THEORY OF COMPUTING, Volume 3 (2007), pp. 61–79
http://theoryofcomputing.org

A Simple PromiseBQP-complete Matrix Problem

Dominik Janzing Pawel Wocjan

Received: October 22, 2006; published: March 30, 2007.

Abstract: Let A be a real symmetric matrix of size N such that the number of non-zero entries in each row is polylogarithmic in N and the positions and the values of these entries are specified by an efficiently computable function. We consider the problem of estimating an arbitrary diagonal entry $(A^m)_{jj}$ of the matrix A^m up to an error of εb^m, where b is an a priori given upper bound on the norm of A and m and ε are polylogarithmic and inverse polylogarithmic in N, respectively.

We show that this problem is PromiseBQP-complete. It can be solved efficiently on a quantum computer by repeatedly applying measurements of A to the jth basis vector and raising the outcome to the mth power. Conversely, every uniform quantum circuit of polynomial length can be encoded into a sparse matrix such that some basis vector $|j\rangle$ corresponding to the input induces two different spectral measures depending on whether the input is accepted or not. These measures can be distinguished by estimating the mth statistical moment for some appropriately chosen m, i. e., by the jth diagonal entry of A^m. The problem remains PromiseBQP-hard when restricted to matrices having only -1, 0, and 1 as entries. Estimating off-diagonal entries is also PromiseBQP-complete.

ACM Classification: F.1.3, G.1.3

AMS Classification: 15A18, 15A60, 65F50, 65F15

Key words and phrases: quantum computation, complexity, BQP, completeness, PromiseBQP, PromiseBQP-complete problem, “non-quantum” characterization of BQP

DOI: 10.4086/toc.2007.v003a004

1 Introduction

It is still not understood well enough which problems are tractable for quantum computers. It is therefore desirable to better understand the class of problems which can be solved efficiently on a quantum computer. In quantum complexity theory, this class is referred to as BQP. Meanwhile, some characterizations of BQP are known [19, 24, 2, 25]. Strictly speaking, these are characterizations of the class PromiseBQP instead of BQP, a difference that is often ignored in the literature (see Section 2 for clarifying remarks). The class BQP, like its classical analogue BPP, is not known to have complete problems. Here we present a new characterization of PromiseBQP which is related to the computation of powers of large matrices.

It should not be too surprising that computational problems can be formulated in terms of "large" matrices. For example, the transformations of a quantum computer can be represented by multiplication of matrices of a certain type. However the matrix problems derived from this representation would usually not be very natural in classical terms (they are, of course, natural, as physical questions about the behavior of quantum systems). For instance, the problem of estimating an entry of products of unitary matrices which are given by a tensor embedding of low-dimensional unitaries, is PromiseBQP-complete, but it is not obvious where problems of this nature could arise in real-life applications referring to the macroscopic world.

It is known that Hamiltonians with finite range interactions can generate sufficiently complex dynamics that can serve as autonomous programmable quantum computers [15, 14]. Therefore, it is not unexpected that problems related to spectra and eigenvectors of self-adjoint operators lead to computationally hard problems. One might think that many of such problems could be solved efficiently on a quantum computer. However, results proving that questions pertaining to the minimal eigenvalue of Hamiltonians are PromiseQMA-complete [18, 17, 21] demonstrate that efficient algorithms are unlikely to exist for this problem.

The situation changes dramatically when we do not aim at deciding whether some Hamiltonian H has an eigenvalue below a certain bound but only whether a given state $|\psi\rangle$ has a considerable component in the eigenspace corresponding to a particular eigenvalue of H. This problem can be answered by (1) applying H-measurements to $|\psi\rangle$ several times and (2) statistically evaluating the obtained samples. It has been shown in [25] that measurements of so-called k-local operators,[1] applied to a basis state, solve all problems in PromiseBQP. This proves that some class of problems concerning the spectral measure of k-local self-adjoint operators associated with a given state characterize the class of problems that can be solved efficiently on a quantum computer. Unfortunately, the requirement of k-locality restricts the applicability of these results since it is not clear where k-local matrices occur apart from in the study of quantum systems. For this reason we have considered sparse matrices that do not require such a k-local structure and show that a very natural problem, namely the computation of diagonal entries of their powers, characterizes the complexity class PromiseBQP.

The paper is organized as follows. In Section 2 we define the complexity classes PromiseBQP and BQP and clarify the difference. In Section 3 we review known characterization of PromiseBQP. In Section 4 we define formally the problem of estimating diagonal entries and in Section 5 we prove that this problem can be solved efficiently on a quantum computer. To this end, we use the quantum phase

[1] An operator on n qubits is called k-local if it can be decomposed into a sum of terms that act on at most k qubits [18]

estimation algorithm to implement a measurement of the observable defined by the sparse matrix. To do this it is necessary that the time evolution defined by the sparse matrix can be implemented efficiently. Since the diagonal entries of the mth powers are the mth statistical moments of the spectral measure, we can estimate them after polynomially many measurements provided that the accuracy is sufficiently high. An appropriate decision problem, namely to decide whether this statistical moment is greater than a certain bound or smaller than another bound (given the *promise* that either of these alternatives is true), is therefore in PromiseBQP.

In Section 6 we show that diagonal entry estimation encompasses PromiseBQP. The proof relies on an encoding of the quantum circuit which solves the computational problem considered into a sparse self-adjoint matrix such that the spectral measure (and hence an appropriately chosen statistical moment) corresponding to the initial state depends on the solution. In Section 7 we show that the problem remains PromiseBQP-hard if restricted to matrices with entries -1, 0, and 1. The idea is that the gates, which are encoded into the constructed Hamiltonian, are not required to be unitary, even though the circuit that then realizes the corresponding measurement is certainly unitary. This fact could be interesting in its own right. For example, it could be possible that there are even more general ways of simulating non-unitary circuits by encoding them into self-adjoint operators. In this context, it would be interesting to clarify the relation to other measurement based schemes of computation [22, 7, 9].

2 Complexity theoretic clarifications: BQP and PromiseBQP

Certain complexity theoretic issues related to BQP are often blurred in the literature; therefore some clarifications seem to be in order. BQP is a class of languages. But in the literature, when people talk about BQP they often mean the promise-problem version (PromiseBQP). Exactly like with BPP and AM, BQP itself is not known to have complete problems, but *PromiseBQP* has complete promise problems, and that is adequate for most purposes.

The notion of promise problems was introduced and initially studied in [10]. Oded Goldreich's article [12] provides a survey of the most important applications that this notion has found in complexity theory. Most importantly, the author of this article argues that in some situations the use of promise problems is indispensable. These include the notion of "unique solutions" (e. g. unique-SAT), the formulation of "gap problems" (e. g. hardness of approximation), the identification of complete problems (e. g. for the class Statistical Zero Knowledge), the indication of separations between certain computational resources (e. g. the study of circuit complexity, derandomization, PCP and zero knowledge).

We refer the reader to this article for more details on promise problems and their applications. Our definition of *PromiseBQP* is modeled after Oded Goldreich's definition of *PromiseBPP* in [12, Definition 1.2]. We use his definition of *Karp-reduction* among promise problems [12, Definition 1.3] for reductions among problems in PromiseBQP.

