Tight Bounds for Quantum Phase Estimation and Related Problems

Nikhil S. Mande[1] and Ronald de Wolf[2]

[1]University of Liverpool, UK
[2]QuSoft, CWI and University of Amsterdam, the Netherlands

Phase estimation, due to Kitaev [arXiv'95], is one of the most fundamental subroutines in quantum computing, used in Shor's factoring algorithm, optimization algorithms, quantum chemistry algorithms, and many others. In the basic scenario, one is given black-box access to a unitary U, and an eigenstate $|\psi\rangle$ of U with unknown eigenvalue $e^{i\theta}$, and the task is to estimate the eigenphase θ within $\pm\delta$, with high probability. The repeated application of U and U^{-1} is typically the most expensive part of phase estimation algorithms, so for us the *cost* of an algorithm will be that number of applications.

Motivated by the "guided local Hamiltonian problem" from quantum chemistry, we tightly characterize the cost of several variants of phase estimation where we are no longer given an arbitrary eigenstate, but are required to estimate the *maximum* eigenphase of U, aided by *advice* in the form of copies of a state (or a unitary preparing that state) that is promised to have at least a certain overlap γ with the top eigenspace. We give algorithms and matching lower bounds (up to logarithmic factors) for all ranges of parameters. We show a crossover point below which advice is not helpful: $o(1/\gamma^2)$ copies of the advice state (or $o(1/\gamma)$ applications of an advice-preparing unitary) are not significantly better than having no advice at all. We also show that having much more advice (more than $O(1/\gamma^2)$ copies or more than $O(1/\gamma)$ applications of the advice-preparing unitary) does not significantly reduce cost, and neither does knowledge of the eigenbasis of U. As an immediate consequence of a key technical component of our proof, we obtain a lower bound on the complexity of the Unitary recurrence time problem, matching an upper bound of She and Yuen [ITCS'23] and resolving one of their open questions.

Lastly, we study how efficiently one can reduce the error probability in the basic phase-estimation scenario. We show that a phase-estimation algorithm with precision δ and error probability ε has cost $\Omega\left(\frac{1}{\delta}\log\frac{1}{\varepsilon}\right)$, matching the obvious way to amplify the basic constant-error-probability phase estimation algorithm. This contrasts with some other scenarios in quantum computing (e.g., search) where error-probability reduction costs only a factor $O(\sqrt{\log(1/\varepsilon)})$. Our lower bound uses a variant of the polynomial method with trigonometric polynomials.

Nikhil S. Mande: mande@liverpool.ac.uk, https://mande-nikhil.github.io/, Part of this work was done while the author was a postdoc at QuSoft and CWI, Amsterdam.

Ronald de Wolf: rdewolf@cwi.nl, https://homepages.cwi.nl/ rdewolf/, Partially supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037).

1 Introduction

1.1 Phase estimation

Kitaev [Kit95] gave an elegant and efficient quantum algorithm for the task of *phase estimation* nearly 30 years ago. The task is easy to state: given black-box access to a unitary and an eigenstate of it, estimate the phase of the associated eigenvalue. Roughly speaking, the standard algorithm for this task sets up a superposition involving many different powers of the unitary to extract many different powers of the eigenvalue, and then uses a quantum Fourier transform to turn that into an estimate of the eigenphase.[1] Many of the most prominent quantum algorithms can either be phrased as phase estimation, or use phase estimation as a crucial subroutine. Some examples are Shor's period-finding algorithm [Sho97] as presented in [CEMM98]; approximate counting [BHMT02] can be done using phase estimation on the unitary of one iteration of Grover's search algorithm [Gro96], which also recovers the $O(\sqrt{N})$ complexity for searching an N-element unordered search space; the HHL algorithm for solving linear systems of equations estimates eigenvalues in order to invert them [HHL09]. Applications of phase estimation in quantum chemistry are also very prominent, as discussed below.

More precisely, in the task of phase estimation, we are given black-box access to an N-dimensional unitary U (and a controlled version thereof) and a state $|\psi\rangle$ that satisfies $U|\psi\rangle = e^{i\theta}|\psi\rangle$. Our goal is to output (with probability at least $2/3$) a $\tilde{\theta} \in [0, 2\pi)$ such that $|\tilde{\theta} - \theta|$ is at most δ in $\mathbb{R}$ mod 2π. In the basic scenario we are given access to one copy of $|\psi\rangle$, and are allowed to apply U and its inverse. Since the repeated applications of U and U^{-1} are typically the most expensive parts of algorithms for phase estimation, the *cost* we wish to minimize is the number of applications of U and U^{-1}. We are additionally allowed to apply arbitrary unitaries that do not depend on U, at no cost. Kitaev's algorithm has cost $O(1/\delta)$ in this measure.

1.2 Phase estimation with advice

One of the core problems in quantum chemistry is the following: given a classical description of some Hamiltonian H (for instance an "electronic structure" Hamiltonian in the form of a small number of local terms), estimate its *ground-state energy*, which is its smallest eigenvalue. If H is normalized such that its eigenvalues are all in $[0, 2\pi)$ and we define the unitary $U = e^{iH}$ (which has the same eigenvectors as H, with an eigenvalue λ of H becoming the eigenvalue $e^{i\lambda}$ for U), then finding the ground state energy of H is equivalent to finding the smallest eigenphase of U. If we are additionally given a *ground state* $|\psi\rangle$ (i.e., an eigenstate corresponding to the smallest eigenphase), then phase estimation is tailor-made to estimate the ground-state energy. However, in quantum chemistry it is typically hard to prepare the ground state of H, or even something close to it. What can sometimes be done is a relatively cheap preparation of some quantum "advice state" that has some non-negligible "overlap" γ with the ground space, for instance the "Hartree-Fock state". This idea is also behind the Variational Quantum Eigensolver (VQE) [PMS+14]. In the complexity-theoretic context, this problem of ground-state energy estimation for a local Hamiltonian given an advice state, is known as the "guided local Hamiltonian problem", and has received quite some attention recently [GG23, CFG+23, WFC24] because of its connections with quantum chemistry as well as deep complexity questions such as the

[1] An added advantage of the standard algorithm for phase estimation is that it can also work with a quantum Fourier transform that is correct on average rather than in the worst case [LW22]. However, there are also approaches to phase estimation that avoid the QFT altogether, see e.g. [Ral21].

PCP conjecture. These complexity-theoretic results typically focus on completeness for BQP or QMA or QCMA of certain special cases of the guided local Hamiltonian problem, and don't care about polynomial overheads of the cost (number of uses of U and U^{-1}) in the number of qubits $\log N$ or in the parameters δ and γ. In contrast, we care here about getting essentially optimal bounds on the cost of phase estimation in various scenarios.

To be more precise, suppose our input unitary is $U = \sum_{j=0}^{N-1} e^{i\theta_j}|u_j\rangle\langle u_j|$ with each $\theta_j \in [0, 2\pi - 2\delta)$.[2] Let $\theta_{\max} = \max_{j\in\{0,1,\ldots,N-1\}} \theta_j$ denote the maximum eigenphase, and let S denote the space spanned by all eigenstates with eigenphase $\theta_{\max}$, i.e., the "top eigenspace". Advice is given in the form of a state $|\alpha\rangle$ whose projection on S has squared norm at least γ^2: $\|P_S|\alpha\rangle\|^2 \geq \gamma^2$. Note that if S is spanned by a single eigenstate $|u_{\max}\rangle$, then this condition is the same as $|\langle\alpha|u_{\max}\rangle| \geq \gamma$, which is why we call γ the *overlap* of the advice state with the target eigenspace. The task $\mathsf{maxQPE}_{N,\delta}$ is to output, with probability at least 2/3, a δ-precise (in $\mathbb{R}$ mod 2π) estimate of $\theta_{\max}$.[3]

We will distinguish between the setting where the advice is given in the form of a number of copies of the advice state $|\alpha\rangle$, or the potentially more powerful setting where we can apply (multiple times) a *unitary* A that prepares $|\alpha\rangle$ from some free-to-prepare initial state, say $|0\rangle$. We would have such a unitary A for instance if we have a procedure to prepare $|\alpha\rangle$ ourselves in the lab. We can also distinguish between the situation where the eigenbasis $|u_0\rangle, \ldots, |u_N\rangle$ of U is known (say, the computational basis where $|u_j\rangle = |j\rangle$) and the potentially harder situation where the eigenbasis is unknown. These two binary distinctions give us four different settings. For each of these settings we determine essentially optimal bounds on the cost of phase estimation, summarized in Table 1 (see Section 2 for the formal setup).

Let us highlight some interesting consequences of our results. First, a little bit of advice is no better than no advice at all: the upper bounds in the odd-numbered rows of Table 1 are actually obtained by algorithms that don't use the given advice ($o(1/\gamma^2)$ copies of $|\alpha\rangle$ or $o(1/\gamma)$ applications of A and A^{-1}) at all, yet their costs essentially match our lower bounds for algorithms that use the given advice.[4]

A second interesting consequence is that too much advice is no better than a moderate amount of advice: the upper bounds in Rows 2 and 4 use $O(1/\gamma^2)$ advice states, and the upper bounds in Rows 6 and 8 use $O(1/\gamma)$ advice unitaries, and using more advice does not reduce the cost further.

