circuit of depth at most $t = O(1)$ acting on n qudits of fixed finite local dimension d. Let $a < b$ be any parameters such that $a(n) = o(1)$. Then the* NIC *problem (a, b) for this restricted class of circuits is in* P*.*

Our success at understanding parameterized QCSAT came from our characterization of optimal witness states of parameterized verifier circuits as small linear combinations of stabilizers. However, the eigenvectors of parameterized NIC circuits U do not have such a characterization. We need to develop new tools to classically describe the eigenvectors of U in order to achieve an equivalent result for the parameterized NIC problem. Our inability to do so might suggest that the parameterized problem is harder than we previously suspected.

We conclude with the following intriguing question regarding the parameterized complexity of NIC problems. Since the algorithm for NIC for Clifford circuits breaks down in the presence of a single non-Clifford gate, the question of parameterized complexity of NIC is left largely open. It may happen that this problem becomes hard (e.g., NP-hard) in the presence of a single non-Clifford gate.

[4] We remark that the operator norm is important here; the problem is in BQP if the goal is to decide if $\|U - \mathbb{I}\|_2$ is small or large, where $\|\cdot\|_2$ is the normalized-2 norm [20].

2 Proofs of Theorems 1 and 2

▶ **Theorem 1.** *For every* QCSAT *instance U with $t \leq n$ T-gates, there is an n-qubit Clifford unitary W and a t-qubit state $|\phi\rangle$ such that $W(|\phi\rangle \otimes |0^{n-t}\rangle)$ is an optimal input state. Furthermore, the Clifford unitary W only depends on the description of U and can be computed in (classical, deterministic) time* $\mathsf{poly}(n, s)$.

Proof. Suppose U is a Clifford+T circuit with $c = s - t$ Clifford gates and t T-gates. We assume that U acts on $n + m$ qubits partitioned into a witness register of n qubits and an ancilla register of m qubits. Let Π_{out} be a projector onto the all-ones state $|1\rangle\langle 1|^{\otimes k}$ applied to some designated output register of k qubits. We assume that Π_{out} acts trivially on the remaining $n + m - k$ qubits. Define the *maximum acceptance probability* of U as

$$\pi(U) = \max_{\psi} \langle \psi \otimes 0^m | U^\dagger \Pi_{\mathsf{out}} U | \psi \otimes 0^m \rangle = \max_{\psi} \| \Pi_{\mathsf{out}} U | \psi \otimes 0^m \rangle \|^2, \tag{4}$$

where the maximum is over all normalized n-qubit witness states $|\psi\rangle$. Let

$$\pi(U, \psi) := \| \Pi_{\mathsf{out}} U | \psi \otimes 0^m \rangle \|^2. \tag{5}$$

We shall implement each T-gate in U by the following well-known post-selection gadget:

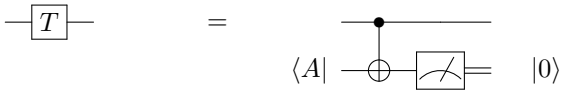

Here, we measure the output qubit in the standard basis and post-select on a measurement outcome of 0. The gadget consumes one copy of a single-qubit magic state

$$|A\rangle \stackrel{\text{def}}{=} \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle) \tag{6}$$

and, therefore, we can rewrite $\pi(U, \psi)$ as

$$\pi(U, \psi) = 2^t \| \Pi^{(1)}_{\mathsf{out}} C | \psi \otimes 0^m \otimes A^{\otimes t} \rangle \|^2. \tag{7}$$

Here C is a Clifford circuit acting on $n+m+t$ qubits with $c+t$ gates (with c gates originating from U and t CNOT gates originating from the gadgets) and $\Pi^{(1)}_{\mathsf{out}}$ is a product of Π_{out} and single-qubit projectors $|0\rangle\langle 0|$ applied to the second qubit of each T-gate gadget. The extra factor 2^t takes into account that each gadget succeeds with the probability $1/2$. Define $\Pi^{(2)}_{\mathsf{out}} = C^\dagger \Pi_{\mathsf{out}} C$. Then

$$\pi(U, \psi) = 2^t \| \Pi^{(2)}_{\mathsf{out}} | \psi \otimes 0^m \otimes A^{\otimes t} \rangle \|^2. \tag{8}$$

Let us say that a projector Π acting on n qubits is a *stabilizer projector* if it can be written as $\Pi = C^\dagger(|0^{n-k}\rangle\langle 0^{n-k}| \otimes \mathbb{I}_k)C$ for some integer $k \in [0, n]$ and some unitary Clifford operator C on n qubits. We shall use the following facts.

▶ **Fact 7** ([14])**.** *Suppose Π is a stabilizer projector on n qubits. Then $\mathrm{Tr}_{\{t\}}((|0\rangle\langle 0| \otimes \mathbb{I}_{n-1})\Pi) = \sigma\Pi'$ for some $\sigma \in \{0, 1, 1/2\}$ and some stabilizer projector Π' on $n-1$ qubits. One can compute σ and Π' in time* $\mathsf{poly}(n)$.

▶ **Fact 8** (Bipartite Stabilizer Projectors [7])**.** *Suppose Π is a stabilizer projector acting on a bipartite system LR where L and R are arbitrary qubit registers. Then there exist unitary Clifford operators C_L and C_R acting on L and R respectively such that*

$$\Pi = (C_L \otimes C_R)^\dagger \Pi' (C_L \otimes C_R) \tag{9}$$

where Π' is a tensor product of the following one-qubit and two-qubit stabilizer projectors:

- *Single-qubit projectors* $|0\rangle\langle 0|$ *and* I.
- *Two-qubit projectors* $|\Phi^+\rangle\langle\Phi^+|$ *where* $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$.
- *Two-qubit projectors* $|00\rangle\langle 00| + |11\rangle\langle 11|$.

Furthermore, each two-qubit projector acts on one qubit in L *and one qubit in* R*. The above decomposition can be computed in time poly*$(|L| + |R|)$.

