

Method For Calculating
Mobius-like Transforms
Via a Quantum Computer

Robert R. Tucci

P.O. Box 226

Bedford, MA 01730

tucci@ar-tiste.com

March 24, 2014

CROSS REFERENCES TO RELATED APPLICATIONS

The following related patent applications are to be filed on the same day as this one:

- “Method For Calculating Symmetrized Functions Via a Quantum Computer”, by R.R. Tucci
- “Method For Calculating Mean Values Via a Quantum Computer”, by R.R. Tucci
- “Method For Discovering Structure of a Bayesian Network Via a Quantum Computer”, by R.R. Tucci

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH AND DEVELOPMENT

Not Applicable

REFERENCE TO COMPUTER PROGRAM LISTING

A computer program listing consisting of a single file entitled `ArQ-Src1-6.txt`, in ASCII format, is included with this patent application.

BACKGROUND OF THE INVENTION

(A) FIELD OF THE INVENTION

The invention relates to a quantum computer; that is, an array of quantum bits (called qubits). More specifically, it relates to methods for using a classical computer to generate a sequence of operations that can be used to operate a quantum computer.

(B) DESCRIPTION OF RELATED ART

Henceforth, we will allude to certain references by codes. Here is a list of codes and the references they will stand for.

Ref.Bra1 is G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, “Quantum amplitude amplification and estimation”, arXiv:quant-ph/0005055

Ref.Bra2 is G. Brassard, F. Dupuis, S. Gambs, and A. Tapp, “An optimal quantum algorithm to approximate the mean and its application for approximating the median of a set of points over an arbitrary distance”, arXiv:1106.4267

Ref.Dev is S. Devitt, Kae Nemoto, and W. Munro, “Quantum error correction for beginners”, arXiv:0905.2794.

Ref.Durr is C. Dürr, P. Hoyer, “A quantum algorithm for finding the minimum”, arXiv:quant-ph/9607014.

Ref.GPat is Lov K. Grover, “Fast Quantum Mechanical Algorithms”, US Patent 6,317,766

Ref.HeTi is Ru He, Jian Tian, “Bayesian Learning in Bayesian Networks of Moderate Size by Efficient Sampling”. Unpublished.

Ref.KeSm is R. Kennes, P. Smets, “Computational aspects of the Mobius transform”, arXiv:1304.1122

Ref.KoSo is M. Koivisto, and K. Sood, “Exact Bayesian structure discovery in Bayesian networks”, *The Journal of Machine Learning Research* 5 (2004): 549-573.

Ref.Tuc-qLis is R.R. Tucci, “qSym, qMobius, qMargi, qMean and qJennings, 5 Code Generators for Generating Quantum Circuits that Perform Some Artificial Intelligence Related Tasks”. Unpublished. Copy included as an appendix to this patent application.

Ref.Tuc-qMobius is R.R. Tucci, “Quantum Circuit for Calculating Mobius-like Transforms Via Grover-like Algorithm”. Unpublished. Copy included as an appendix to this patent application.

Ref.Tuc-qSym is R.R. Tucci, “Quantum Circuit for Calculating Symmetrized Functions Via Grover-like Algorithm”. Unpublished. Copy included as an appendix to this patent application.

Ref.TucAFGA is R.R. Tucci, “An Adaptive, Fixed-Point Version of Grover’s Algorithm”, arXiv:1001.5200

Ref.TucAFGApat is R.R. Tucci, “Method for Driving Starting Quantum State to Target One”, US Patent 8,527,437

Ref.TucQuibbs is R.R. Tucci, “Quibbs, a Code Generator for Quantum Gibbs Sampling”, arXiv:1004.2205

Ref.TucSimAnn is R.R. Tucci, “Code Generator for Quantum Simulated Annealing”, arXiv:0908.1633

This invention deals with quantum computing. A quantum computer is an array of quantum bits (qubits) together with some hardware for manipulating those qubits. Quantum computers with several hundred qubits have not been built yet.