Thirdly, it turns out that knowledge of the eigenbasis of U doesn't really help in reducing the cost: the costs in row 1 and row 3 are the same, and similarly for rows 2 vs. 4, 5 vs. 7 and 6 vs. 8.

[2] By multiplying U with $e^{i\delta}$, we can equivalently assume that all eigenphases are in the interval $[\delta, 2\pi-\delta)$, hence $\notin [-\delta, \delta] \mod 2\pi$. In an earlier version of this paper we did not include this promise in the problem definition. However, then the algorithm of Lemma 4.3 runs into trouble: if $\theta_{\max}$ is, say, π, while all other eigenphases are just above 0, then the unitary V in that proof generates a state with non-negligible weight on values x_k that are just below 2π (since those are good approximations of the just-above-0 eigenphases mod 2π), and the generalized maximum-finding will zoom in on those values, ending up with some number just below 2π as its (very wrong) guess for $\theta_{\max}$. We thank Han-Hsuan Lin (personal communication) for alerting us to this issue.

[3] It doesn't really matter, but we focus on the *maximum* rather than minimum eigenphase of U because eigenphase 0 (i.e., eigenvalue 1) is a natural baseline, and we are looking for the eigenphase furthest away from this baseline.

[4] The proofs in Section 3 yield the same asymptotic lower bounds for algorithms with access to at most c/γ^2 advice states for Theorem 3.2, Rows 1 and 3 of Table 1, and for algorithms with access to at most c/γ advice unitaries for Rows 5 and 7 of Table 1, where c is a suitably small universal constant. We chose to use $o(\cdot)$ here to avoid clutter.

Row	Basis	Access to advice	Number of accesses	Upper bound	Lower bound
1	known	state	$o\left(\frac{1}{\gamma^2}\right)$	$\widetilde{O}\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 4.4	$\Omega\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 3.3
2	known	state	$\Omega\left(\frac{1}{\gamma^2}\right)$	$\widetilde{O}\left(\frac{1}{\gamma\delta}\right)$, Lemma 4.6	$\Omega\left(\frac{1}{\gamma\delta}\right)$, Lemma 3.4
3	unknown	state	$o\left(\frac{1}{\gamma^2}\right)$	$\widetilde{O}\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 4.4	$\Omega\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 3.3
4	unknown	state	$\Omega\left(\frac{1}{\gamma^2}\right)$	$\widetilde{O}\left(\frac{1}{\gamma\delta}\right)$, Lemma 4.6	$\Omega\left(\frac{1}{\gamma\delta}\right)$, Lemma 3.4
5	known	unitary	$o\left(\frac{1}{\gamma}\right)$	$\widetilde{O}\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 4.4	$\Omega\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 3.5
6	known	unitary	$\Omega\left(\frac{1}{\gamma}\right)$	$\widetilde{O}\left(\frac{1}{\gamma\delta}\right)$, Lemma 4.5	$\Omega\left(\frac{1}{\gamma\delta}\right)$, Lemma 3.6
7	unknown	unitary	$o\left(\frac{1}{\gamma}\right)$	$\widetilde{O}\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 4.4	$\Omega\left(\frac{\sqrt{N}}{\delta}\right)$, Lemma 3.5
8	unknown	unitary	$\Omega\left(\frac{1}{\gamma}\right)$	$\widetilde{O}\left(\frac{1}{\gamma\delta}\right)$, Lemma 4.5	$\Omega\left(\frac{1}{\gamma\delta}\right)$, Lemma 3.6

Table 1: Our results for the cost of $\mathsf{maxQPE}_{N,\delta}$. We assume $\gamma \geq 1/\sqrt{N}$ since a random state has overlap $1/\sqrt{N}$ with the target eigenspace with high probability, and such a state can be prepared at no cost. The 'Basis' column indicates whether the eigenbasis of U is known; 'Access to advice' indicates whether we get copies of the advice state or a unitary to prepare it; 'Number of accesses' refers to the number of accesses to advice that we have. The last two columns show our bounds with references to the lemmas where they are stated and proved. The $\widetilde{O}(\cdot)$ in the upper-bound column hides a factor $\log N$ for the odd-numbered rows, and $\log(1/\gamma)$ for the even-numbered rows. The lower bounds assume $\delta \in (0, 1)$.

Our techniques. Our upper bounds use the subroutine of *generalized maximum-finding* of van Apeldoorn, Gilyén, Gribling, and de Wolf [AGGW20] which allows us to find maximum values in the second register of a two-register superposition even when the first of these two registers has an unknown basis. We derive the upper bound of row 4 from the upper bound of row 8 by using roughly $1/\gamma$ copies of $|\alpha\rangle$ to simulate one reflection around the state $|\alpha\rangle = A|0\rangle$, using the techniques of Lloyd, Mohseni, and Rebentrost [LMR13].

Our lower bounds follow via reductions from a fractional version of the Boolean OR function with advice. We show a lower bound for this by a simple modification of the adversary method [Amb02], taking into account the input-dependent advice in the initial state, and the fact that applications of U can be made to correspond to "fractional" queries.

Gate-complexity of our algorithms. We stated the *cost* (i.e., number of applications of U and U^{-1}) of our algorithms here in the upper-bound column of Table 1, not the overall time complexity. However, it is easy to verify that the gate-complexities of our algorithms are only larger than the cost by log-factors in all cases except rows 2 and 4: our algorithms use two main subroutines, which have only small overheads in gate-complexity, namely basic quantum phase estimation [Kit95] and generalized maximum-finding [AGGW20]. In contrast, our upper bound for row 4 (and hence for row 2) additionally uses [LMR13] $O(1/\gamma)$ times to implement a reflection about the state $|\alpha\rangle$, consuming $O(1/\gamma)$ copies of that state for each reflection. One such reflection then has cost 0 but gate-complexity $\widetilde{O}(1/\gamma)$, meaning the overall gate-complexity of our algorithm for row 4 is $\widetilde{O}(1/\gamma\delta + 1/\gamma^2)$ rather than $\widetilde{O}(1/\gamma\delta)$, which makes a difference if $\gamma \ll \delta$.

Comparison with related work. Some of the results in our table were already (partially) known. A cost-$\widetilde{O}(\sqrt{N}/\delta)$ algorithm for the adviceless setting with unknown eigenbasis (implying the upper bounds of rows 1, 3, 5, 7) was originally due to Poulin and Wocjan [PW09], and subsequently improved in the log-factors by van Apeldoorn et al. [AGGW20]; the latter algorithm is basically our proof of Lemma 4.4. Lin and Tong [LT20] (improving upon [GTC19]) studied the situation with an advice-preparing unitary. Their setting is

slightly different from ours, they focus on preparing the ground state[5] of a given Hamiltonian without a known bound on its spectrum, but [LT20, Theorem 8] implies a cost-$O(\log(1/\gamma)\log(1/\delta)\log\log(1/\delta)/\gamma\delta)$ algorithm for our row 8. Their follow-up paper [LT22] further reduces the number of auxiliary qubits with a view to near-term implementation, but does not reduce the cost further. Our cost-$O(\log(1/\gamma)/\gamma\delta)$ algorithm is slightly better in the log-factors than theirs, and uses quite different techniques ([LT20] uses quantum singular value transformation [GSLW19]).

On the lower-bound side, $\Omega(1/\delta)$ for the cost of phase estimation has long been known to hold when the success probability is required to be a constant, this follows for instance from the approximate counting lower bound of Nayak and Wu [NW99] (see also [Bes05]). Lin and Tong [LT20, Theorem 10] proved lower bounds of $\Omega(1/\gamma)$ and $\Omega(1/\delta)$ on the cost for the setting with known eigenbasis and advice unitary (our row 6, and hence also row 8). This is subsumed by our stronger (and essentially optimal) $\Omega(1/\gamma\delta)$ lower bound in row 6. As far as we are aware, ours is the first paper to systematically tie together these different results and to complete the table with tight upper and lower bounds for the cost in all 8 cases.

Let us also mention some recent work that is not directly covered by our results. First, lower bounds for the slightly unusual small-success-probability regime were recently studied by Lin [Lin23]. Second, there has been work to make phase estimation more efficient in the important special case where the unitary $U = e^{iH}$ is induced by a Hamiltonian H given classically as the sum of relatively simple terms, when the cost of phase estimation interacts with the cost of Hamiltonian simulation. See for instance the recent paper by Wan, Berta, and Campbell [WBC22] and references therein.

Application. She and Yuen [SY23, Theorems 1.6 and 1.7] studied the (t,δ)-*Unitary recurrence time problem*, which is to distinguish whether an input unitary U satisfies $U^t = I$ or $\|U^t - I\| \geq \delta$, promised that one of these is the case (see Definition 2.4). They proved non-matching upper and lower bounds for the cost of quantum algorithms for this problem (see Theorem 2.5 in this paper). As an immediate application of our lower bound for fractional OR with advice, we also obtain improved lower bounds for the unitary recurrence time problem that match the upper bound of She and Yuen and answer one of their open problems [SY23, Section 2].