By definition, $\Pi^{(2)}_{\mathsf{out}}$ is a stabilizer projector on $n + m + t$ qubits. Fact 7 implies that

$$\langle 0^m|\Pi^{(2)}_{\mathsf{out}}|0^m\rangle = \gamma 2^{-r}\Pi^{(3)}_{\mathsf{out}} \tag{10}$$

for some $\gamma \in \{0, 1\}$, integer $r \in \{0, \ldots, m\}$, and some stabilizer projector $\Pi^{(3)}_{\mathsf{out}}$ on $n+t$ qubits. From Eqs. (8,10) one gets

$$\pi(U, \psi) = \gamma 2^{t-r}\|\Pi^{(3)}_{\mathsf{out}}|\psi \otimes A^{\otimes t}\rangle\|^2. \tag{11}$$

If $\gamma = 0$ then $\pi(U, \psi) = 0$ for any witness $|\psi\rangle$. Accordingly, one can choose the Clifford unitary W in the statement of the theorem arbitrarily. From now on we assume that $\gamma = 1$. Apply Fact 8 to the stabilizer projector $\Pi^{(3)}_{\mathsf{out}}$ and the partition $[n + t] = LR$ where L is the n-qubit witness register and R is the t-qubit magic state register. We get

$$\Pi^{(3)}_{\mathsf{out}} = (C_L \otimes C_R)^\dagger \Pi^{(4)}_{\mathsf{out}}(C_L \otimes C_R) \tag{12}$$

where $\Pi^{(4)}_{\mathsf{out}}$ is a tensor product of one-qubit and two-qubit stabilizer projectors from Fact 8. Substituting this into Eq. (11) with $\gamma = 1$ one gets

$$\pi(U, \psi) = 2^{t-r}\|\Pi^{(4)}_{\mathsf{out}}C_L|\psi\rangle \otimes C_R|A^{\otimes t}\rangle\|^2. \tag{13}$$

Equivalently,

$$\pi(U, C_L^\dagger\psi) = 2^{t-r}\|\Pi^{(4)}_{\mathsf{out}}|\psi\rangle \otimes C_R|A^{\otimes t}\rangle\|^2. \tag{14}$$

By definition, the register R contains t qubits. Thus $\Pi^{(4)}_{\mathsf{out}}$ may contain at most t two-qubit projectors $|\Phi^+\rangle\langle\Phi^+|$ and $|00\rangle\langle 00| + |11\rangle\langle 11|$. Indeed, each two-qubit projector must have one qubit in R, see Fact 8. Partition $L = L'L''$ such that each two-qubit projector that appears in $\Pi^{(4)}_{\mathsf{out}}$ has one qubit in L'' and the other qubit in R. Then

$$|L''| \le t \qquad \text{and} \qquad \Pi^{(4)}_{\mathsf{out}} = \Gamma_{L'} \otimes \Lambda_{L''R} \tag{15}$$

for some stabilizer projectors Γ and Λ. Here the subscripts indicate the registers acted upon by each projector. Furthermore, Γ is a tensor product of single-qubit projectors $|0\rangle\langle 0|$ and I. Define an operator

$$\Pi^{(5)}_{\mathsf{out}} \stackrel{\mathsf{def}}{=} {}_R\langle A^{\otimes t}|C_R^\dagger\Lambda_{L''R}C_R|A^{\otimes t}\rangle_R \tag{16}$$

acting on the register L''. Note that $\Pi^{(5)}_{\mathsf{out}}$ is Hermitian and positive semi-definite (although $\Pi^{(5)}_{\mathsf{out}}$ might not be a projector). Then

$$\pi(U, C_L^\dagger\psi) = 2^{t-r}\langle\psi|\Gamma_{L'} \otimes \Pi^{(5)}_{\mathsf{out}}|\psi\rangle. \tag{17}$$

Here the tensor product separates L' and L''. It follows that $|\psi\rangle$ is an optimal witness such that $\pi(U) = \pi(U, \psi)$ if

$$C_L|\psi\rangle = |\phi_{L'}\rangle \otimes |\phi_{L''}\rangle, \tag{18}$$

where $|\phi_{L'}\rangle$ is a $(+1)$-eigenvector of the projector $\Gamma_{L'}$ and $|\phi_{L''}\rangle$ is an eigenvector of $\Pi^{(5)}_{\mathsf{out}}$ with the maximum eigenvalue. From Equation (15) one infers that $|\phi_{L''}\rangle$ is a state of at most t qubits. Since $\Gamma_{L'}$ is a product of $|0\rangle\langle 0|$ and $\mathbb{I}$ terms, divide L' into L'_1 and L'_2 such that

$$\Gamma_{L'} = |0^{|L'_1|}\rangle\langle 0^{|L'_1|}|_{L'_1} \otimes \mathbb{I}_{L'_2}. \tag{19}$$

Then, the minimizing $\phi_{L'}$ has the form $\phi_{L'} = \left|0^{|L'_1|}\right\rangle_{L'_1} \otimes |\mathsf{junk}\rangle_{L'_2}$ for any state $|\mathsf{junk}\rangle$. To conclude, one can choose an optimal witness state $|\psi\rangle$ such that

$$|\psi\rangle = C_L^\dagger \left(|0^{|L'_1|}\rangle \otimes |\mathsf{junk}\rangle \otimes |\phi_{L''}\rangle\right) \tag{20}$$

for some $(\leq t)$-qubit state $|\phi_{L''}\rangle$ and some Clifford operators C_L. This is equivalent to the statement of the theorem. Additionally observe that all the above steps necessary to obtain W can be implemented efficiently since they only involve manipulations with Clifford circuits and stabilizer projectors. ◀

▶ **Theorem 2.** *There exists a classical randomized algorithm that takes as input a parameterized instance of* QCSAT *problem, a precision parameter $\delta > 0$, and outputs a real random variable ξ such that*

$$0 \leq \xi \leq \mathsf{Val} \quad \textit{and} \quad \Pr[\xi \geq (1-\delta)\mathsf{Val}] \geq \frac{99}{100}. \tag{2}$$

The algorithm has runtime $\mathsf{poly}(n, m, s, t) + O(\delta^{-1} t 2^t)$.

Proof. First let us introduce some notations. Let $\mathcal{H}(n)$ be the set of all normalized n-qubit pure states. Given a hermitian n-qubit operator M, let $\lambda_{max}(M) = \max_{\psi\in\mathcal{H}(n)} \langle\psi|M|\psi\rangle$ be the maximum eigenvalue of M. Define two-qubit projectors

$$\Gamma^{(1)} = |\Phi^+\rangle\langle\Phi^+| \quad \text{and} \quad \Gamma^{(2)} = |00\rangle\langle 00| + |11\rangle\langle 11|. \tag{21}$$

Recall that $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Given a pair of disjoint k-qubit registers R and L, let $\Gamma^{(i)}_{RL}$ be a $(2k)$-qubit projector that applies $\Gamma^{(i)}$ to the j-th qubit of R and the j-th qubit of L for each $j = 1, \ldots, k$. Below we follow notations introduced in the proof of Theorem 1. Our starting point is the expression for the quantity Val derived in Eqs. (16,17,18) thereof, namely

$$\mathsf{Val} = 2^{t-r} \lambda_{max}(\Pi^{(5)}_{\mathsf{out}}), \tag{22}$$

where $\Pi^{(5)}_{out}$ is a positive semi-definite operator defined in Eq. (16), namely

$$\Pi^{(5)}_{\mathsf{out}} = {}_R\langle A^{\otimes t}| C_R^\dagger \Lambda_{L''R} C_R |A^{\otimes t}\rangle_R. \tag{23}$$

Recall that R and L'' are disjoint qubit registers such that $|R| = t$, $|L''| \leq t$, C_R is some Clifford operator acting on R, and $\Lambda_{L''R}$ is a product of one- and two-qubit stabilizer projectors from Fact 8 such that each one-qubit projector acts on R and each two-qubit projector acts on both R and L''. In other words, $\Lambda_{L''R}$ can be written as

$$\Lambda_{L''R} = \Gamma^{(1)}_{L_1R_1} \Gamma^{(2)}_{L_2R_2} |0\rangle\langle 0|_{R_3} \mathbb{I}_{R_4} \tag{24}$$

for some partitions $L'' = L_1L_2$ and $R = R_1R_2R_3R_4$ with $|R_1| = |L_1|$ and $|R_2| = |L_2|$. Below we use notations $r_i = |R_i|$ and $\ell_i = |L_i|$.