Definition 2.1. A promise problem Π is a pair of non-intersecting sets, denoted $(\Pi_{\mathrm{YES}}, \Pi_{\mathrm{NO}})$, i. e., $\Pi_{\mathrm{YES}}, \Pi_{\mathrm{NO}} \subseteq \{0,1\}^*$ and $\Pi_{\mathrm{YES}} \cap \Pi_{\mathrm{NO}} = \emptyset$. The set $\Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{NO}}$ is called the *promise*.

Standard "language recognition" problems are cast as the special case in which that promise is the set of all strings, i. e., $\Pi_{\mathrm{YES}} \cup \Pi_{\mathrm{NO}} = \{0,1\}^*$. In this case we say that the *promise is trivial*. The standard definitions of complexity classes (i. e., classes of languages) extend naturally to promise problems. Then the set of NO-instances is not necessarily the complement of the set of YES-instances. Instead, the requirement is only that these two sets are non-intersecting. Our definition of PromiseBQP is as follows:

Definition 2.2. PromiseBQP is the class of promise problems $(\Pi_{\text{YES}}, \Pi_{\text{NO}})$ that can be solved by a uniform family[2] of quantum circuits. More precisely, it is required that there is a uniform family of quantum circuits Y_r acting on $\text{poly}(r)$ qubits that decide if a string $\mathbf{x}$ of length r is a YES-instance or NO-instance in the following sense. The application of Y_r to the computational basis state $|\mathbf{x},\mathbf{0}\rangle$ produces the state

$$Y_r|\mathbf{x},\mathbf{0}\rangle = \alpha_{\mathbf{x},0}|0\rangle \otimes |\psi_{\mathbf{x},0}\rangle + \alpha_{\mathbf{x},1}|1\rangle \otimes |\psi_{\mathbf{x},1}\rangle \tag{2.1}$$

such that

1. for every $\mathbf{x} \in \Pi_{\text{YES}}$ it holds that $|\alpha_{\mathbf{x},1}|^2 \geq 2/3$ and

2. for every $\mathbf{x} \in \Pi_{\text{YES}}$ it holds that $|\alpha_{\mathbf{x},1}|^2 \leq 1/3$.

Equation (2.1) has to be read as follows. The input string $\mathbf{x}$ determines the first r bits. Furthermore, ℓ additional ancilla bits are initialized to 0. After Y_r has been applied we interpret the first qubit as the relevant output and the remaining $r+\ell-1$ output values are irrelevant. The size of the ancilla register is polynomial in r.

Note that nothing is required in the definition of PromiseBQP with respect to inputs which violate the promise. For example, the problem of deciding whether a string is contained in the promise could be computationally much harder than the promise problem itself. It is clear that the promise on the probability gap between the instances YES and NO is necessary to decide the problem by measuring the output qubit. This motivates why promise problems appear quite naturally. We are now able to define BQP:

Definition 2.3. The class BQP is the subclass of PromiseBQP consisting of those promise problems for which the promise is trivial, i. e., is equal to the set of all strings $\{0,1\}^*$. In this case, the language $L = \Pi_{\text{YES}}$ is said to be a BQP-language.

One should emphasize that a language is in BQP if and only if the corresponding decision problem is in PromiseBQP since the notion of BQP-language implies that its whole complement is a NO-instance.

Definition 2.4. The promise problem $\Pi = (\Pi_{\text{YES}}, \Pi_{\text{NO}})$ is *Karp-reducible* to the promise problem $\Pi' = (\Pi'_{\text{YES}}, \Pi'_{\text{NO}})$ if there exists a polynomial-time computable function f (i. e., an efficiently computable classical function) such that

1. for every $\mathbf{x} \in \Pi_{\text{YES}}$ it holds that $f(\mathbf{x}) \in \Pi'_{\text{YES}}$ and

2. for every $\mathbf{x} \in \Pi_{\text{NO}}$ it holds that $f(\mathbf{x}) \in \Pi'_{\text{NO}}$.

[2]By "uniform circuit" we mean that there exists a polynomial time classical algorithm that generates the sequence of quantum gates for every desired input length.

3 Prior work on PromiseBQP-complete and PromiseBQP-hard problems

A simple observation showing that PromiseBQP-complete problems exist at all is the following. Any computational model that is universal for quantum computing immediately leads to a complete problem for PromiseBQP; namely, the problem of simulating that model and determining the output. For example, the problem of estimating entries of unitary matrices specified by quantum circuits can clearly be formulated as such a PromiseBQP-complete decision problem. So one can interpret the results proving that models such as adiabatic, topological, or one-way quantum computing are universal to be proving that the associated simulation problems are PromiseBQP-complete. However, such "quantum" problems do not really help us in understanding the difference between quantum and classical computation. An important challenge for complexity theory is therefore to construct problems that seem as classical as possible but characterize nevertheless the power of *quantum* computation.

Reference [19] characterizes the class PromiseBQP by a combinatorial problem, namely the problem to evaluate the so-called quadratically signed weight enumerators. They are given by

$$S(A,B,x,y) = \sum_{b,Ab=0} (-1)^{b^T Bb} x^{|b|} y^{n-|b|} ,$$

where A and B are $0,1$ matrices with B of dimension n by n and A of dimension m by n. The variable b in the sum ranges over column vectors of dimension n having entries $0,1$, b^T denotes the transpose of b, $|b|$ is the Hamming weight of b and all calculations involving A, B, and b are modulo 2. Let $\mathrm{lwtr}(A)$ denote the lower triangular part of A. Then it is PromiseBQP-complete to determine the sign of $S(A,\mathrm{lwtr}(A),k,\ell)$ for integers k,ℓ with a matrix A whose diagonal entries are 1. Here we are given the promise that the modulus of $S(A,\mathrm{lwtr}(A),k,\ell)$ is at least $(k^2+\ell^2)^{n/2}/2$. Since all matrices and vectors have entries $0,1$, this problem can certainly be considered as a *classical* problem.

Another PromiseBQP-complete problem could be formulated in the context of knot theory. In [3] it was shown that the quantum computer can efficiently provide estimations for the values of the Jones polynomial when evaluated at roots of unity. In [24, 2] it was shown that this evaluation can solve every problem in PromiseBQP. The idea is, roughly speaking, that the sequence of gates can be translated into sequences of braids whose unitary representations correspond directly to the gates. These links between quantum computing and knot theory are quite plausible when taking into account that the latter has been successfully applied to topological quantum field theories and that quantum computers can be useful to simulate topological field theories [11].

The complexity of certain quantum measurements was considered in [25] (see also [23]). It was shown that sampling from the spectral measure of 4-local n-qubit observables can solve all problems in PromiseBQP provided that every measurement precision being inverse polynomial in n can be achieved. Even though this problem is completely quantum, it provided one of the key idea of this paper. The essential insight is that measurement for appropriate observables, when applied to basis states, can solve problems in PromiseBQP. We convert the problem of [25] into a quantum-free problem by (1) replacing 4-local operators with general sparse matrices and (2) by replacing direct statements on the distribution of measurement outcomes with statements on the statistical moments of this probability measure. This simplifies the problem considerably since these moments are directly given by diagonal entries of powers of the matrix considered.