Theorem 1.1 (Lower bound for Unitary recurrence time)**.** *Any quantum algorithm solving the (t,δ)-recurrence time problem for N-dimensional unitaries has cost $\Omega(t\sqrt{N}/\delta)$.*

Interestingly, our lower bound uses the adversary method as opposed to She and Yuen's usage of the polynomial method.

1.3 Phase estimation with small error probability

For our results in this subsection we revert to the original scenario of phase estimation, where an algorithm is given the actual eigenstate as input and the goal is to estimate its eigenphase θ. However, we now consider the regime where we want small error probability ε rather than constant error probability 1/3. Let $\mathsf{QPE}_{N,\delta,\varepsilon}$ denote the task of computing, with error probability $\leq \varepsilon$, a δ-approximation of θ. Repeating Kitaev's $O(1/\delta)$-cost phase estimation algorithm [Kit95] $O(\log(1/\varepsilon))$ times and taking the median of the answers, we have the following ε-dependent upper bound.

[5] Because generalized maximum-finding (Lemma 4.1) actually outputs a state in addition to an estimate, our algorithms can be modified to also output a state that is close to the top eigenspace of U.

Theorem 1.2 (Kitaev + standard error-reduction)**.** *For all integers $N \geq 2$, all $\varepsilon \in (0, 1/2)$, and all $\delta \in [0, 2\pi)$, there exists an algorithm that solves $\mathsf{QPE}_{N,\delta,\varepsilon}$ with cost $O\left(\frac{1}{\delta}\log\frac{1}{\varepsilon}\right)$.*

Grover's algorithm [Gro96] can compute the OR_N function with error probability $\leq 1/3$ using $O(\sqrt{N})$ queries to its N input bits. Interestingly, there exists an ε-error quantum algorithm for OR_N with only $O(\sqrt{N\log(1/\varepsilon)})$ queries [BCWZ99], which is asymptotically optimal. Similarly one can reduce error from $1/3$ to ε for all symmetric Boolean functions at the expense of only a factor $\sqrt{\log(1/\varepsilon)}$ in the query complexity [Wol08]. This is a speed-up over the naive $O(\log(1/\varepsilon))$ multiplicative overhead. Since optimal quantum algorithms with error probability $1/3$ for OR_N and for all symmetric functions can be derived from phase estimation, one may ask if one can achieve such an efficient error-reduction for quantum phase estimation as well: is there an algorithm for $\mathsf{QPE}_{N,\delta,\varepsilon}$ of cost $O\left(\frac{1}{\delta}\sqrt{\log(1/\varepsilon)}\right)$? We answer this in the negative, showing Theorem 1.2 is optimal.

Theorem 1.3. *For integers $N \geq 2$ and $\varepsilon \in (0, 1/2), \delta \in (0, 1)$, every algorithm that solves $\mathsf{QPE}_{N,\delta,\varepsilon}$ has cost $\Omega\left(\frac{1}{\delta}\log\frac{1}{\varepsilon}\right)$.*

In particular, this means that the optimal complexity of OR_N with small error probability ε of [BCWZ99] can*not* be derived from a phase estimation routine, in contrast to the case of OR_N (and search) with constant error probability. To show Theorem 1.3 we first argue that a cost-C algorithm for $\mathsf{QPE}_{N,\delta,\varepsilon}$ gives us a cost-C algorithm that distinguishes $U = I$ versus $U = I - (1 - e^{i\theta})|0\rangle\langle 0|$ where $\theta \notin [-3\delta, 3\delta] \bmod 2\pi$. We then note that the acceptance probability of such an algorithm can be written as a degree-$2C$ *trigonometric* polynomial in θ, and invoke a known upper bound on the growth of such trigonometric polynomials in order to lower bound their degree.

2 Preliminaries

We state the required preliminaries in this section. All logarithms are taken base 2. For a positive integer N, $U(N)$ denotes the space of N-dimensional unitaries, and I denotes the N-dimensional Identity matrix (we drop the subscript if the dimension is clear from context).

For a positive integer $N \geq 2$ and a value $\theta \in [0, 2\pi)$, define the N-dimensional unitary U_θ as $U_\theta = I - (1 - e^{i\theta})|0\rangle\langle 0|$. In other words, U_θ is the diagonal matrix with all 1's except the first entry, which is $e^{i\theta}$. For an integer $j \in \{0, 1, \ldots, N-1\}$ and $\delta \in [0, 2\pi)$, define $M_{j,\delta} = I - (1 - e^{i\delta})|j\rangle\langle j|$.

2.1 Model of computation

Here we give a description of our model of computation for all tasks considered in this paper. All problems considered in this paper have the following properties:

- **Input:** An N-dimensional unitary U. We have access to the input as described below.
- **State space:** The state space of an algorithm comprises two registers: the first register is N-dimensional, and the second register is an arbitrarily large workspace.
- **Access to input and allowed operations:** An algorithm $\mathcal{A}$ may apply U and U^{-1} to the first register (possibly controlled by a qubit in the second register), and unitaries independent of U to the whole space. It performs a POVM at the end to determine the classical output.

- **Cost of an algorithm:** Total number of applications of U and U^{-1} (or controlled versions thereof).

Depending on the specific problem under consideration, the following properties are variable.

- **Initial state:** The initial state is assumed to be $|0\rangle|0\rangle$ unless mentioned otherwise.
- **Input promise:** The subset of the N-dimensional unitary group $U(N)$ from which the input is taken (possibly the full set).
- **Output:** The output requirement.
- **Advice:** We may be given access to a specific number of "advice states" $|\alpha\rangle$, or access to a specific number of applications of a unitary A that prepares an advice state (e.g., $A|0\rangle = |\alpha\rangle$).

2.2 Problems of interest

We list our problems of interest here. All problems fit in the framework of the previous subsection, so we skip descriptions of the input, access to the input and allowed operations, and the workspace.

Definition 2.1 (Phase Estimation)**.** *Let $N \geq 2$ be an integer and $\varepsilon, \delta > 0$. The task $\mathsf{QPE}_{N,\delta,\varepsilon}$ is:*

- ***Advice:*** *We are given a single state $|\psi\rangle$ (in other words, our starting state is $|\psi\rangle|0\rangle$) with the promise that $U|\psi\rangle = e^{i\theta}|\psi\rangle$ for some unknown $\theta \in [0, 2\pi)$.*
- ***Output:*** *With probability at least $1 - \varepsilon$, output $\tilde{\theta} \in [0, 2\pi)$ such that $|\tilde{\theta} - \theta| \leq \delta \bmod 2\pi$.*

The following is a decision version of a special case of phase estimation, where $|\psi\rangle = |0\rangle$:

Definition 2.2. *Let $N \geq 2$ be an integer, $\varepsilon, \delta \in (0, 1)$. The task $\mathsf{dist}_{N,\delta,\varepsilon}$ is:*

- ***Input promise:*** *$U \in \{I, \{U_\theta : \theta \notin [-\delta, \delta] \mod 2\pi\}\}$.*
- ***Output:*** *With probability at least $1 - \varepsilon$, output 1 if $U = I$ and output 0 if $U \neq I$.*

We next define the natural variant of phase estimation that we consider when an algorithm need not be given a state from the target eigenspace but only a state $|\alpha\rangle$ that has non-negligible overlap with that eigenspace.

Definition 2.3 (Maximum phase estimation)**.** *Let $N \geq 2$ be an integer and $\delta > 0$. The task $\mathsf{maxQPE}_{N,\delta}$ is:*

- ***Input promise:*** *We consider two cases: one where the eigenbasis of U is known, and the other where it is unknown. In the former case, we may assume $U = \sum_{j=0}^{N-1} e^{i\theta_j}|j\rangle\langle j|$ for $\theta_j \in [0, 2\pi)$. Define $\theta_{\max} = \max_{j \in \{0,1,\ldots,N-1\}} \theta_j$. We are promised that $\theta_{\max} \in [0, 2\pi - 2\delta)$.*
- ***Advice:*** *We consider two cases:*
  - *In one case we are given access to advice in the form of a state $|\alpha\rangle$ such that $\|P_S|\alpha\rangle\|^2 \geq \gamma^2$, where P_S denotes the projection on S, the space of all eigenstates with eigenphase $\theta_{\max}$. If S is spanned by one $|u_{\max}\rangle$, this requirement is the same as $|\langle\alpha|u_{\max}\rangle| \geq \gamma$.*

– *In the other case, we have black-box access to a unitary A that prepares such a state $|\alpha\rangle$. We can apply A and A^{-1}. As before, γ is the* overlap *of $|\alpha\rangle$ with the target eigenspace.*

- ***Number of accesses to advice:*** *We either have 'few' accesses to advice as defined above ($o(1/\gamma^2)$ advice states or $o(1/\gamma)$ advice unitaries), or 'many' accesses to advice ($\Omega(1/\gamma^2)$ advice states or $\Omega(1/\gamma)$ advice unitaries).*
- ***Output:*** *With probability at least $2/3$, output a value in $[\theta_{\max} - \delta, \theta_{\max} + \delta] \bmod 2\pi$.*

Definition 2.4 (Unitary recurrence time, [SY23, Definition 1.5])**.** *For integers $N \geq 2, t \geq 1$ and $\delta \in (0,1)$, the (t,δ)-recurrence time problem is:*

- ***Input promise:*** *Either $U = I$, or $\|U^t - I\| \geq \delta$ in spectral norm.*
- ***Output:*** *With probability at least $2/3$: output 1 if $U = I$, and 0 otherwise.*

The following are the non-matching upper and lower bounds for this problem of She and Yuen [SY23] (which we improve upon in our Theorem 1.1).