We claim that $\Pi^{(5)}_{\mathsf{out}}$ commutes with any Pauli operator Z_j, $j \in L_2$. Indeed, it suffices to check that Z_j commutes with $\Lambda_{L''R}$. The latter acts on the register L_2 by a diagonal operator $\Gamma^{(2)}_{L_2R_2}$ which commutes with Z-type Pauli operators confirming the claim. Thus one can choose the maximum eigenvector ψ of $\Pi^{(5)}_{\mathsf{out}}$ such that $Z_j|\psi\rangle = (-1)^{\sigma_j}|\psi\rangle$ for all $j \in L_2$ and some unknown $\sigma \in \{0,1\}^{\ell_2}$. Equivalently, $|\psi\rangle = |\psi'\rangle_{L_1} \otimes |\sigma\rangle_{L_2} \equiv |\psi'_{L_1} \otimes \sigma_{L_2}\rangle$ for some $\psi' \in \mathcal{H}(\ell_1)$. Using Eq. (22) one gets

$$\mathsf{Val} = 2^{t-r} \max_{\psi \in \mathcal{H}(\ell_1)} \max_{\sigma \in \{0,1\}^{\ell_2}} \langle \psi_{L_1} \otimes \sigma_{L_2} | \Pi^{(5)}_{\mathsf{out}} | \psi_{L_1} \otimes \sigma_{L_2} \rangle. \tag{25}$$

We shall discard the register L_2 using the identity

$${}_{L_2}\langle \sigma | \Gamma^{(2)}_{L_2R_2} | \sigma \rangle_{L_2} = |\sigma\rangle\langle\sigma|_{R_2}. \tag{26}$$

Substituting this identity into Eq. (25) gives

$$\mathsf{Val} = 2^{t-r} \max_{\sigma \in \{0,1\}^{\ell_2}} \lambda_{max}(\Pi^{(6)}_{out}(\sigma)) \tag{27}$$

where $\Pi^{(6)}_{out}(\sigma)$ is a positive semi-definite operator acting on the register L_1 defined as

$$\Pi^{(6)}_{out}(\sigma) = {}_R\langle A^{\otimes t} | C^\dagger_R \Gamma^{(1)}_{L_1R_1} | \sigma \rangle\langle \sigma |_{R_2} |0\rangle\langle 0|_{R_3} \mathbb{I}_{R_4} C_R | A^{\otimes t} \rangle_R. \tag{28}$$

We shall discard the register L_1 using the quantum teleportation identity

$${}_{L_1}\langle \psi | \Gamma^{(1)}_{L_1R_1} | \psi \rangle_{L_1} = 2^{-\ell_1} |\psi^*\rangle\langle\psi^*|_{R_1} \tag{29}$$

which holds for any state $\psi \in \mathcal{H}(\ell_1)$. Here ψ^* is the complex conjugate of ψ in the standard basis of ℓ_1 qubits. Using the teleportation identity and the definition of $\Pi^{(6)}_{out}(\sigma)$ one can check that

$$\langle \psi | \Pi^{(6)}_{out}(\sigma) | \psi \rangle = 2^{-\ell_1} \langle A^{\otimes t} | C^\dagger_R | \psi^* \rangle\langle \psi^* |_{R_1} | \sigma \rangle\langle \sigma |_{R_2} |0\rangle\langle 0|_{R_3} \mathbb{I}_{R_4} C_R | A^{\otimes t} \rangle_R \tag{30}$$

for any state $\psi \in \mathcal{H}(\ell_1)$. At this point both registers L_1 and L_2 have been discarded. After some algebra one can rewrite Eq. (30) as

$$\langle \psi | \Pi^{(6)}_{out}(\sigma) | \psi \rangle = 2^{-\ell_1} \| {}_{R_1}\langle \psi^* | \varphi(\sigma) \rangle_{R_1R_4} \|^2, \tag{31}$$

where $|\varphi(\sigma)\rangle$ is a state of R_1R_4 defined as

$$|\varphi(\sigma)\rangle = {}_{R_2R_3}\langle \sigma_{R_2}, 0_{R_3} | C_R | A^{\otimes t} \rangle_R. \tag{32}$$

Since the set $\mathcal{H}(\ell_1)$ is closed under the complex conjugation, Eqs. (27,31) give

$$\mathsf{Val} = 2^{t-r-\ell_1} \max_{\sigma \in \{0,1\}^{\ell_2}} \max_{\psi \in \mathcal{H}(\ell_1)} \| {}_{R_1}\langle \psi | \varphi(\sigma) \rangle_{R_1R_4} \|^2. \tag{33}$$

At this point the only remaining registers are R_1 and R_4. Clearly, the optimal state $\psi \in \mathcal{H}(\ell_1)$ that achieves the maximum in Eq. (33) coincides with the largest eigenvector of a reduced density matrix

$$\rho_{R_1}(\sigma) = \mathrm{Tr}_{R_4} |\varphi(\sigma)\rangle\langle\varphi(\sigma)|. \tag{34}$$

Thus Eq. (33) is equivalent to

$$\mathsf{Val} = 2^{t-r-\ell_1} \max_{\sigma \in \{0,1\}^{\ell_2}} \lambda_{max}(\rho_{R_1}(\sigma)). \tag{35}$$

Recall that any t-qubit Clifford operator can be efficiently compiled to a Clifford circuit with $O(t^2)$ one- and two-qubit gates [1]. Thus one can compute a t-qubit state $|\phi\rangle := C_R|A^{\otimes t}\rangle_R$ as a vector of complex amplitudes in time $\mathsf{poly}(t)2^t$ using the standard state vector simulator of quantum circuits. We assume that ϕ is stored in a classical RAM memory for the rest of the algorithm such that any amplitude of ϕ can be accessed in time $\mathsf{poly}(t)$. By definition, $|\varphi(\sigma)\rangle$ is obtained from $|\phi\rangle$ by projecting the registers R_2R_3 onto the basis state $|\sigma_{R_2}, 0_{R_3}\rangle$, see Eq. (32). Equivalently, $|\varphi(\sigma)\rangle$ is obtained from $|\phi\rangle$ by selecting a subset of $2^{t-r_2-r_3}$ amplitudes. Thus one can compute $|\varphi(\sigma)\rangle$ as a vector of complex amplitudes in time $poly(t)2^{t-r_2-r_3}$ for any given σ. By definition, $|\varphi(\sigma)\rangle$ is a state of $t-r_2-r_3$ qubits. We shall use the following fact.