However, once they are built, it is expected that they will perform certain calculations much faster than classical computers. A quantum computer follows a sequence of elementary operations. The operations are elementary in the sense that they act on only a few qubits (usually 1, 2 or 3) at a time. Henceforth, we will sometimes refer to sequences as products and to operations as operators, matrices, instructions, steps or gates. Furthermore, we will abbreviate the phrase “sequence of elementary operations” by “SEO”. SEOs for quantum computers are often represented by quantum circuits. In the quantum computing literature, the term “quantum algorithm” usually means a SEO for quantum computers for performing a desired calculation. Some quantum algorithms have become standard, such as those due to Deutsch-Jozsa, Shor and Grover. One can find on the Internet many excellent expositions on quantum computing.

This invention gives quantum circuits for calculating two closely related linear transforms that we refer to jointly as Mobius-like transforms. The first is the Mobius transform of a given function $f(x^n)$ where $x^n \in Bool$. The second transform is a trivial variation of the first: a marginal of a probability distribution $P(y^n)$, where $y^n \in Bool^n$.

Our algorithm utilizes the original Grover’s algorithm (see **Ref.GPat**) or any variant thereof, as long as it accomplishes the task of driving a starting state $|s\rangle$ towards a target state $|t\rangle$. However, we recommend to the users of our algorithm that they use a variant of Grover’s algorithm called AFGA (adaptive fixed point Grover’s algorithm) which was first proposed in **Ref.TucAFGA** and **Ref.TucAFGAp** .

Previously proposed algorithms for calculating the Mobius transform of a structureless function $f(x^n)$ use only a classical computer and take $\mathcal{O}(2^n)$ steps (see **Ref.KeSm** , **Ref.KoSo**). Our algorithm uses a quantum computer and it takes $\mathcal{O}(\sqrt{2^n})$ steps.

Another novel feature of our invention is that our quantum circuits use Grover’s algorithm in conjunction with two new techniques that we call “targeting two hy-

potheses” and “blind targeting”. When targeting two hypotheses, $|t\rangle$ is a superposition $a_0|0\rangle + a_1|1\rangle$ of two orthonormal states or hypotheses $|0\rangle$ and $|1\rangle$. When targeting blindly, the value of $\langle t|s\rangle$ is not known a priori.

The technique of “targeting two hypotheses” can be used in conjunction with Grover’s algorithm or variants thereof to estimate (i.e., infer) the amplitude of one of many states in a superposition. An earlier technique by Brassard et al (**Ref.Bra1** , **Ref.Bra2**) can also be used in conjunction with Grover’s algorithm to achieve the same goal of amplitude inference. However, our technique is very different from that of Brassard et al. They try to produce a ket $|x^n\rangle$, where the bit string x^n encodes the amplitude that they are trying to infer. We, on the other hand, try to infer an amplitude $|a_1|$ by measuring the ratio $|a_1|/|a_0|$ and assuming we know $|a_0|$ a priori.

Another novel feature of this invention is that we show how to use our quantum circuit that calculates Mobius transforms to calculate the minimum value of a function. Previous papers (see **Ref.GPat** , **Ref.Bra1** , **Ref.Bra2** , **Ref.Durr**) have proposed algorithms for finding the minimum value of a function via Grover’s algorithm. Their methods do not use Mobius transforms, or blind targeting of two hypotheses, so they are very different to the method proposed in this patent.

BRIEF SUMMARY OF THE INVENTION

A preferred embodiment of the invention is qMobius, a computer program written in Java. A trivial variation of qMobius is qMargi, another computer program written in Java. Source code for qMobius1.6 and qMargi1.6 is included as an appendix to this patent. qMobius and qMargi are “code generators” for generating quantum circuits. The quantum circuits generated by qMobius can be used to calculate the Mobius transform of a given function. The quantum circuits generated by qMargi are very similar to the circuits generated by qMobius, and they can be used to calculate a marginal of a probability distribution.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG.1 shows a block diagram of a classical computer feeding data to a quantum computer.

FIG.2 shows equations that describe technique of targeting two hypotheses.

FIG.3 shows quantum circuit used by qMobius to generate starting state $|s\rangle$ used in AFGA.