Theorem 2.5 ([SY23, Theorems 1.6 and 1.7])**.** *Let $\delta \leq \frac{1}{2\pi}$. Every quantum algorithm solving the (t,δ)-recurrence time problem for d-dimensional unitaries has cost $\Omega\left(\max\left(t/\delta, \sqrt{d}\right)\right)$. The (t,δ)-recurrence time problem can be solved with cost $O(t\sqrt{d}/\delta)$.*

2.3 Trigonometric polynomials and their growth

Definition 2.6 (Trigonometric Polynomials)**.** *A function $p : \mathbb{R} \to \mathbb{C}$ is said to be a* trigonometric polynomial *of degree d if there exist complex numbers $\{a_k : k \in \{-d, \ldots, d\}\}$ such that for all $\theta \in \mathbb{R}$,*

$$p(\theta) = \sum_{k=-d}^{d} a_k e^{ik\theta}.$$

We will use the following property of low-degree trigonometric polynomials.

Theorem 2.7 ([BE95, Theorem 5.1.2])**.** *Let t be a degree-n real-valued trigonometric polynomial and $s \in (0, \pi/2]$ be such that $\mu(\{\theta \in [-\pi, \pi) : |t(\theta)| \leq 1\}) \geq 2\pi - s$, where μ denotes the Lebesgue measure on $\mathbb{R}$. Then, $\sup_{x \in \mathbb{R}} |t(x)| \leq \exp(4ns)$.*

3 Lower bounds for maximum phase estimation and Unitary recurrence time

In this section we show lower bounds on the quantum complexity of maximum phase estimation obtained by varying all its parameters (see Section 2.1 and Definition 2.3). In this section and the next, we refer to the row numbers of Table 1 when stating and proving our bounds.

Recall that for an integer $j \in \{0, 1, \ldots, N-1\}$ and $\delta \in [0, 2\pi)$ we define $M_{j,\delta} = I - (1 - e^{i\delta})|j\rangle\langle j|$. Our lower bounds will be via reduction from the following "Fractional OR with advice" problem, which fits in the framework of the model described in Section 2.1.

Definition 3.1 (Fractional OR with advice)**.** *Let $N \geq 2$ be integer, $\delta > 0$. The task* $\mathsf{frOR}_{N,\delta,t}$ *is:*

- ***Input promise:*** *$U \in \{I, \{M_{j,\delta} : j \in \{1, 2, \ldots, N-1\}\}\}$.*

- ***Advice:*** *When $U = I$ we are given t copies of $|0\rangle$ as advice. When $U = M_{j,\delta}$, we are given t copies of the state $\gamma|j\rangle + \sqrt{1-\gamma^2}|0\rangle$, i.e., part of our starting state is $(\gamma|j\rangle + \sqrt{1-\gamma^2}|0\rangle)^{\otimes t}$.*
- ***Output:*** *With probability at least $2/3$, output 1 if $U = I$ and output 0 if $U \neq I$.*

We first show a lower bound on the cost of computing $\mathsf{frOR}_{N,\delta,t}$ when $t = o(1/\gamma^2)$. All of our lower bounds in Table 1 as well as our lower bound for the Unitary recurrence time problem will use this lower bound. The proof (given in Appendix A) follows along the same lines as Ambainis' adversary lower bound [Amb02, Theorem 4.1] of $\Omega(\sqrt{N})$ queries for the N-bit Search problem, but now we additionally take into account the initial advice states and the fact that our input unitaries are only *fractional* versions of phase queries.

Theorem 3.2. *For an integer $N \geq 2$, real numbers $\gamma \geq 1/\sqrt{N}$, $\delta \in (0,1)$ and $t = o(1/\gamma^2)$, every algorithm solving $\mathsf{frOR}_{N,\delta,t}$ has cost $\Omega(\sqrt{N}/\delta)$.*

In the following four lower-bound lemmas we assume $\delta \in (0,1)$ as well.

Lemma 3.3 (Lower bound for Rows 1,3)**.** *Row 1 (and hence Row 3) has a lower bound of $\Omega(\sqrt{N}/\delta)$.*

Proof. A cost-C algorithm $\mathcal{A}$ for $\mathsf{maxQPE}_{N,\delta}$ with t advice states and known eigenbasis of U immediately yields a cost-C algorithm $\mathcal{A}'$ for $\mathsf{frOR}_{N,3\delta,t}$: run $\mathcal{A}$ on the input unitary, output 1 if the output phase is in $[-\delta,\delta]$ modulo 2π, and output 0 otherwise. When $U = I$, the correctness of $\mathcal{A}$ guarantees that with probability at least $2/3$, the value output by $\mathcal{A}$ is in $[-\delta,\delta]$ mod 2π. When $U = M_{j,3\delta}$, the correctness of $\mathcal{A}$ guarantees that with probability at least $2/3$, the value output by $\mathcal{A}$ is in $[2\delta,4\delta]$. For $\delta \in (0,1)$, we have $[-\delta,\delta]$ mod $2\pi \cap [2\delta,4\delta]$ mod $2\pi = \emptyset$. Thus, $\mathcal{A}'$ solves $\mathsf{frOR}_{N,3\delta,t}$ and has cost C. Theorem 3.2 yields the bound $C = \Omega\left(\sqrt{N}/\delta\right)$ when $t = o(1/\gamma^2)$, giving the desired result. □

Lemma 3.4 (Lower bound for Rows 2,4)**.** *Row 2 (and hence Row 4) has a lower bound of $\Omega(1/\gamma\delta)$.*

Proof. We prove the required lower bound for $\mathsf{maxQPE}_{N,\delta}$ with inputs satisfying the promise that $U \in \{I_N, \{M_{j,3\delta} : j \in \{1,2,\ldots,1/\gamma^2 - 1\}\}\}$. Because of this assumption, we may take the uniform superposition over the first $1/\gamma^2$ computational basis states as our advice state: the algorithm should work with such an advice state, since it has overlap at least γ with the top eigenspace for each of the possible U. However, an algorithm can prepare such advice states at no cost, so we may assume that the algorithm has no access to advice at all. As in the previous proof, this gives an algorithm of the same cost for $\mathsf{frOR}_{1/\gamma^2,3\delta,0}$ (ignoring all other dimensions). Theorem 3.2 with $N = 1/\gamma^2$ and $t = 0$ yields the required lower bound of $\Omega(1/\gamma\delta)$. □

Lemma 3.5 (Lower bound for Rows 5,7)**.** *Row 5 (and hence Row 7) has a lower bound of $\Omega(\sqrt{N}/\delta)$.*

Proof. Towards the required lower bound, consider a cost-C algorithm $\mathcal{A}$ solving $\mathsf{maxQPE}_{N,\delta}$ with inputs satisfying the promise $U \in \{I_N, \{M_{j,3\delta} : j \in \{1,2,\ldots,N-1\}\}\}$, and with $t = o(1/\gamma)$ accesses to a unitary that prepares an advice state that has overlap at least γ with the target eigenspace. We want to construct an algorithm $\mathcal{A}'$ for $\mathsf{maxQPE}_{N,\delta}$ with the same promised inputs that uses *no* advice, and with cost not much larger than that of $\mathcal{A}$. Note that we may assume $\gamma = o(1)$, since otherwise $t = 0$, so then $\mathcal{A}$ itself already uses no advice.

We first show how an algorithm can itself implement a good-enough advice unitary A quite cheaply. Assuming without loss of generality that $k = \pi/(3\delta)$ is an integer, U^k is actually a "phase query": if $U = M_{j,3\delta}$, then U^k is the diagonal matrix with $(e^{i3\delta})^k = e^{i\pi} = -1$ in the jth entry and 1s elsewhere; and if $U = I$ then $U^k = I$. Thus A can start by mapping $|0\rangle$ to a uniform superposition over all indices, and then use Grover's algorithm with U^k as the phase-query operator to amplify the amplitude of $|j\rangle$ to $\geq \gamma$. We know that $O(\gamma\sqrt{N})$ "Grover iterations" suffice for this (see, for example, [Wol19, Section 7.2] for details). Each Grover iteration would use one phase-query U^k, so the overall cost (number of applications of U and U^{-1}) of this advice unitary is $k \cdot O(\gamma\sqrt{N}) = O(\gamma\sqrt{N}/\delta)$. If $U = I$, the state just remains the uniform superposition that Grover's algorithm starts with.