▶ **Lemma 9.** *Suppose $|\varphi\rangle$ is a pure n-qubit state specified as a vector of complex amplitudes and $\delta > 0$ is a precision parameter. Consider a partition $[n] = AB$, where A and B are disjoint qubit registers. Let $\rho_A = \mathrm{Tr}_B|\varphi\rangle\langle\varphi|$ be the reduced density matrix of A. There exists a classical randomized algorithm that runs in time $O(\delta^{-1}n2^n)$ and outputs a real random variable ξ such that*

$$0 \le \xi \le \lambda_{max}(\rho_A) \quad \textit{and} \quad \Pr[\xi \ge (1-\delta)\lambda_{max}(\rho_A)] \ge \frac{99}{100}. \tag{36}$$

Proof. Assume wlog that $|A| \le |B|$ (otherwise switch A and B). We have $\rho_A = \eta\eta^\dagger$, where η is a matrix of size $2^{|A|} \times 2^{|B|}$ with matrix elements $\langle x|\eta|y\rangle = \langle x_A y_B|\varphi\rangle$. Given a list of amplitudes of $|\varphi\rangle$, one can can compute the matrix of η in time $O(2^n)$ since $|A| + |B| = n$. For any $|A|$-qubit state $|v\rangle$ the matrix-vector product $|v\rangle \to \eta^\dagger|v\rangle$ can be computed in time $O(2^n)$. Likewise, for any $|B|$-qubit state $|w\rangle$ the matrix-vector product $|w\rangle \to \eta|w\rangle$ can be computed in time $O(2^n)$. We conclude that the matrix-vector product $|v\rangle \to \rho_A|v\rangle$ can be computed in time $O(2^n)$.

We shall compute an estimator ξ satisfying Eq. (36) using the power method with a random starting state [19]. Namely, let $|\phi\rangle \in \mathcal{H}(|A|)$ be a Haar-random state of A. Given a number of iterations $q \ge 2$, the power method outputs an estimator

$$\xi_q = \frac{\langle\psi_q|\rho_A|\psi_q\rangle}{\langle\psi_q|\psi_q\rangle}, \qquad |\psi_q\rangle := \rho_A^{q-1}|\phi\rangle. \tag{37}$$

Clearly, computing ξ_q requries q matrix-vector multiplications for the matrix ρ_A. Thus the runtime scales as $O(q2^n)$. Note that $0 \le \xi_q \le \lambda_{max}(\rho_A)$ with certainty. Theorem 3.1 of [19] guarantees that the relative error

$$\epsilon_q := \frac{\lambda_{max}(\rho_A) - \xi_q}{\lambda_{max}(\rho_A)} \tag{38}$$

obeys

$$\mathbb{E}(\epsilon_q) \le \frac{0.871\log(2^{|A|})}{q-1} \le \frac{n}{q} \tag{39}$$

for all $q \ge 2$. Here the expectation value is taken over the random starting state ϕ and we use the natural logarithm. Since ϵ_q is a non-negative random variable, Markov inequality implies that $\epsilon_q \le 100 \cdot \mathbb{E}(\epsilon_q)$ with the probability at least 99/100. Thus the desired estimator ξ satisfying Eq. (36) can be chosen as $\xi = \xi_q$ with $q = \lceil 100n/\delta \rceil$. ◀

Applying Lemma 9 to the state $|\varphi(\sigma)\rangle$ of $n = t - r_2 - r_3$ qubits with the registers $A = R_1$ and $B = R_4$ one obtains an estimator $\xi(\sigma)$ satisfying

$$0 \le \xi(\sigma) \le \lambda_{max}(\rho_{R_1}(\sigma)) \quad \text{and} \quad \Pr[\xi(\sigma) \ge (1-\delta)\lambda_{max}(\rho_{R_1}(\sigma))] \ge 99/100. \tag{40}$$

The runtime required to compute the estimator $\xi(\sigma)$ for any fixed σ is $O(\delta^{-1}t2^{t-r_2-r_3})$. Thus computing the estimators $\xi(\sigma)$ for all $\sigma \in \{0,1\}^{\ell_2}$ takes time $O(\delta^{-1}t2^t)$. Here we noted that $r_2 = \ell_2$. We choose the desired estimator ξ approximating the quantity Val as

$$\xi = 2^{t-r-\ell_1} \max_{\sigma\in\{0,1\}^{\ell_2}} \xi(\sigma). \tag{41}$$

Let $\sigma^* \in \{0,1\}^{\ell_2}$ be the optimal bit string that achieves the maximum in Eq. (35) such that

$$\mathsf{Val} = 2^{t-r-\ell_1}\lambda_{max}(\rho_{R_1}(\sigma^*)). \tag{42}$$

From Eq. (40) one infers that $\xi \leq \mathsf{Val}$ with certainty and

$$\Pr[\xi \geq (1-\delta)\mathsf{Val}] \geq \Pr[\xi(\sigma^*) \geq (1-\delta)\lambda_{max}(\rho_{R_1}(\sigma^*)] \geq 99/100. \tag{43}$$

We conclude by noting that all manipulations performed in the proof of Theorem 1 to compute the Clifford circuit C_R and the stabilizer projector $\Lambda_{L''R}$ can be implemented in time $\mathsf{poly}(n,m,s,t)$ using the standard stabilizer formalism [14]. Thus the total runtime required to compute the desired estimator ξ is

$$\mathsf{poly}(n,m,s,t) + O(\delta^{-1}t2^t). \tag{44}$$

◀

In Appendix A (Theorem 16) we give an alternative algorithm for solving the same problem as Theorem 2 but using slightly different techniques. It has some advantages over Theorem 2 which we elaborate in Appendix A.

3 Implied lower bounds from the Exponential-Time Hypothesis

The Exponential-Time Hypothesis (ETH), introduced by Impagliazzo and Paturi [15], is the conjecture that, informally, (classical) k-SAT requires exponential (classical) time.

▶ **Definition 10** (Exponential-Time Hypothesis [15]). *Let*

$$s_k \stackrel{\text{def}}{=} \inf\left\{\delta : \textit{there exists } 2^{\delta n}\textit{-time algorithm for solving } k\textit{-SAT}\right\}. \tag{45}$$

Then $s_k > 0$ is a constant for all $k \geq 3$.

It is a stronger assumption than $\mathsf{P} \neq \mathsf{NP}$, which implies just that k-SAT requires superpolynomial-time. We show that ETH, together with Theorem 2, implies T-count lower bounds, which is Corollary 3.