FIG.4 shows equations that describe properties of starting state $|s\rangle$ generated by the qMobius circuit of **FIG.3**,

FIG.5 shows **Control Window** of qMobius.

FIG.6 shows quantum circuit used by qMargi to generate starting state $|s\rangle$ used in AFGA.

FIG.7 shows equations that describe properties of starting state $|s\rangle$ generated by the qMargi circuit of **FIG.6**,

FIG.8 shows **Control Window** of qMargi.

DETAILED DESCRIPTION OF THE INVENTION

This section describes in detail a preferred embodiment of the invention and other possible embodiments of the invention. For a more detailed description of possible embodiments of this invention, see **Ref.Tuc-qMobius** , **Ref.Tuc-qLis** , **Ref.TucAFGApat** and references therein.

A preferred embodiment of the invention is qMobius, a computer program written in Java. A trivial variation of qMobius is qMargi, another computer program written in Java. Source code for qMobius1.6 and qMargi1.6 is included as an appendix to this patent. qMobius and qMargi are “code generators” for generating quantum

circuits. The quantum circuits generated by qMobius can be used to calculate the Mobius transform of a given function. The quantum circuits generated by qMargi are very similar to the circuits generated by qMobius, and they can be used to calculate a marginal of a probability distribution.

FIG.1 is a block diagram of a classical computer feeding data to a quantum computer. Box **100** represents a classical computer. qMobius1.6 and qMargi1.6 software runs inside Box **100**. Box **100** comprises sub-boxes **101**, **102**, **103**. Box **101** represents input devices, such as a mouse or a keyboard. Box **102** comprises the CPU, internal and external memory units. Box **102** does calculations and stores information. Box **103** represents output devices, such as a printer or a display screen. Box **105** represents a quantum computer, comprising an array of quantum bits and some hardware for manipulating the state of those bits.

The remainder of this section is divided into 5 subsections. Subsection (A) describes the quantum circuit generated by qMobius. Subsection (B) describes qMobius’s user interface. Subsection (C) describes the quantum circuit generated by qMargi. Subsection (D) describes qMargi’s user interface. Subsection (E) discusses other possible embodiments of the invention.

(A)qMobius: Quantum Circuit

In this section, we describe the quantum circuit generated by qMobius. For a more detailed description of the circuit, see **Ref.Tuc-qMobius** .

The full algorithm utilizes the original Grover’s algorithm or any variant thereof, as long as the algorithm drives a starting state $|s\rangle$ to a target state $|t\rangle$. For concreteness, we will assume for the preferred embodiment of this invention that we are using a variant of Grover’s algorithm called AFGA, described in **Ref.TucAFGA** and **Ref.TucAFGApat** .

Consider FIG.2.

FIG.2 describes what we will call “targeting two hypotheses”. Targeting two

hypotheses is a trick that can sometimes be used when applying Grover’s original algorithm or some variant thereof. Sometimes it is possible to arrange things so that the target state is a superposition $a_0 |0\rangle + a_1 |1\rangle$ of two orthonormal states $|0\rangle$ and $|1\rangle$, so that if we know a_0 , we can infer a_1 , a type of hypothesis testing with 2 hypotheses. If the target state were just proportional to say $|0\rangle$, then its component along $|0\rangle$ would be 1 after normalization so one wouldn’t be able to do any type of amplitude inference.

Suppose z_0, z_1 are complex numbers and $|\chi\rangle$ is an unnormalized state that satisfy **201**. Define p and q by **202**.

Let μ, ν and ω label subsystems. Assume the states $|\psi_0\rangle_\mu$ and $|\psi_1\rangle_\mu$ are orthonormal, the states $|0\rangle_\nu$ and $|1\rangle_\nu$ are orthonormal, and the states $|0\rangle_\omega$ and $|1\rangle_\omega$ are orthonormal.

We want to use AFGA with a starting state given by **203** and a target state given by **204**.