We now have all components to describe $\mathcal{A}'$: Run $\mathcal{A}$, and whenever $\mathcal{A}$ invokes an advice unitary, use the above A. Since $\mathcal{A}$ uses at most t advice unitaries, the cost of $\mathcal{A}'$ is at most $C + t \cdot O(\gamma\sqrt{N}/\delta)$. Note that $\mathcal{A}'$ uses no advice at all anymore, and solves $\mathsf{maxQPE}_{N,\delta}$ under the promise that the input unitary satisfies $U \in \{I_N, \{M_{j,3\delta} : j \in \{1, 2, \ldots, N-1\}\}\}$. Again, this immediately yields an algorithm of the same cost for $\mathsf{frOR}_{N,3\delta,0}$ as in the previous two proofs. Theorem 3.2 now implies

$$C + O(t\gamma\sqrt{N}/\delta) = \Omega(\sqrt{N}/\delta),$$

and hence $C = \Omega(\sqrt{N}/\delta)$ since $t = o(1/\gamma)$ ($t \leq c/\gamma$ for sufficiently small constant c also suffices). □

Lemma 3.6 (Lower bound for Rows 6,8)**.** *Row 6 (and hence Row 8) has a lower bound of* $\Omega(1/\gamma\delta)$.

Proof. Just as in the proof of Lemma 3.4, we may assume $N = 1/\gamma^2$ by only allowing input unitaries of the form $U \in \{I_N, \{M_{j,3\delta} : j \in \{1, 2, \ldots, 1/\gamma^2 - 1\}\}\}$. With this assumption, we may assume that we have no access to advice (i.e., $t = 0$) since an algorithm can prepare a good-enough advice state at no cost, namely the uniform superposition over all $N = 1/\gamma^2$ basis states. Lemma 3.5 now yields the required lower bound of $\Omega(1/\gamma\delta)$. □

Finally we prove an optimal lower bound for the Unitary recurrence time problem, matching She and Yuen's upper bound (Theorem 2.5), resolving one of their open problems [SY23, Section 2].

Proof of Theorem 1.1. Consider an algorithm $\mathcal{A}$ solving the (t, δ)-recurrence time problem. Restrict to inputs of the form $U \in \left\{I_N, \left\{M_{j,3\delta/t} : j \in \{1, 2, \ldots, N-1\}\right\}\right\}$. When $U = I$ we have $U^t = I$. When $U = M_{j,3\delta/t}$, we have $\|U^t - I\| = |1 - e^{3i\delta}| \geq \delta$ for all $\delta \in [0, 1]$. Thus, $\mathcal{A}$ solves $\mathsf{frOR}_{N,3\delta/t,0}$. Theorem 3.2 yields the required lower bound of $\Omega(t\sqrt{N}/\delta)$. □

4 Upper bounds for maximum phase estimation

In this section we show upper bounds on the quantum complexity of our 8 variants of maximum phase estimation (see Section 2.1, Definition 2.3 and Table 1). We require the following generalized maximum-finding procedure, adapted from [AGGW20, Theorem 49]; we changed their wording a bit and modified it from minimum-finding to maximum-finding.

Lemma 4.1 (Generalized maximum-finding [AGGW20, Theorem 49])**.** *There exists a quantum algorithm* $\mathcal{M}$ *and constant* $C > 0$ *such that the following holds. Suppose we have a* q*-qubit unitary* V *such that*

$$V|0\rangle = \sum_{k=0}^{K-1} |\psi_k\rangle|x_k\rangle,$$

where $x_0 > x_1 > \cdots > x_{K-1}$ are distinct real numbers (written down in finite precision), and the $|\psi_k\rangle$ are unnormalized states. Let X be the random variable obtained if we were to measure the last register, so $\Pr[X = x_k] = \||\psi_k\rangle\|^2$. Let $M \geq C/\sqrt{\Pr[X \geq x_j]}$ for some j. Then $\mathcal{M}$ uses at most M applications of V and V^{-1}, and $O(qM)$ other gates, and outputs a state $|\psi_i\rangle|x_i\rangle$ (normalized) such that $x_i \geq x_j$ with probability at least $3/4$ (in particular, if $j = 0$ then $\mathcal{M}$ outputs the maximum).

Remark 4.2. *It may be verified by going through [AGGW20, Lemma 48 & Theorem 49] that the only applications of V and V^{-1} used by $\mathcal{M}$ are to prepare $V|0\rangle$ starting from $|0\rangle$, and to reflect about the state $V|0\rangle$.*

We can use generalized maximum-finding to approximate the largest eigenphase starting from the ability to prepare a superposition of eigenstates (possibly with some additional workspace qubits):

Lemma 4.3. *There exists a quantum algorithm $\mathcal{B}$ such that the following holds. Suppose we have an N-dimensional unitary U with (unknown) eigenstates $|u_0\rangle, \ldots, |u_{N-1}\rangle$ and associated eigenphases $\theta_0, \ldots, \theta_{N-1} \in [0, 2\pi - 2\delta)$. Suppose we also have a unitary A such that*

$$A|0\rangle = \sum_{j=0}^{N-1} \alpha_j |u_j\rangle|\phi_j\rangle,$$

where $\sum_{j:\theta_j=\theta_{\max}} |\alpha_j|^2 \geq \gamma^2$ and the $|\phi_j\rangle$ are arbitrary (normalized) states. Then $\mathcal{B}$ uses at most $O(1/\gamma)$ applications of A and A^{-1}, and $O(\log(1/\gamma)/\gamma\delta)$ applications of U and U^{-1}, and with probability at least $2/3$ outputs a number $\theta \in [\theta_{\max} - \delta, \theta_{\max} + \delta]$.

Proof. As mentioned in footnote 2, by multiplying U with the known phase $e^{i\delta}$, we may assume all eigenphases θ_j are in the interval $[\delta, 2\pi - \delta)$. Let $\tilde{V}$ be the unitary that applies phase estimation with unitary U, precision δ, and small error probability η (to be determined later), on the first register of the state $A|0\rangle$, writing the estimates of the phase in a third register. Then

$$\tilde{V}|0\rangle = \sum_{j=0}^{N-1} \alpha_j |u_j\rangle|\phi_j\rangle|\tilde{\theta_j}\rangle,$$

where, for each j, $|\tilde{\theta_j}\rangle$ is a superposition over estimates of θ_j, most of which are δ-close to θ_j.

For the purposes of analysis, we would like to define a "cleaned up" unitary V (very close to $\tilde{V}$) that doesn't have any estimates with error $> \delta$ in $V|0\rangle$ Let $|\tilde{\theta_j}'\rangle$ be the state obtained from $|\tilde{\theta_j}\rangle$ by removing the estimates that are more than δ-far from θ_j, and renormalizing. Because we ran phase estimation with error probability $\leq \eta$, it is easy to show that $\left\||\tilde{\theta_j}'\rangle - |\tilde{\theta_j}\rangle\right\| = O(\sqrt{\eta})$. Then there exists[6] a unitary V such that $\left\|\tilde{V} - V\right\| = O(\sqrt{\eta})$ and

$$V|0\rangle = \sum_{j=0}^{N-1} \alpha_j |u_j\rangle|\phi_j\rangle|\tilde{\theta_j}'\rangle = \sum_{k=0}^{K-1} |\psi_k\rangle|x_k\rangle,$$

where the x_k are the distinct estimates that have support in the last register, and the $|\psi_k\rangle$ are (unnormalized) superpositions of the $|u_j\rangle|\phi_j\rangle$'s that are associated with those estimates.

[6] This is fairly easy to show, see e.g. [CW23, proof of Theorem 2.4 in Appendix A].

Algorithm $\mathcal{B}$ now applies the maximum-finding algorithm $\mathcal{M}$ of Lemma 4.1 with the unitary $\tilde{V}$. Let us first analyze what would happen if $\mathcal{B}$ used the cleaned-up V instead of $\tilde{V}$. Because we assumed that $\theta_j \in [\delta, 2\pi - \delta)$ for all j, and we have removed all estimates of θ_j with error $> \delta$, the largest x_k's in $V|0\rangle$ are good estimates of $\theta_{\max}$ (with no worries about potential "wrapping around" 2π of the estimates). Let X denote the random variable obtained if we measure the last register of $V|0\rangle$, and note that $\Pr[X \geq \theta_{\max} - \delta] \geq \sum_{j:\theta_j=\theta_{\max}} |\alpha_j|^2 \geq \gamma^2$ because all estimates in $V|0\rangle$ have error $\leq \delta$. Hence $\mathcal{B}$ would use $O(1/\gamma)$ applications of V and V^{-1} to find a $\theta \in [\theta_{\max} - \delta, \theta_{\max} + \delta]$ with success probability $\geq 3/4$.