▶ **Corollary 3.** ETH *implies that any* QMA*-complete family of Clifford+T verifier circuits must include circuits with $\Omega(n)$ T-gates, where n is the size of the witness.*

Proof. Consider a family of Clifford + T-gate QMA verifier circuits and let $f(n)$ be the number of T gates in the circuits for n-qubit witnesses. Suppose that $f(n) = o(n)$. Then by Theorem 2, each instance can be solved in $O\left(\mathsf{poly}(n)2^{o(n)}\right)$ deterministic classical time, which contradicts ETH. The theorem follows from the fact that $\mathsf{NP} \subseteq \mathsf{QMA}$. ◀

Second, we show that Theorem 2 along with ETH implies a linear lower-bound on the T-count complexity of generating a W state. This is Corollary 4.

▶ **Corollary 4.** ETH *implies that any Clifford+T circuit V that, when applied to the all-zero state, outputs $|W_m\rangle \otimes |\mathsf{junk}\rangle$ must include $\Omega(m)$ T gates.*

The proof will use the following fact, due to Barahona [4]:

▶ **Theorem 11** ([4])**.** *Estimating, to within inverse polynomial additive error, the ground state energy of the following class of Hamiltonians is NP-hard:*

$$H = \sum_{\{i,j\}\in E} Z_i Z_j + \sum_{i\in V} Z_i, \tag{46}$$

where $G = (V, E)$ is a planar graph with maximum degree 3.

Proof of Corollary 4. Consider a circuit V starting from all-zeroes that outputs $|W_m\rangle\otimes|\mathsf{junk}\rangle$ with t T gates. We will show that this implies a $2^{O(t)}$-time classical algorithm for k-SAT, thus proving the Corollary by contradiction.

Let H' be a diagonal Hamiltonian of the form of Equation (46) with $m' = O(n)$ terms, where n is the number of vertices in the graph. Let $H = H' - m' \leq 0$ be H' shifted by a constant so that it's negative semidefinite. It will be convenient to write H as a sum of $m = 2m'$ terms:

$$H = \sum_{i=1}^{m} H_i = \sum_{i=1}^{m'} Z_{S_i} - \sum_{i=m'}^{2m'} 1 \tag{47}$$

where S_i is a set of one or two indices and $Z_S = \prod_{i\in S} Z_i$.

Let C be the circuit consisting of m gates constructed in the following way. It acts on an m-qubit "control" register A and an n-qubit "computational" register B. For $1 \leq i \leq m'$, each gate C_i implements the i-th term of H (either Z or ZZ) on the corresponding qubits of the computational register, controlled on the i-th qubit of the control register. Because each term of the Hamiltonian is Pauli, the controlled version is Clifford, and so C is a Clifford operator. For $m' < i \leq 2m'$, the corresponding gate is simply Z_i (in the control register).

Let $U \stackrel{\text{def}}{=} (V^\dagger \otimes I) \cdot C \cdot (V \otimes I)$; U has $2t$ T gates. The probability of measuring all zeros on the control register given the input state $U|0^m, \psi\rangle_{A,B}$ is

$$\mathrm{Tr}\left[|0^m\rangle\langle 0^m|_A \otimes I_B V^\dagger C V |0^m, \psi\rangle\langle 0^m, \psi|_{A,B} V^\dagger C^\dagger V\right] \tag{48}$$

$$= \mathrm{Tr}\left[|W^m\rangle\langle W^m|_A \otimes I_B C |W^m, \psi\rangle\langle W^m, \psi|_{A,B} C^\dagger\right] \tag{49}$$

$$= \frac{1}{m^2} \sum_{i,j,k,\ell} \mathrm{Tr}\left[|e_k\rangle\langle e_\ell|_A I_B C |e_i, \psi\rangle\langle e_j, \psi|_{A,B} C^\dagger\right] \tag{50}$$

$$= \frac{1}{m^2} \sum_{i,j,k,\ell} \mathrm{Tr}\left[|e_k\rangle\langle e_\ell|_A I_B H_i |e_i, \psi\rangle\langle e_j, \psi|_{A,B} H_j\right] \tag{51}$$

$$= \frac{1}{m^2} \sum_{i,j} \mathrm{Tr}\left[H_i |\psi\rangle\langle\psi|_B H_j\right] \tag{52}$$

$$= \frac{1}{m^2} \sum_{i,j} \langle\psi|H_i H_j|\psi|\psi|H_i H_j|\psi\rangle \tag{53}$$

$$= \frac{1}{m^2} \langle\psi|H^2|\psi|\psi|H^2|\psi\rangle\,. \tag{54}$$

Because H is negative semidefinite, the state ψ that maximizes the probabality of measuring all zeros as above also minimizes $\langle\psi|H|\psi|\psi|H|\psi\rangle$. By Theorem 2, such a W-state preparation circuit V with t T-gates implies a $2^{O(t)}$-time classical algorithm for solving k-SAT. ◀

4 The complexity of quantum non-identity check

▶ **Definition 12** (Non-Identity Check [16]). *An instance of the non-identity check (*NIC*) problem is a classical description of a quantum circuit U and two real numbers a, b such that $b > a$ with the promise that*

$$d_{\mathbb{I}}(U) \stackrel{\text{def}}{=} \min_{\phi} \left\| U - e^{i\phi} \cdot \mathbb{I} \right\| \tag{55}$$

is either at most a or at least b. The instance is called a yes *instance if $d_{\mathbb{I}}(U) \geq b$ and a* no *instance if $d_{\mathbb{I}}(U) \leq a$.*

In this section, we present two scenarios in which the non-identity check problem becomes trivial to solve. It was proved by Ji and Wu [17], that the problem is in general QMA-hard for depth 2 circuits over qudits when $b - a = 1/\mathsf{poly}(n)$ and the quantum gates are specified to $\Omega(\log n)$ bits of precision. We first show that NIC is solvable in P when the entire circuit is Clifford regardless of the depth of the circuit. Second, we show that if $a = o(1)$ is a sub-constant function, then the problem is in P for constant-depth circuits built from a finite gate set.

4.1 Clifford circuits

▶ **Theorem 5.** NIC *for Clifford circuits is contained in* P.

Proof. In order to see this, let C be the unknown Clifford circuit for which we need to determine whether $\|C - \mathbb{I}\| \geq \alpha$ or $\|C - \mathbb{I}\| \leq \beta$ for $\alpha - \beta \geq 1/\mathsf{poly}(n)$. To this end, consider a Hamiltonian $H = (\mathbb{I} - C)^\dagger(\mathbb{I} - C) = 2\mathbb{I} - C - C^\dagger$, whose norm satisfies $\|H\| = \|\mathbb{I} - C\|^2 \leq 4$. In order to solve NIC$(C)$, we need to decide if $\|H\| \geq \alpha^2$ or $\|H\| \leq \beta^2$. In the former case observe that $\mathrm{Tr}(H^p) \geq \alpha^{2p}$ and in the latter case we have $\mathrm{Tr}(H^p) \leq 2^n \cdot \beta^{2p}$. Since $\alpha - \beta \geq 1/\mathsf{poly}(n)$, it suffices to pick $p = \mathsf{poly}(n)$, in order to satisfy $\alpha^{2p} \geq 2 \cdot 2^n \cdot \beta^{2p}$. Hence if an algorithm could estimate $\mathrm{Tr}(H^p)$ well enough, then it can distinguish if C satisfies the Yes or No instance of NIC problem.