It’s easy to check that **205** and **206** are true. $|t\rangle$ only appears in AFGA within the projection operator $|t\rangle\langle t|$, and this projection operator always acts solely on the space spanned by $|t\rangle$ and $|s\rangle$. But $|t\rangle\langle t|$ and $|0\rangle\langle 0|_\omega$ act identically on that space. Hence, for the purposes of AFGA, we can replace $|t\rangle\langle t|$ by $|0\rangle\langle 0|_\omega$. We will call $|0\rangle_\omega$ the “sufficient” target state to distinguish it from the full target state $|t\rangle_{\mu,\nu,\omega}$.

Recall that AFGA converges in order $1/|\langle t|s\rangle|$ steps. From the definitions of $|s\rangle$ and $|t\rangle$, one finds **207**.

Once system (μ, ν, ω) has been driven to the target state $|t\rangle_{\mu,\nu,\omega}$, one can measure the subsystem ν while ignoring the subsystem (μ, ω) . If we do so, the outcome of the measurements of ν can be predicted from the partial density matrix **208**. From this density matrix, one gets **209** and **210**.

At first sight, it seems that Grover-like algorithms and AFGA in particular require knowledge of $|\langle t|s\rangle|$. Next, we will describe a technique called “blind targeting” for bypassing that onerous requirement.

For concreteness, we will assume in our discussion below that we are using AFGA and that we are targeting two hypotheses, but the idea of this technique could be carried over to other Grover-like algorithms in a fairly obvious way.

According to **207**, when targeting two hypotheses, $|\langle t|s\rangle| = \sqrt{p}$. Suppose we guess-estimate p , and use that estimate and the AFGA formulas of **Ref. TucAFGA** to calculate the various rotation angles α_j for $j = 0, 1, \dots, N_{Gro} - 1$, where N_{Gro} is the number of Grover steps. Suppose N_{Gro} is large enough. Then, in the unlikely event that our estimate of p is perfect, as $j \rightarrow N_{Gro} - 1$, \hat{s}_j will converge to \hat{t} . On the other hand, if our estimate of p is not perfect but not too bad either, we expect that as $j \rightarrow N_{Gro} - 1$, the point \hat{s}_j will reach a steady state in which, as j increases, \hat{s}_j rotates in a small circle in the neighborhood of \hat{t} . After steady state is reached, all functions of \hat{s}_j will vary periodically with j .

Suppose we do AFGA with p fixed and with $N_{Gro} = (N_{Gro})_0 + r$ Grover steps where $r = 0, 1, \dots, N_{tail} - 1$. Call each r a “tail run”, so p is the same for all N_{tail} tail runs, but N_{Gro} varies for different tail runs. Suppose that steady state has already been reached after $(N_{Gro})_0$ steps. For any quantity Q_r where $r = 0, 1, \dots, N_{tail} - 1$, let $\langle Q \rangle_{LP}$ denote the outcome of passing the N_{tail} values of Q_r through a low pass filter that takes out the AC components and leaves only the DC part. For example, $\langle Q \rangle_{LP}$ might equal $\sum_r Q_r / N_{tail}$ or $[\max_r Q_r + \min_r Q_r] / 2$. By applying the SEO of tail run r to a quantum computer several times, each time ending with a measurement of the quantum computer, we can obtain values $P_r(0)$ and $P_r(1)$ of $P(0)$ and $P(1)$ for tail run r . Then we can find $\left\langle \sqrt{P(1)/P(0)} \right\rangle_{LP} = \langle |z_1|/|z_0| \rangle_{LP}$. But we also expect to know $|z_0|$, so we can use $\langle |z_1|/|z_0| \rangle_{LP} |z_0|$ as an estimate of $|z_1|$. This estimate of $|z_1|$ and the known value of $|z_0|$ yield a new estimate of $p = |z_1|^2 + |z_0|^2$, one that is much better than the first estimate we used. We can repeat the previous steps using this new estimate of p . Every time we repeat this process, we get a new estimate of p that is better than our previous estimate. Call a “trial” each time we repeat the process of N_{tail} tail runs. p is fixed during a trial, but p varies from trial to trial.

Let $\theta(S)$ equal 1 if statement S is true and 0 if it's false.