4 Definition of diagonal entry estimation

In this section, we formulate a quantum-free PromiseBQP-complete problem. Before we define the problem "diagonal entry estimation" we introduce the notion of sparse matrices and the spectral measure. Here we call an $N \times N$ matrix A row-sparse (column-sparse) if it has no more than $s = \text{polylog}(N)$ non-zero entries in each row (column) and there is an efficiently computable function f that specifies for a given row (column) the non-zero entries and their positions (compare [4, 8, 6]). Here and in the following the term "efficiency" is used in the sense that the computation time is polylogarithmic in N.

Let A be self-adjoint with spectral decomposition

$$A = \sum_{\lambda} \lambda\, Q_{\lambda}\,, \tag{4.1}$$

and $|\psi\rangle$ be some normalized vector of size N. The spectral measure induced by A and $|\psi\rangle$ is a probability distribution on the spectrum of A such that the eigenvalue λ occurs with probability $\|Q_{\lambda}|\psi\rangle\|^2$. Noting that $A^m = \sum_{\lambda} \lambda^m Q_{\lambda}$, we infer the following fact which will be repeatedly used in the sequel. The expected value of A^m in the state $|\psi\rangle$ is given by

$$\langle\psi|A^m|\psi\rangle = \sum_{\lambda} \lambda^m \langle\psi|Q_{\lambda}|\psi\rangle\,, \tag{4.2}$$

i. e., by the mth statistical moment of the spectral measure. The operator norm $\|A\|$ of A is given by the maximum over all $|\lambda|$ in Equation (4.1). In [25], eigenvalue sampling is defined to be a quantum process that allows us to sample from a probability distribution that coincides with the spectral measure induced by A and $|\psi\rangle$. Throughout the paper we refer to such a procedure as *measuring* the observable A in the state $|\psi\rangle$. Now we are ready to formalize the problem of estimating diagonal entries of powers of sparse symmetric matrices.

Definition 4.1. An instance of the promise problem "diagonal entry estimation" is specified by a tuple (A,b,m,j,ε,g), where A is a symmetric sparse matrix of size $N \times N$ with real entries, b is an upper bound on the norm $\|A\|$, $m = \text{polylog}(N)$ is a positive integer, $j \in [1,\dots,N]$, $\varepsilon = 1/\text{polylog}(N)$, and $g \in [-b^m, b^m]$.

The task is to decide if such a tuple (A,b,m,j,ε,g) is an element of Π_{YES} or an element of Π_{NO}. These instances are defined as follows:

1. (A,b,m,j,ε,g) is a YES-instance iff
$$(A^m)_{jj} \geq g + \varepsilon\, b^m\,,$$

2. (A,b,m,j,ε,g) is a NO-instance iff
$$(A^m)_{jj} \leq g - \varepsilon\, b^m\,.$$

The promise is that only tuples in $\Pi = \Pi_{\text{YES}} \cup \Pi_{\text{NO}}$ are considered. Any answer is acceptable when the entry is not between the stated bounds.

Recall that N is exponential in the input length and the description of the entries of A is given by a function that can be computed with running time in $O(\text{polylog}(N))$.

Special instances of this kind (matrices with entries $0,1$) arise in graph theory. Let A be the adjacency matrix of a graph with N vertices and degree bounded from above by s. Then the diagonal entry $(A^m)_{jj}$ of the mth power of A is equal to the number of walks of length m that start and end at the vertex j. Here sparseness means that for every node the number of neighbors is polynomial and that there is an efficiently computable function specifying the set of neighbors for each node. A natural setting satisfying these requirements is the following. Let the nodes represent the set of strings of length $n = \lceil \log_k N \rceil$ over some finite alphabet $\{\alpha_1, \ldots, \alpha_k\}$ with constant k. The edges are implicitly specified by a given equivalence relation on substrings of length l for some constant l in the sense that two strings are adjacent if they can be obtained from each other by replacing one substring with an equivalent one. There is certainly no promise that would here occur in a natural way. However, the promise is only needed to formulate a *decision* problem. Even without the promise, we can efficiently *estimate* the number of walks up to an accuracy that is specified by the gap in the promise decision problem.

The main contribution of this paper is the proof that diagonal entry estimation is PromiseBQP-complete.

Theorem 4.2. *The problem "diagonal entry estimation" is PromiseBQP-complete.*

We emphasize that this result also provides an understanding of the complexity class BQP, not only PromiseBQP because of the following observation:

Corollary 4.3. *A language L belongs to BQP if and only if L is Karp-reducible to the problem "diagonal entry estimation." Here Karp-reduction is meant in the sense of reduction between promise problems; a language recognition problem is simply a promise problem with trivial promise.*

First, we prove that the problem "diagonal entry estimation" is in PromiseBQP. Second, we prove that it is PromiseBQP-hard.

It is important to note that the scale on which the estimation has reasonable precision is given by b^m. If the a priori known bound on the norm is, for instance, $b' := 2b$ instead of b, then the accuracy is changed by the exponential factor 2^m. Our results show that quantum computation outperforms classical computation in estimating the diagonal entries (provided that PromiseBQP$\neq$ PromiseBPP). But one has to be very careful on which scale this result remains true.

5 Diagonal entry estimation is in PromiseBQP

We now describe how to construct a circuit that solves diagonal entry estimation. Without loss of generality we may assume $b = 1$ and rescale the measurement results later. The main idea is as follows. We measure the observable A in the state $|j\rangle$. We obtain an eigenvalue λ as result and compute λ^m. The average over these values over large sampling converges to the desired entry. The measurement is done by (1) considering A as a Hamiltonian of a quantum system and simulating the corresponding dynamics $U_t = \exp(-iAt)$ and (2) applying the phase estimation algorithm to U_t. The proof that this works follows from a careful analysis of possible error sources. These are

1. errors due to the statistical nature of the phase estimation algorithm,

2. statistical errors due to estimation of the expected value from the empirical mean, and

3. errors caused by the imperfect simulation of the Hamiltonian time evolution.

We show that all these errors can be made sufficiently small with polynomial resources only.

5.1 Inaccuracies of phase estimation

We embed A into the Hilbert space of n qubits, where $n = \lceil \log_2 N \rceil$. Let us first assume that the unitary matrix $U := \exp(iA)$ can be implemented exactly. We apply the phase estimation procedure which works as follows [20]. We start by adding p ancillas to the qubits on which U acts. The idea is to control the implementation of the 2^ℓth power of U by the ℓth control qubit. More precisely, we have the controlled gates

$$W_\ell := |0\rangle\langle 0|^{(\ell)} \otimes \mathbf{1} + |1\rangle\langle 1|^{(\ell)} \otimes U^{2^\ell},$$

where the superscript (ℓ) indicates that the projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ act on the ℓth control qubit, respectively. Note that the decomposition of W_1 into elementary gates is obtained by replacing each elementary gate in the circuit implementing U with a corresponding controlled gate. Similarly, W_ℓ is realized by applying the quantum circuit implementing the corresponding controlled U-gate 2^ℓ times. Set $W := W_1 W_2 \cdots W_p$. The phase estimation circuit consists of applying Hadamard gates on all control qubits, the circuit W, and the inverse Fourier transform on the control qubits. The desired value a is obtained by measuring the control qubits in the computational basis. Let $|\psi\rangle$ be an arbitrary eigenvector of U with unknown eigenvalue $e^{i2\pi\varphi}$ for some phase $\varphi \in [0,1)$. In order to ensure that the phase estimation algorithm outputs a random value $a \in \{0,\dots,2^p - 1\}$ such that

$$\Pr(|\varphi - a/2^p| < \eta) > 1 - \theta, \tag{5.1}$$

for some $\theta, \eta > 0$ it is sufficient [20] to set

$$p := \lceil \log(1/\eta) \rceil + \lceil \log\big(2 + (1/(2\theta))\big) \rceil.$$