We now show how to compute $\mathrm{Tr}(H^p)$ for $p = \mathsf{poly}(n)$ exactly and efficiently. In this direction, observe that we can express H^p in terms of sum of Clifford powers, i.e.,

$$\mathrm{Tr}(H^p) = \mathrm{Tr}\left(\sum_{i=-p}^{p} a_i C^i \right) = \sum_{i=-p}^{p} a_i \, \mathrm{Tr}\left(C^i\right) \tag{56}$$

for some coefficients $a_i \in \mathsf{R}$. Furthermore, since H assumes a simple form $H = \mathbb{I} - C$, the $(2p+1)$ coefficients a_i can be computed explicitly in $\mathsf{poly}(n)$ time. Now, it remains to compute $\mathrm{Tr}\left(C^i\right)$ for each one of the $2p+1$ terms. The trace of a Clifford power can be computed exactly by observing that $\mathrm{Tr}\left(C^i\right) = 2^n \langle \Phi | \mathbb{I} \otimes C^i | \Phi \rangle$ where $|\Phi\rangle = \frac{1}{\sqrt{2^n}} \sum_i |i, i\rangle$ is the maximally entangled state on $(2n)$-qubits. Observe that since $|\Phi\rangle$ is a stabilizer state, we have that $|\psi\rangle = \mathbb{I} \otimes C^i |\Phi\rangle$ is also a $(2n)$-qubit stabilizer state. It remains to estimate inner product between two stabilizer states $|\Phi\rangle$ and $|\psi\rangle$. It is known that the inner product between any n-qubit stabilizer states can be computed exactly (including the overall phase) in time $O(n^3)$, see [12, 8]. Therefore, one can compute each one of the $\mathrm{Tr}\left(C^i\right)$ and a_i in Equation (56) in time $\mathsf{poly}(n)$ and overall since $p = \mathsf{poly}(n)$, we can exactly compute $\mathrm{Tr}(H^p)$ in time $\mathsf{poly}(n)$. This suffices to decide if $\mathrm{Tr}(H^p) \geq \alpha^{2p}$ (the YES instance of NIC) or $\mathrm{Tr}(H^p) \leq 2^n \cdot \beta^{2p}$ (the NO instance of NIC) for $\alpha - \beta \geq 1/\mathsf{poly}(n)$. ◀

4.2 Constant-sized gate sets

In this section, we are going to show that the NIC problem for $a = o(1)$, is in P if we restrict ourselves to circuits of constant depth and a constant-sized gate set. Curiously, the problem was shown to be QMA-hard when either we use an arbitrary gate set or constant-sized gate sets, but we allow circuits of $\Omega(\log n)$-depth [17].

▶ **Theorem 6.** *Let $\mathcal{G}$ be any constant-sized gate set of 1 and 2 qudit gates and U a quantum circuit of depth at most $t = O(1)$ acting on n qudits of fixed finite local dimension d. Let $a < b$ be any parameters such that $a(n) = o(1)$. Then the* NIC *problem (a, b) for this restricted class of circuits is in* P*.*

For this proof, we will need a few definitions and preliminary lemmas which we list here first and prove after the proof of the theorem. First, we will need a wonderful fact about low-depth circuits that the reduced density matrix $\mathrm{tr}_{-i} U\psi U^\dagger$ only depends on the lightcone of the ith qudit.

▶ **Fact 13.** *Consider a quantum state ψ on n qudits and U a quantum circuit. For any subset A of the qudits, let L_A be the support of the lightcone of A with respect to U. Then,*

$$\mathrm{tr}_{-A}(U\psi U^\dagger) = \mathrm{tr}_{-A}\left(U_{L_A}\left(\psi_{L_A}\otimes\nu_{-L_A}\right)U_{L_A}^\dagger\right) \tag{57}$$

where ν is the maximally mixed quantum state and U_{L_A} the circuit restricted to gates contained in L_A.

Second, we notice that if a quantum circuit U is close to identity overall, then the reduced action of the circuit on any region must also be close to identity. We will often use the contrapositive of this statement: if the reduced action of a circuit on any small region is far from identity, then the circuit overall is far from identity.

▶ **Fact 14.** *Let U be a quantum circuit on n qudits and let a be a constant such that $d_{\mathbb{I}}(U) < a$. Then for all states ψ and all regions A,*

$$\left\|\mathrm{tr}_{-A}(U\psi U^\dagger) - \psi_A\right\| \leq a. \tag{58}$$

Proof of Theorem 6. Let $\mathcal{C}(\ell, h)$ be the collection of all quantum circuits acting on ℓ qudits and of depth $\leq h$ consisting of gates only from $\mathcal{G}$. Let us define the *increment-distance* $\eta_{\ell,h}$ as

$$\eta_{\ell,h} \stackrel{\text{def}}{=} \min_{\substack{V\in\mathcal{C}(\ell,h)\\ V\neq e^{i\phi}\mathbb{I}}} d_{\mathbb{I}}(V). \tag{59}$$

Since $\mathcal{G}$ is a finite gate set and $\mathcal{C}(\ell, h)$ has a bounded cardinality, $\eta_{\ell,h} > 0$ is a well-defined constant independent on n and represents the closest a circuit can be to being identity without being identity itself. This is formalized in the following fact.

▶ **Fact 15.** *Let V be a circuit $\in \mathcal{C}(\ell, h)$ such that $d_{\mathbb{I}}(V) < \eta_{\ell,h}$. There exists an angle ϕ_V such that $V = e^{i\phi_V}\mathbb{I}$.*

In order to construct a P algorithm for this problem, we notice that since $a = o(1)$, for some sufficiently large N_0, if $n > N_0$

$$a(n) < \eta_{2^{t+1},t}. \tag{60}$$

Our algorithm will solve only instance of this size or larger. Assume that an instance U of the problem is a False instance, so U is near-identity. Then for each qubit i and every state ψ,

$$\left\| \mathrm{tr}_{-i} \left(U_{L_i} \left(\psi_{L_i} \otimes \nu_{-L_i} \right) U_{L_i}^\dagger \right) - \psi_i \right\| \leq a \tag{61}$$

as a consequence of the prior stated facts. Since, this holds for all states ψ, then $d_{\mathbb{I}}(U_{L_i}) < a$. However, the circuit U_{L_A} acts on at most 2^{t+1} qubits and has depth at most t. Since $a < \eta_{2^{t+1},t}$, then we can conclude that the action of U_{L_i} on the ith qubit must be $\mathbb{I}$ (up to phase) in every False instance. Since this holds for every qubit i, in a False instance, the circuit U must exactly be $\mathbb{I}$ (up to phase). Therefore, the P algorithm is simple: test if each circuit U_{L_i} is exactly identity and if so report False or otherwise report True. ◀