The goal of qMobius is to give a method whereby a user can calculate

$$f(x^n) = \sum_{x^{-n}} \theta(x^{-n} \leq x^n) f^-(x^{-n}),$$

with $x^n, x^{-n} \in Bool^n$, $f^-(x^{-n}) = |A^-(x^{-n})|^2$, and $A^-(x^{-n}) = \langle x^{-n} | \psi^- \rangle$. Using linearity, the type of functions $f^-(x^{-n})$ for which one can find, via our method, the Mobius transform, can be extended to functions which don't necessarily satisfy $\sum_{x^{-n}} f^-(x^{-n}) = 1$.

We will assume that we know how to compile $|\psi^- \rangle_{\alpha^n}$ (i.e., that we can construct it starting from $|0^n \rangle_{\alpha^n}$ using a sequence of elementary operations. Elementary operations are operations that act on a few (usually 1,2 or 3) qubits at a time, such as qubit rotations and CNOTS.) Multiplexor techniques for doing such compilations are discussed in **Ref.TucSimAnn** . If n is very large, our algorithm will be useless unless such a compilation is of polynomial efficiency, meaning that its number of elementary operations grows as $\text{poly}(n)$.

Next consider **FIG.3**.

Our preferred method for calculating the Mobius transform of $f^-(x^{-n})$ consists of applying AFGA using the techniques of targeting two hypotheses and blind targeting. When we apply AFGA, we will use a sufficient target $|0 \rangle_{\omega}$. All that remains for us to do to fully specify our circuit for calculating the Mobius transform of $f^-(x^{-n})$ is to give a circuit for generating $|s \rangle$. That is what **FIG.3** does. **FIG.3** assumes $n = 3$ for concreteness. That figure uses fairly standard quantum circuit notation. The H stands for Hadamard matrix, σ_X for the X Pauli matrix, $P_c = |c \rangle \langle c|$ for $c \in Bool$, etc. More detailed explanations of the symbols in **FIG.3** can be found in **Ref.Tuc-qMobius** and references therein. Note that every horizontal line of **FIG.3** is a qubit.

Let α^- include all alpha minus qubits in **FIG.3**. Let α include all alpha (without the minus superscript) qubits in **FIG.3**. Let β include all beta qubits in **FIG.3**.

Next consider FIG.4.

Assuming that the circuit of FIG.3 is correct, then that circuit will generate the state $|s\rangle$ given by 401, where $|\chi\rangle$ is an unnormalized state and where 402 through 405 are satisfied. (If there is some small mistake in the circuit of FIG.3, then we should be able to find that mistake in the future and make small amendments to FIG.3 with the goal of generating an $|s\rangle$ that satisfies 401).

(B)qMobius: User Interface

In this section, we describe qMobius's user interface. For a more detailed description of the interface, see **Ref.Tuc-qLis** .

(B1)Input Parameters

qMobius expects the following inputs:

n: The number of qubits. $x^n, x^{-n} \in Bool^n$.

the vector c^n : We desire to calculate $f(c^n)$, the Mobius transform $f(x^n)$ of $f^-(x^{-n})$ at $x^n = c^n \in Bool^n$.

a circuit that generates state $|\psi^-\rangle$: The state $|\psi^-\rangle_{\alpha^{-n}}$ acts on n qubits. The demonstration version of qMobius uses a trivial, inconsequential circuit for $|\psi^-\rangle$, but this can be changed easily by subclassing the class of qMobius that defines $|\psi^-\rangle$.

an estimate of $p = |\langle t|s\rangle|^2$

(B2)Output Files

qMobius outputs 3 types of files: a Log File, an English File and a Picture File.

A Log File records all the input and output parameters displayed in the **Control Window** (see section entitled "Control Window"), so the user won't forget them.

An English File gives an “in English” description of a quantum circuit. It completely specifies the output SEO. Each line in it represents one elementary operation, and time increases as we move downwards in the file.

A Picture File partially specifies the output SEO. It gives an ASCII picture of the quantum circuit. Each line in it represents one elementary operation, and time increases as we move downwards in the file. There is a one-to-one onto correspondence between the rows of corresponding English and Picture Files.