Let $|\psi\rangle$ be an eigenvector of A with unknown eigenvalue $\lambda \in [-1,1]$. In order to determine λ approximately using the outcome a in a phase estimation with $U = \exp(iA)$ we proceed as follows. First, we have to take into account that $\varphi > 1/2$ corresponds to negative values λ. Second, we have to consider that the scaling differs by the factor 2π. Finally, we may use the additional information that not all λ in $[-\pi,\pi)$ are possible, but only those in $[-1,1]$. All outputs that would actually correspond to eigenvalues λ in $[1,\pi]$ and $[-\pi,-1)$ are therefore interpreted as $+1$ or -1, respectively. Therefore, we compute values z from the output a by

$$z := \begin{cases} a(2\pi/2^p) & \text{for} \quad 0 \le a < 2^p/(2\pi), \\ 1 & \text{for} \quad 2^p/(2\pi) \le a < 2^p/2, \\ -1 & \text{for} \quad 2^p/2 \le a < 2^p - 2^p/(2\pi), \\ a(2\pi/2^p) - 2\pi & \text{for} \quad 2^p - 2^p/(2\pi) \le a < 2^p. \end{cases}$$

This defines the random variable Z whose values z are approximations for λ that satisfy the following error bound:

$$\Pr\big(|\lambda - Z| < 2\pi\eta\big) > 1 - \theta.$$

This bound follows from Equation (5.1) by appropriate rescaling (note that our reinterpretation of values in $[-\pi,-1]$ and $[1,\pi)$ explained above can only decrease the error unless it was already greater than $\pi-1$). Consequently, we have for every eigenstate $|\psi_i\rangle$ of A with eigenvalue λ_i the statement

$$|E_{|\psi_i\rangle}(Z^m)-\lambda_i^m| \leq 2\,\theta+2\pi m\,\eta\,, \tag{5.2}$$

where $E_{|\psi_i\rangle}(Z^m)$ denotes the expected value of Z^m in the state $|\psi_i\rangle$. The first term on the right-hand side corresponds to the unlikely case that the measurement outcome deviates by more than $2\pi\eta$ from the true value. Since we do not have outcomes z smaller than -1 or greater than 1 the maximal error is at most 2. This leads to the error term 2θ. The second term corresponds to the case $|\lambda_i-z|\leq 2\pi\eta$, which implies $|\lambda_i^m-z^m|\leq(2\pi\eta)m$ because $\lambda_i,z\in[-1,+1]$.

We make the error in Equation (5.2) smaller than $\varepsilon/3$ by choosing the parameters θ and η such that $\theta<\varepsilon/12$ and $\eta<\varepsilon/(12\,\pi\,m)$. The number of control qubits can be chosen to be

$$p := 2\lceil\log((48\,m)/\varepsilon)\rceil\,. \tag{5.3}$$

This is sufficient since

$$\lceil\log(1/\eta)\rceil+\lceil\log\left(2+(1/(2\theta)\right)<2\lceil\log\left((48\,m)/\varepsilon\right)\rceil\,. \tag{5.4}$$

We decompose $|j\rangle$ into A-eigenstates

$$|j\rangle=\sum_i \beta_i|\psi_i\rangle\,,$$

and obtain the statement

$$E_{|j\rangle}(Z^m)=\sum_i|\beta_i|^2\,E_{|\psi_i\rangle}(Z^m)$$

by linearity arguments and

$$(A^m)_{jj}=\langle j|A^m|j\rangle=\sum_i\langle j|\psi_i\rangle\langle\psi_i|j\rangle\lambda_i^m=\sum_i|\beta_i|^2\lambda_i^m\,.$$

Using the triangle inequality and the fact that the right-hand side of Equation (5.2) is smaller than $\varepsilon/3$ for each i we obtain

$$\left|E_{|j\rangle}(Z^m)-(A^m)_{jj}\right|<\varepsilon/3\,. \tag{5.5}$$

5.2 Errors caused by finite sampling

Now we sample the measurement k times in order to estimate the expected value $E_{|j\rangle}(Z^m)$. Since we will later also consider the simulation error we want to estimate $(A^m)_{jj}$ up to an error of $2\varepsilon/3$. To achieve this, it is sufficient to estimate $E_{|j\rangle}(Z^m)$ up to an error of $\varepsilon/3$.

Let $\overline{Z^m}$ denote the average value of the random variable Z^m after sampling k times. Since the values of Z^m are between -1 and 1 we can give an upper bound for the probability that the average is not $\varepsilon/3$-close to the expected value. By Hoeffding's inequality [13, Theorem 2] we get

$$\Pr\left(\left|\overline{Z^m}-E_{|j\rangle}(Z^m)\right|\geq\frac{\varepsilon}{3}\right) \quad\leq\quad 2\exp\left(\frac{-\varepsilon^2}{18}k\right).$$

In summary, we have shown for $b = 1$ that we can distinguish between the two cases in Definition 4.1 with exponentially small error probability. For $b \neq 1$ we have to rescale the inaccuracy of the estimation by b^m. The whole procedure including repeated measurements and averaging can certainly be performed by a single quantum circuit Y_r in the sense of Definition 2.2.

5.3 Inaccuracies of the simulation of $\exp(-iAt)$

We now take into account that $U = \exp(iA)$ cannot be implemented exactly. It is known that the dynamics generated by A can be simulated efficiently if A is sparse [4, 8, 6]. More precisely, for each t we can construct a circuit V which is δ-close to $U_t = \exp(-iAt)$ with respect to the operator norm such that the required number of gates increases only polynomially in the parameters n, t, and $1/\delta$. We analyze the error resulting from using V instead of U, where $\|V - U\| \leq \delta$.

The phase estimation contains $2^{p+1} - 1$ copies of the controlled-V gate. Therefore the circuit F_V implementing the phase estimation procedure with V instead of U deviates from F_U by at most $2^{p+1}\,\delta$ with respect to the operator norm, that is, $\|F_U - F_V\| \leq 2^{p+1}\,\delta$.

Let q and $\tilde{q}$ denote the probability distributions of outcomes when measuring the control register after the phase estimation procedure has been implemented with V and U, respectively. The ℓ_1- distance between q and $\tilde{q}$ is then defined by

$$\|q - \tilde{q}\|_1 := \sum_{a \in \{0,\ldots,2^p-1\}} |q(a) - \tilde{q}(a)|$$

where $q(a)$ and $\tilde{q}(a)$ denote the probabilities of obtaining the outcome a according to the measure q and $\tilde{q}$, respectively. To upper bound $\|q - \tilde{q}\|_1$ we define a function s by $s(a) := 1$ if $q(a) > \tilde{a}$ and $s(a) := -1$ otherwise. Let Q be the observable defined by measuring the ancillas and applying s to the outcome a. Then we can write $\|q - \tilde{q}\|_1$ as a difference of expected values:

$$\langle \psi | F_U^\dagger Q F_U - F_V^\dagger Q F_V | \psi \rangle \leq 2 \|F_U - F_V\| \, \|Q\| \leq 2^{p+2}\,\delta\,.$$

This implies that the corresponding expected values of Z^m can differ by at most $2^{p+2}\,\delta$ because Z takes values only in the interval $[-1, 1]$. We choose the simulation accuracy such that $\delta = \varepsilon/(3 \cdot 2^{p+2})$ and obtain an additional error term of at most $\varepsilon/3$ in Equation (5.5). Using that we have chosen p as in Equation (5.3) we obtain that $\delta \in O(\varepsilon^3 m^2)$.