References

1. Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.
2. Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 115–128, 2007. doi:10.1109/CCC.2007.27.
3. Dorit Aharonov and Tomer Naveh. Quantum NP-a survey. *arXiv quant-ph/0210077*, 2002.
4. F Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253, October 1982. doi:10.1088/0305-4470/15/10/028.
5. Sergei Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71, March 2004. doi:10.1103/PhysRevA.71.022316.
6. 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.
7. Sergey Bravyi, David Fattal, and Daniel Gottesman. GHZ extraction yield for multipartite stabilizer states. *Journal of Mathematical Physics*, 47(6):062106, 2006.
8. Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by Clifford gates. *Phys. Rev. Lett.*, 116:250501, June 2016. doi:10.1103/PhysRevLett.116.250501.
9. Sergey Bravyi, Graeme Smith, and John A. Smolin. Trading classical and quantum computational resources. *Phys. Rev. X*, 6:021043, June 2016. doi:10.1103/PhysRevX.6.021043.
10. David Fattal, Toby S Cubitt, Yoshihisa Yamamoto, Sergey Bravyi, and Isaac L Chuang. Entanglement in the stabilizer formalism. *arXiv preprint quant-ph/0406168*, 2004.
11. Bill Fefferman and Shelby Kimmel. Quantum vs. classical proofs and subset verification. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS*, volume 117 of *LIPIcs*, pages 22:1–22:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
12. Hector J Garcia, Igor L Markov, and Andrew W Cross. Efficient inner-product algorithm for stabilizer states, 2012. arXiv:1210.6646.
13. Hector J Garcia-Ramirez. *Hybrid Techniques for Simulating Quantum Circuits using the Heisenberg Representation*. PhD thesis, University of Michigan, 2014.
14. Daniel Gottesman. The heisenberg representation of quantum computers. *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, 1998. arXiv:arXiv:quant-ph/9807006.
15. Russell Impagliazzo and Ramamohan Paturi. On the Complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, March 2001. doi:10.1006/jcss.2000.1727.

16 Dominik Janzing, Pawel Wocjan, and Thomas Beth. "non-identity-check" is QMA-complete. *International Journal of Quantum Information*, 03(03):463–473, 2005. doi:10.1142/S0219749905001067.

17 Zhengfeng Ji and Xiaodi Wu. Non-identity check remains QMA-complete for short circuits. *arXiv preprint*, 2009. arXiv:0906.5416.

18 Alexei Yu Kitaev, Alexander Shen, Mikhail N Vyalyi, and Mikhail N Vyalyi. *Classical and quantum computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Soc., 2002.

19 Jacek Kuczyński and Henryk Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM journal on matrix analysis and applications*, 13(4):1094–1122, 1992.

20 Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. *Theory Comput.*, 7:1–81, 2016. doi:10.4086/toc.gs.2016.007.

21 Tomoyuki Morimae, Masahito Hayashi, Harumichi Nishimura, and Keisuke Fujii. Quantum Merlin-Arthur with Clifford Arthur. *arXiv preprint*, 2015. arXiv:1506.06447.

22 Bryan O'Gorman. Parameterization of Tensor Network Contraction. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:19, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

23 Hammam Qassim, Hakop Pashayan, and David Gosset. Improved upper bounds on the stabilizer rank of magic states. *arXiv preprint*, 2021. arXiv:2106.07740.

A Alternative algorithm to Theorem 2

In Theorem 16, we give an alternative algorithm for solving the same problem as Theorem 2. While this algorithm will have a inferior worst-case runtime than Theorem 2, it may run significantly faster depending on the structure of the problem instance. Furthermore, it has the added advantage that its runtime can be efficiently calculated in time $\mathsf{poly}(n, s)$. Therefore, one can quickly compute the runtime of Theorem 16 and compare it to that of Theorem 2 and run the faster algorithm. While not faster in a worst-case sense, it may prove optimal for many physical instances. In addition, Theorem 16 can handle not just T gates but all single-qubit phase gates $G = \mathrm{diag}(1, e^{i\theta})$ which may be an advantage for some problems[5]. [6]

▶ **Theorem 16.** *For $t \leq n$, there exists a classical algorithm for parametrized* QCSAT *instance with a single-qubit output register consisting of Clifford and phase gates $G = \mathrm{diag}(1, e^{i\theta})$, running in worst-case time $O(\mathsf{poly}(n, s)2^{3t}\log(t/\delta))$, which decides if* $\mathsf{Val} > c$ *or* $< c - \delta$. *Furthermore, there exists a classical* $\mathsf{poly}(n, s)$ *routine to calculate the runtime of this algorithm without running the algorithm itself.*

Proof. Let us notice that our goal is to compute the largest eigenvalue of $\langle 0^m| U^\dagger |1\rangle\langle 1|_1 U |0^m\rangle$ due to Equation (1). Since $|1\rangle\langle 1| = \frac{\mathbb{I}}{2} - \frac{Z}{2}$, this is equivalent to computing the smallest eigenvalue of $\langle 0^m| U^\dagger Z_1 U |0^m\rangle$. In the case that U is a completely Clifford circuit, $Q^{(s)} = U^\dagger Z_1 U$ is a Pauli matrix $Q^{(s)} = \alpha P_1 \otimes P_2 \otimes \dots P_{n+m}$ for $\alpha \in \{\pm 1, \pm i\}$ and $P_i \in \{\mathbb{I}, X, Y, Z\}$.

[5] It is also easy to extend this algorithm to all k-qubit non-Clifford gates. However, the runtime will now scale as 4^{t+kt}. This follows directly from the fact that any k-qubit non-Clifford gate can be expressed as the linear combination of 4^k Clifford gates.

[6] Theorem 2 can also be extended to general phase gates, but at the cost of potentially weaker upper bounds on the constant α in the exponent.

Therefore,

$$\langle 0^m | U^\dagger Z U | 0^m \rangle = P_1 \otimes \ldots \otimes P_n \cdot \prod_{j=n+1}^{n+m} \langle 0 | P_j | 0 \rangle . \tag{62}$$

Then, the smallest eigenvalue of this matrix is easy to calculate as it is in tensor product; this is effectively the Gottesman-Knill theorem [14]. Furthermore, the Pauli $P^{(s)}$ can be computed efficiently. More specifically, if $U = g_1 \ldots g_s$ with each gate g_i a Clifford matrix, we can define and compute the sequence of Paulis $Q^{(k)} \stackrel{\text{def}}{=} g_k Q^{(k-1)} g_k^\dagger$ for $Q^{(0)} = Z_1$ from $k = 1, \ldots, s$.