Algorithm $\mathcal{B}$ will actually use $\tilde{V}$ and $\tilde{V}^{-1}$ instead of V and V^{-1}, which (because errors in quantum circuits add at most linearly) incurs an overall error in operator norm of $\leq O(\sqrt{\eta}) \cdot O(1/\gamma)$. Choosing $\eta \ll \gamma^2$, this overall error can be made an arbitrarily small constant. The success probability of the algorithm can drop slightly below $3/4$ now due to this error, but is still $\geq 2/3$.

It remains to analyze the cost of $\mathcal{B}$. Each $\tilde{V}$ uses 1 application of A, and $O(\log(1/\eta)/\delta) = O(\log(1/\gamma)/\delta)$ applications of U and U^{-1} for phase estimation (Theorem 1.2), so $\mathcal{B}$ uses $O(1/\gamma)$ applications of A and A^{-1}, and $O(\log(1/\gamma)/\gamma\delta)$ applications of U and U^{-1} in total. □

The upper bounds for our 8 variants of phase estimation (see Table 1) will all follow from this. To derive these upper bounds, we start with the 4 odd-numbered rows, where it turns out the advice is not actually needed to meet our earlier lower bounds. The next proof is basically the same as [AGGW20, Lemma 50] about estimating the minimal eigenvalue of a Hamiltonian (this improved slightly upon the earlier result of [PW09]; see also [Gil19, Lemma 3.A.4]).

Lemma 4.4 (Upper bound for Rows 1, 3, 5, 7)**.** *There is an algorithm that uses no advice and solves the case in Row 3 (and hence in Rows 1, 5, and 7 as well) with cost* $O(\sqrt{N}\log(N)/\delta)$.

Proof. Let A be the unitary that maps $|0\rangle$ to the maximally entangled state in N dimensions. This state can be written in any orthonormal basis, including the (unknown) eigenbasis of U:

$$A|0\rangle = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} |j\rangle|j\rangle = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} |u_j\rangle|\overline{u_j}\rangle,$$

where $|\overline{u_j}\rangle$ denotes the entry-wise conjugated version of $|u_j\rangle$. Applying Lemma 4.3 with this A, $|\phi_j\rangle = |\overline{u_j}\rangle$, and $\gamma = 1/\sqrt{N}$ gives the result. □

The next two lemmas cover the 4 upper-bound cases where advice states/unitaries *are* helpful.

Lemma 4.5 (Upper bound for Rows 6, 8)**.** *There is a quantum algorithm that uses* $O(1/\gamma)$ *applications of the advice unitary (and its inverse) and solves the case in Row 8 (and hence the case in Row 6 as well) with cost* $O(\log(1/\gamma)/\gamma\delta)$.

Proof. Apply Lemma 4.3 with the unitary A that maps $|0\rangle$ to $|\alpha\rangle$, with empty states $|\phi_j\rangle$. □

Lemma 4.6 (Upper bound for Rows 2, 4)**.** *There is a quantum algorithm that uses* $O(1/\gamma^2)$ *copies of the advice state and solves the case in Row 4 (and in Row 2) with cost* $O(\log(1/\gamma)/\gamma\delta)$.

Proof. We will build upon the algorithm for Row 8 of Lemma 4.5. By Remark 4.2 and the algorithm in Lemma 4.5, its $O(1/\gamma)$ applications of the advice unitary A and its inverse A^{-1} are only used there for two purposes: (1) to prepare a copy of the advice state $A|0\rangle = |\alpha\rangle$, and (2) to reflect about $|\alpha\rangle$. We now want to replace these applications of A by using copies of the advice state. For (1) this is obvious. Assume the algorithm for Row 8 uses (2) at most C/γ times, for some constant C. To implement these reflections, we will invoke the result of Lloyd, Mohseni, and Rebentrost [LMR13] (see also [KLL+17]), who showed that given a number $t > 0$ and $O(t^2/\eta)$ copies of a mixed quantum state ρ, one can implement the unitary $e^{it\rho}$ up to error η (in diamond-norm difference between the intended unitary and the actually-implemented channel). We will use this result with $\rho = |\alpha\rangle\langle\alpha|$, $t = \pi$, $\eta = \gamma/(100C)$, noting that the implemented unitary $e^{i\pi|\alpha\rangle\langle\alpha|} = I - 2|\alpha\rangle\langle\alpha|$ is a reflection about $|\alpha\rangle$ (up to a global minus sign that doesn't matter).

Accordingly, we can implement the $\leq C/\gamma$ reflections used by the algorithm for Row 8 using $O(1/\gamma^2)$ copies of $|\alpha\rangle$, each reflection implemented with error $\leq \eta$. Because errors in quantum circuits add at most linearly, the overall error between the algorithm of Row 8 and our simulation of it (using copies of $|\alpha\rangle$) is at most $\eta \cdot C/\gamma \leq 1/100$. Hence we obtain an algorithm for Row 4 that uses $O(1/\gamma^2)$ copies of $|\alpha\rangle$ and has the same cost $O(\log(1/\gamma)/\gamma\delta)$ as the algorithm of Row 8. □

5 Tight bounds for phase estimation with small error probability

Here we prove our lower bound for quantum algorithms solving phase estimation with precision δ and error probability at most ε, Theorem 1.3, which follows from Claims 5.1 and 5.2 below.

Claim 5.1. *For all integers $N \geq 2$, all $\varepsilon \in (0, 1/2)$ and $\delta \in (0, 1)$, if there is a cost-d algorithm solving $\mathsf{QPE}_{N,\delta,\varepsilon}$, then there is a cost-$d$ algorithm solving $\mathsf{dist}_{N,\delta,\varepsilon}$.*

Proof. Consider an algorithm $\mathcal{A}$ of cost d that solves $\mathsf{QPE}_{N,\delta,\varepsilon}$. We construct below an algorithm $\mathcal{A}'$ of cost d solving $\mathsf{dist}_{N,\delta,\varepsilon}$. Let $U \in U(N)$ be the input. The following is the description of $\mathcal{A}'$:

1. Run $\mathcal{A}$ with inputs U and $|0\rangle$.
2. Output 1 if the output of $\mathcal{A}$ is in $[-\delta, \delta] \mod 2\pi$, and output 0 otherwise.

Clearly $\mathcal{A}'$ is a valid algorithm, as far as access to input and allowed operations are concerned, since its initial state is $|0\rangle$, it applies U, U^{-1}, and some unitaries independent of U, and finally performs a two-outcome projective measurement to determine the output bit. The cost of $\mathcal{A}'$ is d.

The correctness follows along the same lines as the proofs in Section 3. We prove correctness here for completeness. First note that the state $|0\rangle$ is an eigenstate of all $U \in \{I\} \cup \{U_\theta : \theta \notin [-3\delta, 3\delta] \mod 2\pi\}$. When $U = I$, the correctness of $\mathcal{A}$ guarantees that with probability at least $1 - \varepsilon$, the value output by $\mathcal{A}$ is in $[-\delta, \delta] \mod 2\pi$. When $U = U_\theta$, the correctness of $\mathcal{A}$ guarantees that with probability at least $1 - \varepsilon$, the value output by $\mathcal{A}$ is in $[\theta - \delta, \theta + \delta] \mod 2\pi$. For $\theta \notin [-3\delta, 3\delta] \mod 2\pi$ we have $[-\delta, \delta] \mod 2\pi \cap [\theta - \delta, \theta + \delta] \mod 2\pi = \emptyset$ since $\delta < 1 < 2\pi/5$, and hence $\mathcal{A}'$ solves $\mathsf{dist}_{N,\delta,\varepsilon}$. □

We next show a lower bound for the cost of algorithms computing $\mathsf{dist}_{N,\delta,\varepsilon}$.

Claim 5.2. *For all integers $N \geq 2$, all $\varepsilon \in (0, 1/2)$ and $\delta \in (0, 1)$, every algorithm for $\mathsf{dist}_{N,\delta,\varepsilon}$ has cost $\Omega\left(\frac{1}{\delta}\log\frac{1}{\varepsilon}\right)$.*

In order to prove Claim 5.2, we first show that amplitudes of basis states in low-cost algorithms that run on U_θ are low-degree trigonometric polynomials in θ. This is analogous to the fact that amplitudes of basis states in query algorithms for Boolean functions are low-degree (algebraic) polynomials in the input variables [BBC+01, Lemma 4.1], and our proof is inspired by theirs.