Putting everything together we obtain a total error of at most ε. Furthermore, this can be achieved by using time and space resources which are polynomial in n, m, and $1/\varepsilon$. This completes the proof that diagonal entry estimation is in PromiseBQP.

It should be mentioned that off-diagonal entries $(A^m)_{ij}$ can also be estimated efficiently on the same scale using superpositions $|i\rangle \pm |j\rangle$ since the values $\langle i|A^m|j\rangle$ can be expressed in terms of differences of the statistical moments of the spectral measure induced by those states. The scale on which the estimation can be done efficiently is then also given by $\varepsilon\, b^m$ with an appropriately modified ε which is still inverse polynomial in n. However, since PromiseBQP-hardness requires only diagonal entries we have focused our attention on the latter.

It is natural to ask whether the above result extends to non-symmetric matrices (note that the generalization to complex-valued hermitian matrices is obvious since we did not make use of the fact that the entries are real). The central part of the algorithm is a measurement of "the observable A," i. e., an procedure that samples from its spectral measure. The most general set of matrices that define a spectral measure is the set of normal operators, i. e., operators that commute with their adjoints. If we decompose normal matrices into

$$\mathrm{Re}(A) := \frac{1}{2}(A + A^\dagger) \qquad \text{and} \qquad \mathrm{Im}(A) := \frac{1}{2i}(A - A^\dagger),$$

the "real" and the "imaginary part" commute. We can thus implement both corresponding measurement procedures (by quantum phase estimation as above) one after another *without* preparing a new state. Then the output pair (μ, ν) defines a complex number $z := \mu + i\nu$ which is close to an eigenvalue of A. Because of inaccuracies in the measurement procedure we may obtain values with $|z| > \|A\|$. In analogy to the methods above we would then instead take the closest value on the circle of radius $\|A\|$. This ensures that we keep the estimation errors, again, small compared to $\|A\|^m$.

6 Diagonal entry estimation is PromiseBQP-hard

To show that the estimation of diagonal entries can solve all problems in PromiseBQP we prove that a particular instance is already PromiseBQP-hard. It is given by the problem to determine the sign of the diagonal entry when an appropriate lower bound on its modulus is provided by the promise.

We assume that we are given a quantum circuit Y_r that is able to decide whether a string $\mathbf{x}$ is in YES or NO in the sense of Definition 2.2. Using Y_r we construct a self-adjoint operator A such that the corresponding spectral measure induced by an appropriate initial state depends on whether $\mathbf{x}$ is in YES or in NO. Note that the proofs for PromiseQMA-completeness of eigenvalue problems for Hamiltonians have already used this idea to construct a self-adjoint operator whose spectral properties encode a given quantum circuit [18, 17, 21]. In these constructions, the *existence* of eigenvalues of a given Hamiltonian depends on whether or not an input state exists that is accepted by a certain circuit. In PromiseBQP, the problem is only to decide whether a *given* state is accepted and not whether such a state exists. Likewise, the problem is not to decide whether an eigenvalue of the constructed observable *exists* which lies in a certain interval. Instead, it refers only to the spectral measure induced by a *given* state. This difference changes the complexity from PromiseQMA to PromiseBQP.

For these reasons, our construction is based on [25] and not on work related to PromiseQMA. Reference [25] established the PromiseBQP-hardness of approximate k-local measurements. This result relied on the ideas in [23] where the PSPACE-hardness of k-local measurements was proved provided that exponentially small error is desired.

However, our description below will only at one point refer to these results since the observable we construct here is only required to be sparse, in contrast to the k-locality assumed in [25, 23]. In some analogy to [16, 25] we construct a circuit U that is obtained from Y_r as follows: First, apply the circuit Y_r. Second, apply a σ_z-gate. Third, implement $Y_r^\dagger$. The resulting circuit U is shown in Figure 1. We denote the dimension of the Hilbert space U acts on by $\tilde{N}$.

Let U be generated by a concatenation of the M elementary gates $U_0, \ldots, U_{M-1}$. The results in [5] imply that we can simulate U by gates having only real entries. To this end, a qubit is added that is used to represent the real and imaginary part of the quantum state. Then the real circuit reproduces exactly the

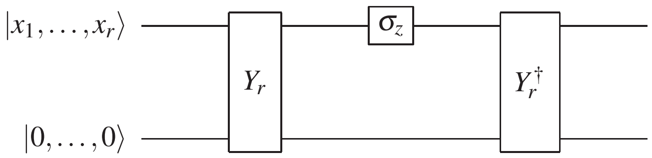

Figure 1: Circuit U constructed from the original circuit Y_r. Whenever the answer to the BQP problem is no, the output state of U is close to the input state $|\mathbf{x},\mathbf{0}\rangle \equiv |x_1,\dots,x_n,0,\dots,0\rangle$. Otherwise, the state $|\mathbf{x},\mathbf{0}\rangle$ is only restored after applying U twice.

output probabilities of the original circuit. We assume furthermore that M is odd, which is automatically satisfied if we decompose $Y_r^\dagger$ in analogy to Y_r and implement a σ_z-gate between Y_r and $Y_r^\dagger$. We define the unitary matrix

$$W := \sum_{\ell=0}^{M-1} |\ell+1\rangle\langle\ell| \otimes U_\ell\,, \tag{6.1}$$

acting on $\mathbb{C}^M \otimes \mathbb{C}^{\tilde{N}}$. Here the $+$ sign in the index has always to be read modulo M. We obtain

$$W^M = \sum_{\ell=0}^{M-1} |\ell\rangle\langle\ell| \otimes U_{\ell+M}\cdots U_{\ell+1}U_\ell\,.$$

Because $U^2 = \mathbf{1}$ we have $(W^M)^2 = \mathbf{1}$. Thus, W^M can only have the eigenvalues ± 1. This defines a decomposition of the space $\mathbb{C}^M \otimes \mathbb{C}^{\tilde{N}}$ into symmetric and antisymmetric W-invariant subspaces $\mathcal{S}^+$ and $\mathcal{S}^-$, respectively with corresponding projectors

$$Q^\pm := \frac{1}{2}(\mathbf{1} \pm W^M)\,.$$

In the following we use the definition $|s_\mathbf{x}\rangle := |0\rangle \otimes |\mathbf{x},\mathbf{0}\rangle$ for the initial state and restrict attention to the span of the orbit

$$\left\{ W^\ell |s_\mathbf{x}\rangle \right\} \quad \text{with } \ell \in \mathbb{N}\,. \tag{6.2}$$

Moreover, we use the abbreviations $\alpha_0 = \alpha_{\mathbf{x},0}$ and $\alpha_1 = \alpha_{\mathbf{x},1}$. We consider first the two extreme cases $|\alpha_1| = 0$ and $|\alpha_1| = 1$. If $|\alpha_1| = 0$ the orbit (6.2) is M-periodic and the action of W is isomorphic to the action of a cyclic shift in M dimensions, i. e., the mapping $|\ell\rangle \mapsto |(\ell+1) \bmod M\rangle$, where $|\ell\rangle$ corresponds to $W^\ell|s_\mathbf{x}\rangle$ with $\ell = 0,1,\dots,M-1$.