We now extend this algorithm to the case that U contains t non-Clifford gates. Consider first the case that there is one non-Clifford qubit rotation gate g_k, with g_k

$$g_k \stackrel{\text{def}}{=} R(\theta_k) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta_k} \end{pmatrix} \tag{63}$$

and acts (without loss of generality) on the first qubit. Let $Q^{(k-1)}$ be the Pauli calculated up to gate g_{k-1}. Now notice that there are 2 cases to consider; namely if the action of $Q^{(k-1)}$ on the first qubit is $\in \{\mathbb{I}, Z\}$ or is $\in \{X, Y\}$. Since $R(\theta)$ commutes with $\mathbb{I}$ and Z and the following commutation relations hold:

$$R(\theta) X R(\theta)^\dagger = (\cos\theta) X + (\sin\theta) Y, \quad R(\theta) Y R(\theta)^\dagger = (-\sin\theta) X + (\cos\theta) Y, \tag{64}$$

we can express

$$Q^{(k)} = g_k Q^{(k-1)} g_k^\dagger = P^{(k,1)} + P^{(k,2)} \tag{65}$$

where $P^{(k,1)}$ and $P^{(k,2)}$ are Pauli matrices scaled by a complex number. For every subsequent Clifford gate $g_{k'}$ we can then recursively define and compute

$$P^{(k',\ell)} \stackrel{\text{def}}{=} g_{k'} P^{(k'-1,\ell)} g_{k'}^\dagger \tag{66}$$

which will remain a Pauli matrix. It is easy to note that this bifurcation from one Pauli matrix to two Pauli matrices when commuting past a non-Clifford gate generalizes to multiple non-Clifford gates with

$$U Z_1 U^\dagger = Q^{(s)} = \sum_{\ell=1}^{\leq 2^t} P^{(s,\ell)} \tag{67}$$

being expressible as the linear combination of $\leq 2^t$ Pauli matrices each acting on $n+m$ qubits. We next show that although there are $\leq 2^t$ Pauli matric, there exists an efficiently computable basis of size $b \leq t+1$ such that each Pauli matrix can be expressed as a product of the basis matrices.

▶ **Lemma 17.** *Let U be a quantum circuit consisting of s gates of which at most t gates are non-Clifford qubit rotation gates. Then $Q^{(s)} = U Z_1 U^\dagger$ can be expressed as a sum of $\leq 2^b$ Pauli matrices which are each products of at most $b \leq t+1$ basis Pauli matrices. Furthermore, the basis can be computed in time $O(\mathsf{poly}(s))$ and the collection of Pauli matrices can be computed in time $O(2^b \cdot \mathsf{poly}(s))$.*

This lemma is proved after the description of the rest of the algorithm. Given the basis $\mathcal{B} = \{B^{(1)}, \ldots, B^{(b)}\}$ for $b \leq t+1$, define $\gamma(k', k) \stackrel{\text{def}}{=} 1$ if $B^{(k')}$ and $B^{(k)}$ commute and $\stackrel{\text{def}}{=} 0$ otherwise. Observe then the Pauli matrices

$$A^{(k)} \stackrel{\text{def}}{=} \prod_{k'<k} X_{k'}^{\gamma(k',k)} \cdot Z_k \tag{68}$$

observe the same commutation relations as $\mathcal{B}$ do. However, $\mathcal{A} = \{A^{(1)}, \ldots, A^{(b)}\}$ act on at most b qubits. For each Pauli matrix $P^{(\ell)}$ defined as a product of elements from $\mathcal{B}$, let us define $O^{(\ell)}$ as the same product, except using the corresponding matrices from $\mathcal{A}$. Then the spectrum of

$$H' \stackrel{\text{def}}{=} \sum_{\ell=1}^{\leq 2^t} O^{(\ell)} \tag{69}$$

is the same as that of $Q^{(s)}$ from Equation (67). This is because there exists a unitary mapping $\mathcal{A}$ to $\mathcal{B}$ which therefore maps $O^{(\ell)}$ to $P^{(\ell)}$ and likewise maps H to $Q^{(s)}$. Therefore, it suffices to compute the minimum eigenvalue of H'. Here H' is a square matrix of dimension $2^b \times 2^b$ with $b \leq t+1$. Computing H' and its minimum eigenvalue to accuracy δ can be done in time $O(\mathsf{poly}(s)2^{3b}\log(t/\delta))$.

Lastly, notice that a convenient quality of this algorithm is that the basis $\mathcal{B}$ and its size b can be computed in time $O(\mathsf{poly}(s))$. Therefore, one can calculate b and calculate[7] the runtime of the algorithm without running the algorithm itself. ◀

Proof of Lemma 17. We proceed by induction on the gates $g_1, \ldots, g_s$ of U. Initially, the only basis matrix is $B^{(0,1)} = Z_1$ and $P^{(0,1)} = B^{(0,1)}$. Then in the inductive step, we assume a basis of $\{B^{(k,\lambda)}\}$ and a collection of Pauli matrices $P^{(k,\ell)}$ such that each Pauli is expressible as a product of the basis matrices. When gate g_k is a Clifford, we define $B^{(k-1,\lambda)} \stackrel{\text{def}}{=} g_k B^{(k-1,\lambda)} g_k^\dagger$. Conveniently, $P^{(k,\ell)} = g_k P^{(k-1,\ell)} g_k^\dagger$ is the product of the corresponding set of new basis matrices.

In the case that $g_k = R(\theta)$ which (without loss of generality) acts on the first qubit, we first transform the basis $\{B^{(k-1,\lambda)}\}$ by multiplying basis terms we ensure that at most 2 basis terms act non-trivially on the first qubit. The terms $P^{(k-1,\ell)}$ can be adjusted in polynomial-time to reflect the new basis. If there are no basis terms acting non-trivially or one basis term acting as Z, then we set $P^{(k,\ell)}$ equal to $P^{(k-1,\ell)}$. In the case that the one basis term (without loss of generality, $B^{(k-1,1)}$ acts as X, then we introduce a new basis term defined as $B^{(k,\text{new})} \stackrel{\text{def}}{=} B^{(k,1)} \cdot XY$. Any Pauli term $P^{(k-1,\ell)}$ involving $B^{(k-1,1)}$ after commuting by g_k now a linear combination of said term and $B^{(k,\lambda_{\text{new}})}$ according to Equation (64). A similar argument holds when the one basis $B^{(k-1,1)}$ term acts as Y. In the case that two basis terms act non-trivially on the first qubit, we can also enforce that one of the basis terms acts as Z and the other acts as either X or Y. Then, a similar argument enforces that it suffices to introduce a single additional basis term.

Therefore, at the end of the induction, the total basis has size at most $t+1$ and $UZ_1U^\dagger$ can be expressed as a linear combination of $\leq 2^t$ Pauli terms. ◀

[7] Namely, this is to check in time $O(\mathsf{poly}(s))$ if $b \ll t$ to see if this algorithm will be more efficient at computing H' than the stabilizer-rank of magic states algorithm (Theorem 2) derived from [23].