English and Picture Files are used in many of my previous computer programs. I’ve explained those files in detail in previous papers so I won’t do so again here. See, for example, **Ref.TucQuibbs** for a detailed description of the content of those files and how to interpret that content.

(B3)Control Window

FIG.5 shows the **Control Window** for qMobius. This is the main and only window of qMobius (except for the occasional error or advice message window). This window is open if and only if qMobius is running.

The **Control Window** allows the user to enter the following inputs:

File Prefix: Prefix to the 3 output files that are written when the user presses the **Write Files** button. For example, if the user inserts `test` in this text field, the following 3 files will be written:

- `test_qMob_log.txt` This is a Log File.
- `test_qMob_eng.txt` This is an English File
- `test_qMob_pic.txt` This is a Picture File.

Number of $|\psi\rangle$ Qubits: This equals n .

c vector radio buttons: These radio buttons allow the user to specify the vector $c^n = (c_j) \in Bool^n$. The j index of the radio buttons grows downward and is

indicated by **bit 0**, **bit 1**, **bit 2**, **bit 3**. For demonstration purposes, this Java applet only allows a maximum number of 4 $|\psi\rangle$ qubits. However, the applet is based on a class called `MobMain` which does not have these limitations.

The number of rows of radio buttons that are visible equals the number chosen in the **Number of $|\psi\rangle$ Qubits** menu.

Estimate of $|z_{-1}|^2/|z_{-0}|^2$: This equals the user's initial estimate of $|z_1|^2/|z_0|^2$.

Maximum Number of Grover Steps: `qMobius` will stop iterating the AFGA if it reaches this number of iterations.

Gamma Tolerance (degs): This is an angle given in degrees. `qMobius` will stop iterating the AFGA if the absolute value of γ_j becomes smaller than this tolerance. (γ_j is an angle in AFGA that tends to zero as the iteration index j tends to infinity. γ_j quantifies how close the AFGA is to reaching the target state).

Delta Lambda (degs): This is the angle $\Delta\lambda$ of AFGA, given in degrees.

The **Control Window** displays the following output text boxes.

$|z_{-0}|^2$: This equals $|z_0|^2$, the probability of the “null” hypothesis of the two hypotheses being targeted.

Starting Gamma (degs): This is γ_0 , the first γ_j , the γ_j for the first Grover iteration, given in degrees.

Final Gamma (degs): This is the γ_j for the final Grover iteration, given in degrees.

Number of Grover Steps: This is N_{Gro} , the total number of Grover iterations that were performed. It must be smaller or equal to the **Maximum Number of Grover Steps**. It will be smaller if the **Final Gamma (degs)** reached the **Gamma Tolerance (degs)** before the **Maximum Number of Grover Steps** was reached.

Number of Qubits: This is the total number of qubits for the output quantum circuit.

Number of Elementary Operations: This is the number of elementary operations in the output quantum circuit. Since there are no LOOPS in qMobius v1.6, this is the number of lines in the English File, which equals the number of lines in the Picture File.

(C)qMargi: Quantum Circuit

In this section, we describe the quantum circuit generated by qMargi. For a more detailed description of the circuit, see **Ref.Tuc-qMobius** . As we shall see, the quantum circuit generated by qMargi is a trivial variation of the one generated by qMobius.

FIG.1 and FIG.2, and everything we've said about them, applies to qMargi just like it does to qMobius.

The goal of qMargi is to give a method whereby a user can calculate

$$P(x^{n_0}) = \sum_{x^{-n}} \theta(x^{n_0} = x^{-n_0})P(x^{-n}),$$

where $n > n_0 > 0$, $x^{-n} = (x^{-(n-n_0)}, x^{-n_0}) \in Bool^n$, $P^-(x^{-n}) = |A^-(x^{-n})|^2$, and $A^-(x^{-n}) = \langle x^{-n} | \psi^- \rangle$.

As with qMobius, we will assume for qMargi that we know how to compile $|\psi^- \rangle_{\alpha^n}$.

Next consider FIG.6.