If $|\alpha_1| = 1$ the action of W corresponds to a cyclic shift with an additional phase -1, i. e., the mapping $|\ell\rangle \mapsto |\ell+1\rangle$ for $\ell = 0,1,\dots,M-2$ and $|M-1\rangle \mapsto -|0\rangle$. In the first case, the state $|s_\mathbf{x}\rangle$ induces a spectral measure $R^{(0)}$ being equal to the uniform distribution on the Mth roots of unity, i. e., the values $\exp(-i\pi 2\ell/M)$ for $\ell = 0,\dots,M-1$. In the second case, $|s_\mathbf{x}\rangle$ induces the measure $R^{(1)}$ being equal to the uniform distribution on the values $\exp(-i\pi(2\ell+1)/M)$ for $\ell = 0,\dots,M-1$. We observe that $R^{(1)}$ and $R^{(0)}$ coincide up to a reflection of the real axis in the complex plane.

In the general case, the orbit defines an $2M$-dimensional space whose orthonormal basis vectors are obtained by renormalizing the vectors

$$W^\ell Q^+ |s_\mathbf{x}\rangle \quad \text{and} \quad W^\ell Q^- |s_\mathbf{x}\rangle \quad \text{with } \ell = 0,1,\dots,M-1\,.$$

We obtain then a convex sum of $R^{(0)}$ and $R^{(1)}$ as spectral measures induced by W and $|s_{\mathbf{x}}\rangle$. The following calculation shows that $|\alpha_0|^2$ and $|\alpha_1|^2$ define the corresponding weights:

$$\langle s_{\mathbf{x}}|Q^+|s_{\mathbf{x}}\rangle = \frac{1}{2}\langle s_{\mathbf{x}}|\mathbf{1}+W^M|s_{\mathbf{x}}\rangle = \frac{1}{2}\langle \mathbf{x},\mathbf{0}|\mathbf{1}+U|\mathbf{0},\mathbf{x}\rangle = \frac{1}{2}(1+\langle \mathbf{x},\mathbf{0}|Y_r^\dagger \sigma_z Y_r|\mathbf{0},\mathbf{x}\rangle) = |\alpha_0|^2\,,$$

where the last equality follows easily by replacing $Y_r|\mathbf{x},\mathbf{0}\rangle$ and its adjoint with the expression in Equation (2.1) and its adjoint. Thus, we obtain the spectral measure

$$R := |\alpha_0|^2 R^{(0)} + |\alpha_1|^2 R^{(1)}\,.$$

We define the self-adjoint operator

$$A := \frac{1}{2}(W+W^\dagger)\,. \tag{6.3}$$

The support of the spectral measure corresponding to A is directly given by the real part of the support of R. To obtain the corresponding probabilities one has to take into account that in many cases two different eigenvalues of W lead to the same eigenvalue of A.

To calculate the distribution of outcomes for A-measurements we observe that $R^{(0)}$ leads to a distribution $P^{(0)}$ on the $(M-1)/2$ eigenvalues

$$\lambda_\ell^{(0)} = \cos\frac{2\pi\ell}{M} \quad \text{for } \ell = 0,\ldots,(M-1)/2$$

with probabilities $P_1^{(0)} := 1/M$ and $P_\ell^{(0)} := 2/M$ for $\ell > 1$. Likewise, $R^{(1)}$ leads to a distribution $P^{(1)}$ on the $(M-1)/2$ values

$$\lambda_\ell^{(1)} = \cos\frac{\pi(2\ell+1)}{M} \quad \text{for } \ell = 0,\ldots,(M-1)/2$$

with probabilities $P^{(1)}_{(M-1)/2} = 1/M$ and $P_\ell^{(1)} = 2/M$ for $\ell < (M-1)/2$. As it was true for $R^{(0)}$ and $R^{(1)}$, the measures $P^{(0)}$ and $P^{(1)}$ coincide up to a reflection.

We now set $|j\rangle := |s_{\mathbf{x}}\rangle$, i. e., the input state is considered as the jth basis vector of $\mathbb{C}^M \otimes \mathbb{C}^{\tilde{N}}$. Then the diagonal entry $(A^m)_{jj}$ coincides with the mth moment of the spectral measure:

$$(A^m)_{jj} = \langle j|A^m|j\rangle = \sum_\lambda \lambda^m P(\lambda)\,,$$

where λ runs over all eigenvalues of the restriction of A to the smallest A-invariant subspace containing $|j\rangle$, and $P(\lambda)$ denotes its probability according to the spectral measure corresponding to A. Since the latter is a convex sum of $P^{(0)}$ and $P^{(1)}$ we may write $(A^m)_{jj}$ as the convex sum

$$\begin{aligned} (A^m)_{jj} &= (1-|\alpha_1|^2)\sum_\ell \left(\lambda_\ell^{(0)}\right)^m P_\ell^{(0)} + |\alpha_1|^2 \sum_\ell \left(\lambda_\ell^{(1)}\right)^m P_l^{(1)} \\ &=: (1-|\alpha_1|^2)E_0 + |\alpha_1|^2 E_1\,. \end{aligned} \tag{6.4}$$

The values E_0 and E_1 can be considered as the mth statistical moments of random variables on $[-1,1]$ whose distributions are given by $P^{(0)}$ and $P^{(1)}$, respectively.

In order to see how the value $(A^m)_{jj}$ changes with $|\alpha_1|$ we observe that,

$$E_0 = \sum_{\ell=0}^{(M-1)/2} \left(\lambda_\ell^{(0)}\right)^m P_\ell^{(0)} \geq P_0^{(0)} + \left(\lambda_{(M-1)/2}^{(0)}\right)^m = \frac{1}{M} + \left(\lambda_{(M-1)/2}^{(0)}\right)^m .$$

Here we have used that $\lambda_0^{(0)} = 1$ and that the eigenvalues are numbered in a decreasing order. Thus, $\lambda_{(M-1)/2}$ is the smallest one. Because of the reflection symmetry of the measures we have $E_1 = -E_0$. Now we choose m sufficiently large such that the term $(\lambda_{(M-1)/2}^{(0)})^m$ is negligible compared to $1/M$ since we have then $E_0 - E_1 \approx 2/M$ which is a sufficient difference for our purpose.