Claim 5.3. *Let $t > 0$ be a positive integer and let $\theta \in [0, 2\pi]$. Consider a quantum circuit that has starting state $|0\rangle$, uses an arbitrary number of θ-independent unitaries, uses t applications of controlled-U_θ and controlled-U_θ^{-1} in total, and performs no intermediate measurements. Then the amplitudes of basis states before the final measurement are degree-t trigonometric polynomials in θ.*

Proof. We prove this by induction on t. The claim is clearly true when $t = 0$ since all amplitudes are constants in this case. For the inductive step, suppose the claim is true for $t = d$. Let $|\psi_d\rangle$ denote the state of the circuit just before the application of the $(d+1)$th application of U_θ (the argument for U_θ^{-1} is similar, and we skip it). By the inductive hypothesis, we have

$$|\psi_d\rangle = \sum_w \sum_{b\in\{0,1\}} \sum_{j=0}^{N-1} p_{j,b,w}(\theta)|j\rangle|b\rangle|w\rangle,$$

where the first register is where U_θ and U_θ^{-1} act, the second register is the control qubit, and the last register represents the workspace (i.e., U_θ and U_θ^{-1} do not act on this register), and each $p_{j,b,w}$ is a trigonometric polynomial of degree at most d in θ. For a basis state $|j\rangle|b\rangle|w\rangle$, we have

$$U_\theta|j\rangle|b\rangle|w\rangle = \begin{cases} e^{i\theta}|0\rangle|b\rangle|w\rangle & \text{if } j = 0 \text{ and } b = 1 \\ |j\rangle|b\rangle|w\rangle & \text{otherwise.} \end{cases}$$

In both cases, the amplitudes of the basis states after the application of U_θ are degree-$(d+1)$ trigonometric polynomials in θ. After the last application of U_θ the algorithm will apply an input-independent unitary. The amplitudes after that unitary are linear combinations of the amplitudes before, which won't increase degree. This concludes the inductive step, and hence the theorem. □

Proof of Claim 5.2. Consider a cost-t algorithm $\mathcal{A}'$ solving $\mathsf{dist}_{N,\delta,\varepsilon}$. Claim 5.3 implies that on input U_θ, the amplitudes of the basis states before the final measurement are degree-t trigonometric polynomials in θ. The acceptance-probability polynomial $p : \mathbb{R} \to \mathbb{R}$ given by $p(\theta) := \Pr[\mathcal{A}'(U_\theta) = 1]$ is a degree-$2t$ trigonometric polynomial, because it is the sum of squares of moduli of certain amplitudes, and each of these squares is a degree-$2t$ trigonometric polynomial. The correctness of the algorithm ensures that $p(0) \in [1-\varepsilon, 1]$ and $p(\theta) \in [0, \varepsilon]$ for all $\theta \notin [-3\delta, 3\delta] \mod 2\pi$. See Figure 1 for a visual depiction of the behaviour of p for $\theta \in [-\pi, \pi)$.

Scaling by a global factor of $1/\varepsilon$, we obtain a trigonometric polynomial q of degree $2t$ satisfying:

- $q(0) \geq (1-\varepsilon)/\varepsilon > 1/(2\varepsilon)$, and
- $q(\theta) \in [0, 1]$ for all $\theta \in [-\pi, \pi) \setminus [-3\delta, 3\delta]$.

Thus, Theorem 2.7 is applicable with $s = 6\delta$ and $n = 2t$, which implies $1/(2\varepsilon) \leq \sup_{x\in\mathbb{R}} |q(x)| \leq \exp(48t\delta)$. By taking logarithms and rearranging we get $t = \Omega\left(\frac{1}{\delta}\log\frac{1}{\varepsilon}\right)$, proving the theorem. □

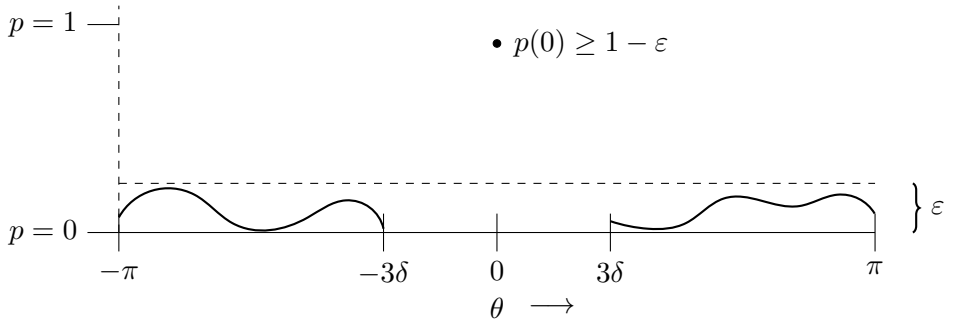

Figure 1: Acceptance probability p of $\mathcal{A}'$ as a function of θ in the proof of Claim 5.2

6 Conclusion

In this paper we considered several natural variants of the fundamental phase estimation problem in quantum computing, and proved essentially tight bounds on their cost in each setting. As an immediate application of one of our bounds, we resolved an open question of [SY23, Section 2].

We mention some interesting questions in the first variant of phase estimation we considered, where an algorithm is given a number of copies of advice states/unitaries instead of black-box access to a perfect eigenstate as in the basic phase estimation setup. First, are the logarithmic overheads for the cost in the input dimension N and the inverse of the overlap γ in our upper bounds (see Table 1) necessary, or can we give tighter upper bounds? Second, what is the optimal gate-complexity for rows 2 and 4? Third, can we show the $\log(1/\varepsilon)$-dependence on the error probability also in the advice-guided case, like we did for basic phase estimation (Theorem 1.3)?

Acknowledgements. We thank Jordi Weggemans for useful comments and for a pointer to [LT20], and Han-Hsuan Lin for pointing out an error in an earlier version which we corrected here (see Footnote 2).

References

[AGGW20] Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. arXiv:1705.01843. Earlier version in FOCS'17. doi:10.22331/q-2020-02-14-230.

[Amb02] Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier version in STOC'00. doi:10.1006/jcss.2002.1826.

[BBC+01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. quant-ph/9802049. Earlier version in FOCS'98. doi:10.1145/502090.502097.

[BCWZ99] Harry Buhrman, Richard Cleve, Ronald de Wolf, and Christof Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proceedings of 40th IEEE FOCS*, pages 358–368, 1999. arXiv:cs/9904019. doi:10.1109/SFFCS.1999.814607.

[BE95] Peter Borwein and Tamás Erdélyi. *Polynomials and polynomial inequalities*, volume 161. Springer Science & Business Media, 1995. doi:10.1007/978-1-4612-0793-1.

[Bes05] Arvid J. Bessen. Lower bound for quantum phase estimation. *Physical Review A*, 71(4):042313, 2005. doi:10.1103/PhysRevA.71.042313.

[BHMT02] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74. 2002. arXiv:quant-ph/0005055.

[CEMM98] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. In *Proceedings of the Royal Society of London*, volume A454, pages 339–354, 1998. quant-ph/9708016. doi:10.1098/rspa.1998.0164.

[CFG+23] Chris Cade, Marten Folkertsma, Sevag Gharibian, Ryu Hayakawa, François Le Gall, Tomoyuki Morimae, and Jordi Weggemans. Improved hardness results for the guided local hamiltonian problem. In *Proceedings of 50th ICALP*, pages 32:1–32:19, 2023. arXiv:2207.10250. doi:10.4230/LIPICS.ICALP.2023.32.

[CW23] Yanlin Chen and Ronald de Wolf. Quantum algorithms and lower bounds for linear regression with norm constraints. In *Proceedings of 50th ICALP*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:21, 2023. arXiv:2110.13086. doi:10.4230/LIPIcs.ICALP.2023.38.

[GG23] Sevag Gharibian and François Le Gall. Dequantizing the quantum singular value transformation: Hardness and applications to quantum chemistry and the quantum PCP conjecture. *SIAM Journal on Computing*, 52(4):1009–1038, 2023. Earlier version in STOC'22. doi:10.1137/22M1513721.

[Gil19] András Gilyén. *Quantum Singular Value Transformation & Its Algorithmic Applications*. PhD thesis, University of Amsterdam, 2019. URL: https://pure.uva.nl/ws/files/35292358/Thesis.pdf.

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th ACM STOC*, pages 212–219, 1996. quant-ph/9605043. doi:10.1145/237814.237866.

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of 51st ACM STOC*, pages 193–204, 2019. arXiv:1806.01838. doi:10.1145/3313276.3316366.

[GTC19] Yimin Ge, Jordi Tura, and J. Ignacio Cirac. Faster ground state preparation and high-precision ground energy estimation with fewer qubits. *Journal of Mathematical Physics*, 60(2):022202, 2019. arXiv:1712.03193. doi:10.1063/1.5027484.

[HHL09] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for solving linear systems of equations. *Physical Review Letters*, 103(15):150502, 2009. arXiv:0811.3171. URL: https://doi.org/10.1103/PhysRevLett.103.150502.

[Kit95] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem, 1995. arXiv:quant-ph/9511026.

[KLL+17] Shelby Kimmel, Cedric Yen-Yu Lin, Guang Hao Low, Maris Ozols, and Theodore J. Yoder. Hamiltonian simulation with optimal sample complexity. *npj Quantum Information*, 3(13), 2017. arXiv:1608.00281. doi:10.1038/s41534-017-0013-7.

[Lin23] Yao-Ting Lin. A note on quantum phase estimation. 2023. arXiv:2304.02241.

[LMR13] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10:631–633, 2013. arXiv:1307.0401. doi:10.1038/nphys3029.

[LT20] Lin Lin and Yu Tong. Near-optimal ground state preparation. *Quantum*, 4(372), 2020. arXiv:2002.12508. doi:10.22331/q-2020-12-14-372.

[LT22] Lin Lin and Yu Tong. Heisenberg-limited ground state energy estimation for early fault-tolerant quantum computers. *PRX Quantum*, 3(010318), 2022. arXiv:2102.11340. doi:10.1103/PRXQuantum.3.010318.

[LW22] Noah Linden and Ronald de Wolf. Average-case verification of the quantum Fourier transform enables worst-case phase estimation. *Quantum*, 6(872), 2022. arXiv:2109.10215. doi:10.22331/q-2022-12-07-872.

[NW99] Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Proceedings of 31st ACM STOC*, pages 384–393, 1999. doi:10.1145/301250.301349.

[PMS+14] Alberto Peruzzo, Jerrad McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter Love, Alán Aspuru-Guzik, and Jeremy O'Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 3:24, 2014. arXiv:1304.3061. doi:10.1038/ncomms5213.

[PW09] David Poulin and Pawel Wocjan. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters*, 103(22):220502, 2009. arXiv:0905.2199. doi:10.1103/PhysRevLett.103.220502.

[Ral21] Patrick Rall. Faster coherent quantum algorithms for phase, energy, and amplitude estimation. *Quantum*, 5(566), 2021. arXiv:2103.09717. doi:10.22331/q-2021-10-19-566.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. Earlier version in FOCS'94. quant-ph/9508027. doi:10.1137/S0097539795293172.

[SY23] Adrian She and Henry Yuen. Unitary property testing lower bounds by polynomials. In *Proceedings of 14th ITCS*, volume 251, pages 96:1–96:17, 2023. arXiv:2210.05885. doi:10.4230/LIPIcs.ITCS.2023.96.

[WBC22] Kianna Wan, Mario Berta, and Earl T. Campbell. A randomized quantum algorithm for statistical phase estimation. *Physical Review Letters*, 129(030503), 2022. arXiv:2110.12071. doi:10.1103/PhysRevLett.129.030503.

[WFC24] Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable Local Hamiltonian problems with implications to heuristic ansatz state preparation and the quantum PCP conjecture. In *Proceedings of 19th TQC*, pages 10:1–10:24, 2024. arXiv:2302.11578. doi:10.4230/LIPIcs.TQC.2024.10.

[Wol08] Ronald de Wolf. A note on quantum algorithms and the minimal degree of ε-error polynomials for symmetric functions. *Quantum Information and Computation*, 8(10):943–950, 2008. arXiv:0802.1816.

[Wol19] Ronald de Wolf. Quantum computing: Lecture notes, 2019. arXiv:1907.09415, version 5. URL: http://arxiv.org/abs/1907.09415.

A Proof of Theorem 3.2

In this section we prove Theorem 3.2 by a simple modification of the adversary method [Amb02].

Proof of Theorem 3.2. Let $\mathcal{A}$ be a cost-C algorithm solving $\mathsf{frOR}_{N,\delta,t}$. Define $U_j = I - (1-e^{i\delta})|j\rangle\langle j|$ for $j \in \{1,\ldots,N-1\}$, and set $U = I$ if $j = 0$. We will use the following advice states: $|\alpha_0\rangle = |0\rangle$ and $|\alpha_j\rangle = \gamma|j\rangle + \sqrt{1-\gamma^2}|0\rangle$ for $j \in \{1,\ldots,N-1\}$. For $j \in \{0,1,\ldots,N-1\}$ let $|\psi_j^0\rangle = |\alpha_j\rangle^{\otimes t}|0\rangle$ be the initial state, which includes t copies of the advice state, and for $T \in \{1,\ldots,C\}$ let $|\psi_j^T\rangle$ be the state of the algorithm (on input $U = U_j$ with initial state $|\psi_j^0\rangle$) just before the Tth application of U or its inverse.

Define the following progress measure P as a function of the timestep T:

$$P(T) = \sum_{j=1}^{N-1} |\langle\psi_0^T|\psi_j^T\rangle|.$$

We have

$$P(0) = \sum_{j=1}^{N-1} |\langle\psi_0^0|\psi_j^0\rangle| = \sum_{j=1}^{N-1} (1-\gamma^2)^{t/2} = (N-1)(1-\gamma^2)^{t/2}. \tag{1}$$

Since the output of $\mathcal{A}$ is different with high probability for $U = I$ on the one hand or for one of the other U_j on the other hand, it is easy to show that $|\langle\psi_0^C|\psi_j^C\rangle|$ is bounded below 1, say ≤ 0.99. Thus,

$$P(C) \leq 0.99(N-1). \tag{2}$$

The assumption $t = o(1/\gamma^2)$ implies $(1-\gamma^2)^{t/2} - 0.99 = \Omega(1)$, so we see that $P(C)$ is significantly smaller than $P(0)$. We now want to upper bound how much $P(T)$ can shrink in one step, in order to lower bound the number of steps.

For all $j \in \{0,1,\ldots,N-1\}$, define real amplitudes α_{jk} and normalized workspace states $|w_{jk}\rangle$ (depending on T, but we suppress this dependence in our notation) such that

$$|\psi_j^T\rangle = \sum_{k=0}^{N-1} \alpha_{jk}|k\rangle|w_{jk}\rangle.$$

Then

$$\langle\psi_0^T|\psi_j^T\rangle = \sum_{k=0}^{N-1} \alpha_{0k}\alpha_{jk}\langle w_{0k}|w_{jk}\rangle \quad \text{for all } j \in \{1,\ldots,N-1\}. \tag{3}$$

We also have, after applying U:

$$\begin{aligned} I|\psi_0^T\rangle &= \sum_{k=0}^{N-1} \alpha_{0k}|k\rangle|w_{0k}\rangle, \\ U_j|\psi_j^T\rangle &= \sum_{k\in\{0,\ldots,N-1\}\setminus\{j\}} \alpha_{jk}|k\rangle|w_{jk}\rangle \;+\; e^{i\delta}\alpha_{jj}|j\rangle|w_{jj}\rangle \quad \text{for all } j \in \{1,\ldots,N-1\}. \end{aligned}$$

Using the fact that inner products are not changed by the input-independent unitary that follows the application of U, we have for all $j \in \{1,\ldots,N-1\}$

$$\langle\psi_0^{T+1}|\psi_j^{T+1}\rangle = \sum_{k\in\{0,\ldots,N-1\}\setminus\{j\}} \alpha_{0k}\alpha_{jk}\langle w_{0k}|w_{jk}\rangle \;+\; e^{i\delta}\alpha_{0j}\alpha_{jj}\langle w_{0j}|w_{jj}\rangle. \tag{4}$$

Using Equations (3) and (4),

$$\begin{aligned}|\langle\psi_0^T|\psi_j^T\rangle| - |\langle\psi_0^{T+1}|\psi_j^{T+1}\rangle| &\leq |\langle\psi_0^T|\psi_j^T\rangle - \langle\psi_0^{T+1}|\psi_j^{T+1}\rangle| \\ &= |(1-e^{i\delta})\alpha_{0j}\alpha_{jj}\langle w_{0j}|w_{jj}\rangle| \\ &= |1-e^{i\delta}|\cdot|\alpha_{0j}|\cdot|\alpha_{jj}|\cdot|\langle w_{0j}|w_{jj}\rangle| \\ &\leq |1-e^{i\delta}|\cdot|\alpha_{0j}|,\end{aligned}$$

where the first inequality uses the triangle inequality and the last inequality holds because $|\alpha_{jj}|\cdot|\langle w_{0j}|w_{jj}\rangle| \leq 1$ for all j, T. Summing over all $j \in \{1,\ldots,N-1\}$, using the Cauchy-Schwarz inequality and the fact that $\sum_{j=1}^{N-1}|\alpha_{0j}|^2 \leq 1$, we can bound the change in the progress measure in one step by:

$$P(T)-P(T+1) = \sum_{j=1}^{N-1}|\langle\psi_0^T|\psi_j^T\rangle| - |\langle\psi_0^{T+1}|\psi_j^{T+1}\rangle| \leq \sum_{j=1}^{N-1}|1-e^{i\delta}|\cdot|\alpha_{0j}| \leq |1-e^{i\delta}|\sqrt{N-1}. \tag{5}$$

By Equations (1), (2), (5) and a telescoping sum, we have

$$(N-1)((1-\gamma^2)^{t/2}-0.99) \leq P(0)-P(C) = \sum_{T=0}^{C-1}(P(T)-P(T+1)) \leq C|1-e^{i\delta}|\sqrt{N-1}. \tag{6}$$

We have $|1-e^{i\delta}| = 2\sin(\delta/2) \leq \delta$ for all $\delta \in [0,\pi]$. As already mentioned, the assumption $t = o(1/\gamma^2)$ implies $(1-\gamma^2)^{t/2} - 0.99 = \Omega(1)$. Hence we obtain the desired lower bound $C = \Omega(\sqrt{N}/\delta)$ by rearranging Equation (6). □