In order to achieve this we set $m := M^3$. We have

$$\lambda_{(M-1)/2}^{(0)} = -\cos(\pi/M) > -1 + \frac{\pi^2}{2M^2} - \frac{\pi^4}{4!\,M^4} > -1 + \frac{\pi^2}{4M^2} ,$$

where the last inequality holds for sufficiently large M. Due to

$$\lim_{M\to\infty} (1 - \frac{\pi^2}{4M^2})^{M^2} = e^{\frac{-\pi^2}{4}}$$

we conclude that

$$(\cos(\pi/M))^{M^3} < (e^{-\frac{\pi^2}{4}})^M ,$$

and hence

$$E_0 > \frac{1}{M} - (e^{-\frac{\pi^2}{4}})^M > \frac{3}{4M} , \tag{6.5}$$

where we have, again, assumed M to be sufficiently large. To see how $(A^m)_{jj}$ changes with $|\alpha_1|$ we recall

$$(A^m)_{jj} = (1 - |\alpha_1|^2)E_0 + |\alpha_1|^2 E_1 = (1 - 2|\alpha_1|^2)\, E_0 ,$$

by Equation (6.4) and the reflection symmetry. Using the worst cases $|\alpha_1|^2 = 1/3$ for $x \in \Pi_{\mathrm{NO}}$ and $|\alpha_1|^2 = 2/3$ for $x \in \Pi_{\mathrm{YES}}$ we obtain

$$(A^m)_{jj} = \frac{1}{3}E_0 \qquad \text{and} \qquad (A^m)_{jj} = -\frac{1}{3}E_0 .$$

Using $E_0 > 3/(4M)$ from Equation (6.5) we obtain

$$(A^m)_{jj} > \frac{1}{4M} ,$$

if the answer is no and

$$(A^m)_{jj} < -\frac{1}{4M}$$

otherwise. Setting $g := 0$ (see Definition 4.1) we may define $\varepsilon := 1/(4M)$. Then the diagonal entry is greater than $g + \varepsilon$ if $x \in \Pi_{\mathrm{YES}}$ and smaller than $g - \varepsilon$ if $x \in \Pi_{\mathrm{NO}}$. The construction of A as the real part of a unitary matrix ensures that $\|A\| \leq 1 =: b$. This shows that we can find an inverse polynomial accuracy ε such that an estimation of the diagonal entry up to an error $\varepsilon\, b^m$ allows us to check whether $x \in \Pi_{\mathrm{YES}}$.

It remains to show that A is row-sparse and column-sparse in the sense defined in Section 4. First we observe the following. For a gate U_ℓ that only acts on k qubits non-trivially, the matrix describing U_ℓ contains only non-zero entries $\langle b|U_\ell|b'\rangle$ for those pairs b,b' of binary words for which b and b' differ at most at these k positions. For a given b, one can efficiently check which one of these possible 2^k entries is non-zero and similarly for a given b'. Hence we have shown sparseness of U_ℓ. It is easy to show that sparseness is also true for W as defined in Equation (6.1) using the gates U_ℓ and also for A as defined in Equation (6.3).

Note that estimating of off-diagonal entries only is also PromiseBQP-hard. This can be seen by replacing the original matrix A with

$$A' := A \otimes \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix},$$

which has obviously the same norm as A since the right-hand matrix is an orthogonal projector. Then we have

$$(A')^m = A^m \otimes \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}.$$

We obtain

$$(A^m)_{jj} = 2\langle j,0|(A')^m|j,1\rangle\,,$$

where we have used the short-hand notation $|j,i\rangle := |j\rangle \otimes |i\rangle$ for $i = 0,1$. Hence we have reproduced the diagonal entry of A^m by an off-diagonal entry of $(A')^m$.

7 Restriction to matrices with entries 0,1,-1

So far we have allowed for general real-valued matrix entries. We may strengthen the result of the preceding section in the sense that diagonal entry estimation remains PromiseBQP-hard if only the entries $0,\pm 1$ are possible.

It is known that Toffoli and Hadamard gates are universal for quantum computation [1] (in the sense of encoded universality) . For our purposes, the following modified universal set is useful. Let T and H denote the set of Toffoli gates and the set of Hadamard gates, respectively. We consider $T_{\text{left}} \cup T_{\text{right}} \cup H$, where we have defined $T_{\text{left}} := TH$ and $T_{\text{right}} := HT$. In words, T_{left} is, for instance, the set of gates that are obtained by applying an arbitrary Toffoli-gate followed by a Hadamard gate on an arbitrary qubit. One checks easily that all gates in the universal set $T_{\text{right}} \cup T_{\text{left}} \cup H$ have only entries $0,\pm 1/\sqrt{2}$. To construct our new version of the matrix A we replace the gates of Y_r with gates taken from our universal set. The problem is that we should simulate σ_z using an odd number of gates. Since it is by no means obvious whether and how this could be achieved we replace σ_z with a gate from $T_{\text{left}} \cap T_{\text{right}}$. The latter is a tensor product of a Hadamard and a Toffoli gate. The Hadamard gate acts on some extra qubit (prepared in the state $|0\rangle$) that is not used in the computation. The Toffoli gate copies the output to an additional qubit (this is done by initializing a third additional qubit to $|1\rangle$). One verifies easily that the unitary matrix W defined by these replacements acts as a shift in $2M$ dimensions whenever the output is 1. We obtain then a uniform mixture of the two spectral measures $P^{(0)}$ and $P^{(1)}$ defined in Section 6 instead of the measure $P^{(1)}$. To see what happens when the output is 0 we observe that the eigenvectors of the Hadamard gate define a decomposition of the initial state into two components. On the first component W acts as an M-dimenisonal shift and on the second as a M-dimensional shift with phase factor -1. Therefore we obtain also a mixture of the form $r_0 P^{(0)} + (1-r_0)P^{(1)}$. But now the weight r_0 is equal to $1/(4-2\sqrt{2}) = |\langle 0|\psi^+\rangle|^2$ where $|\psi^+\rangle$ is the eigenvector of the Hadamard gate for the eigenvalue 1. The difference between the diagonal entries of A^m for YES-instances and NO-instances is thus only reduced by a constant factor and the problem remains PromiseBQP-hard.

By rescaling with $\sqrt{2}$ we obtain a matrix A with entries $0, \pm 1$. The rescaling is clearly irrelevant for the diagonal entry estimation problem since we now have spectral values within the interval $[-\sqrt{2}, \sqrt{2}]$ and the accuracy required by Definition 4.1 changes by the factor $(\sqrt{2})^m$ accordingly.

8 Conclusions

We have shown that the estimation of diagonal entries of powers of symmetric sparse matrices is PromiseBQP-complete when the demanded accuracy scales appropriately with the powers of the operator norm.

The quantum algorithm proposed here for solving this problem uses the fact that measurements of the corresponding observable allow us to obtain enough information on the probability measure defined by the eigenvector decomposition of the considered basis state. Given the assumption that PromiseBQP$\neq$ PromiseBPP, i. e., that a quantum computer is more powerful than a classical computer, the required information on the spectral measure cannot be obtained by any efficient classical algorithm. This is remarkable since the determination of spectral measures is a problem whose relevance is not restricted to applications in quantum theory only.

Acknowledgements

P.W. was supported by the Army Research Office under Grant No. W911NF-05-1-0294 and by the National Science Foundation under Grant No. PHY-456720. This work was begun during D.J.'s visit to the Institute for Quantum Information at the California Institute of Technology. The hospitality of the IQI members is gratefully acknowledged. We would also like to thank Shengyu Zhang and Andrew Childs for helpful discussions.

References

[1] * D. AHARONOV: A simple proof that Toffoli and Hadamard are quantum universal. 2003. [arXiv:quant-ph/0301040]. 7

[2] * D. AHARONOV AND I. ARAD: The BQP-hardness of approximating the Jones polynomial. 2006. [arXiv:quant-ph/0605181]. 1, 3

[3] * D. AHARONOV, V. JONES, AND Z. LANDAU: A polynomial quantum algorithm for approximating the Jones polynomial. 2005. [arXiv:quant-ph/0511096]. 3

[4] * D. AHARONOV AND A. TA-SHMA: Adiabatic quantum state generation and statistical zero knowledge. In *Proc. 35th Annual ACM Symp. on Theory of Computing*, pp. 20–29, 2003. [STOC:10.1145/780542.780546]. 4, 5.3

[5] * E. BERNSTEIN AND U. VAZIRANI: Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. [SICOMP:10.1137/S0097539796300921]. 6

[6] * D. W. BERRY, G. AHOKAS, R. CLEVE, AND B. C. SANDERS: Efficient quantum algorithms for simulating sparse Hamiltonians. *Comm. Math. Phys.*, 270(2):359–371, 2007. [Springer:hk7484445j37r228]. 4, 5.3

[7] * D. E. BROWNE AND H. BRIEGEL: One-way quantum computation - a tutorial introduction. 2006. [arXiv:quant-ph/0603226]. 1

[8] * A. CHILDS: *Quantum information processing in continuous time*. PhD thesis, Massachusetts Institute of Technology, 2004. 4, 5.3

[9] * A. M. CHILDS, D. W. LEUNG, AND M. A. NIELSEN: Unified derivations of measurement-based schemes for quantum computation. *Phys. Rev. A*, 71:032318, 2005. [PRA:10.1103/PhysRevA.71.032318]. 1

[10] * S. EVEN, A. L. SELMAN, AND Y. YACOBI: The complexity of promise problems with applications to public-key cryptography. *Inform. and Control*, pp. 159–173, 1984. 2

[11] * M. FREEDMAN, A. KITAEV, AND Z. WANG: Simulation of topological field theories by quantum computers. *Comm. Math. Phys.*, 227(3):587–603, 2002. [Springer:btldwt3g5t0308da]. 3

[12] * O. GOLDREICH: On promise problems. Technical Report 18, Electr. Colloquium Computational Complexity, 2005. [ECCC:TR05-018]. 2

[13] * W. HOEFFDING: Probability inequalities for sums of bounded random variables. *Journ. Am. Stat. Ass.*, 58(301):13–30, 1963. 5.2

[14] * D. JANZING: Spin-1/2 particles moving on a 2d lattice with nearest-neighbor interactions can realize an autonomous quantum computer. *Physical Review*, A(75):012307, 2007. [PRA:10.1103/PhysRevA.75.012307]. 1

[15] * D. JANZING AND P. WOCJAN: Ergodic quantum computing. *Quant. Inf. Process.*, 4(2):129–158, 2005. [Springer:wq1g61v1236574t4]. 1

[16] * D. JANZING, P. WOCJAN, AND T. BETH: "Non-Identity check" is QMA-complete. *Int. Journ. Quant. Inf.*, 3(3):463–473, 2005. [doi:10.1142/S0219749905001067]. 6

[17] * J. KEMPE, A. KITAEV, AND O. REGEV: The complexity of the local Hamiltonian problem. *SIAM J. Computing*, 35(5):1070–1097, 2006. [SICOMP:10.1137/S0097539704445226]. 1, 6

[18] * A. KITAEV, A. SHEN, AND M. VYALYI: *Classical and Quantum Computation*. Am. Math. Soc., Providence, Rhode Island, 2002. 1, 6

[19] * E. KNILL AND R. LAFLAMME: Quantum computation and quadratically signed weight enumerators. *Inf. Process. Lett.*, 79(4), 2001. [IPL:10.1016/S0020-0190(00)00222-2]. 1, 3

[20] * M. NIELSEN AND I. CHUANG: *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 5.1, 5.1

[21] * R. OLIVEIRA AND B. TERHAL: The complexity of quantum spin systems on a two-dimensional square lattice. 2005. [arXiv:quant-ph/0504050]. 1, 6

[22] * R. RAUSSENDORF AND H. BRIEGEL: A one-way quantum computer. *Phys. Rev. Lett.*, 86(22):5188–5191, 2001. [PRL:10.1103/PhysRevLett.86.5188]. 1

[23] * P. WOCJAN, D. JANZING, TH. DECKER, AND TH. BETH: Measuring 4-local n-qubit observables could probabilistically solve PSPACE. In *Proc. Winter International Symp. on Information and Communication Technologies*, 2004. (Proceedings contain only the abstract). [arXiv:quant-ph/0308011]. 3, 6

[24] * P. WOCJAN AND J. YARD: The Jones polynomial: quantum algorithms and applications in quantum complexity theory. 2006. [arXiv:quant-ph/0603069]. 1, 3

[25] * P. WOCJAN AND S. ZHANG: Several natural BQP-complete problems. 2006. [arXiv:quant-ph/0606179]. 1, 3, 4, 6

AUTHORS

Dominik Janzing [About the author]
Fakultät für Informatik
Universität Karlsruhe (TH)
Am Fasanengarten 5
76 131 Karlsruhe
Germany
janzing@ira.ura.de
http://iaks-www.ira.uka.de/home/janzing/

Pawel Wocjan [About the author]
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando
FL 32816, USA
wocjan@cs.ucf.edu
http://www.eecs.ucf.edu/~wocjan/

ABOUT THE AUTHORS

DOMINIK JANZING studied physics and mathematics in Tübingen (Germany) and Cork (Ireland). His main interests were the foundations of quantum theory and thermodynamics. In 1998, he completed his Ph. D. thesis under the supervision of Manfred Wolff on the relation between quantum and classical dynamics in infinite quantum spin chains. When he joined the quantum computing group of Thomas Beth in the Faculty of Computer Science at the Universität Karlsruhe he did not expect that he would ever write a paper on complexity classes since his goal was "only" to understand the limits of quantum control. But while he was biking in the forests around Karlsruhe he realized that Pawel's and his results on the complexity of certain quantum measurements allow us to define a PromiseBQP-complete problem. When he was visiting Pawel at Caltech, the California sun gave both of them the strength to improve this idea. Dominik enjoys hiking, especially when his girlfriend Steffi joins him. He also likes fancy furniture.

PAWEL WOCJAN obtained his Ph. D. in CS from the University of Karlsruhe in 2003 under the supervision of Thomas Beth. In his Ph. D. thesis, entitled "Computational Power of Hamiltonians in Quantum Computing," he focused mainly on problems in quantum control theory such as designing efficient decoupling, time-inversion schemes, and Hamiltonian simulation schemes. As a postdoctoral scholar in CS at the Institute for Quantum Information at the California Institute of Technology from August 2004 till August 2006, he became more and more interested in classical and quantum complexity theory and the design of efficient quantum algorithms. He joined the School of Electrical Engineering at the University of Central Florida as an Assistant Professor in August 2006, where he continues his work on the computer-science challenges in quantum information science. He enjoys discussions with Dominik (not only about work!), spending time with his wife Ania, traveling to Europe to see his family and his friends, and many other things which make life so great.