Our preferred method for calculating the marginal probability distribution $P(x^{-n_0})$ consists of applying AFGA using the techniques of targeting two hypotheses and blind targeting. When we apply AFGA, we will use a sufficient target $|0 \rangle_{\omega}$. All that remains for us to do to fully specify our circuit for calculating the marginal $P(x^{-n_0})$ is to give a circuit for generating $|s \rangle$. That is what FIG.6 does. FIG.6

assumes $n_0 = 3$ for concreteness. Note that every horizontal line of FIG.6 is a qubit except for the thick line labelled $\alpha^{-(n-3)}$ which represents $n - 3$ qubits.

Let α^- include all alpha minus qubits in FIG.6. Let α include all alpha (without the minus superscript) qubits in FIG.6. Let β include all beta qubits in FIG.6. Note that in the qMobius case, the number of α^- , α , β qubits were all the same, whereas in this case, there are n α^- qubits but only 3 α and β ones.

Note that FIG.3 for qMobius and FIG.6 for qMargi are very similar.

Next consider FIG.7.

Assuming that the circuit of FIG.6 is correct, then that circuit will generate the state $|s\rangle$ given by 701, where $|\chi\rangle$ is an unnormalized state and where 702 through 705 are satisfied. (If there is some small mistake in the circuit of FIG.6, then we should be able to find that mistake in the future and make small amendments to FIG.6 with the goal of generating an $|s\rangle$ that satisfies 701).

(D)qMargi: User Interface

In this section, we describe qMargi's user interface. For a more detailed description of the interface, see **Ref.Tuc-qLis** .

(D1)Input Parameters

qMargi expects the following inputs:

n : The number of $|\psi^-\rangle$ qubits. $x^{-n} \in Bool^n$.

n_0 : The number of qubits in the marginalized probability distribution. Must have $n > n_0 > 0$.

the vector c^{n_0} : We desire to calculate $\sum_{x^{-(n-n_0)}} P(x^{-n})$ evaluated at $x^{-n_0} = c^{n_0} \in Bool^{n_0}$.

a circuit that generates state $|\psi^-\rangle$: The state $|\psi^-\rangle_{\alpha^{-n}}$ acts on n qubits. The demonstration version of qMargi uses a trivial, inconsequential circuit for $|\psi^-\rangle$,

but this can be changed easily by subclassing the class of qMargi that defines $|\psi^-\rangle$.

an estimate of $p = |\langle t|s\rangle|^2$

(D2)Output Files

Everything we said about qMobius output files applies almost verbatim to qMargi output files.

(D3)Control Window

FIG.8 shows the **Control Window** for qMargi. Everything we said about the qMobius **Control Window** applies almost verbatim to the qMargi **Control Window**. The main difference is that for the qMargi window, there is an extra input called **Number of marginal qubits**. In terms of the notation of section (D1), **Number of $|\psi\rangle$ qubits** equals n and **Number of marginal qubits** equals n_0 .

(E)Other Embodiments

In this section, we describe other possible embodiments of the invention.

A standard definition in the field of quantum computation is that a qu(d)it is a quantum state that belongs to a d dimensional vector space and a qubit is a qu(d)it with $d = 2$. In quantum error correction (see **Ref.Dev** for an introduction), one distinguishes between 2 types of qu(d)its, physical and logical. A logical qu(d)it consists of a number of physical qu(d)its. It goes without saying that the qu(d)its in the quantum circuit FIG.3 or FIG.6 (or variant thereof) can always be interpreted as logical qu(d)its, and additional gates can be added to FIG.3 or FIG.6 (or variant thereof) with the purpose of performing error correction.

Our method for calculating the Mobius transform of a function can be used for calculating the minimum value of a function as follows. Suppose $x^n \in Bool^n$, and

$E(x^n) > 0$ is the function we wish to minimize. Define a secondary function $D^-(\cdot)$ which is sharply peaked at the minimum of the function $E(\cdot)$. If $E(x^n)$ is minimum when $x^n = X^n$, then assume $D^-(x^n)$ is almost equal to the Kronecker delta function $\delta(x^n, X^n)$. Let $D(\cdot)$ denote the Mobius transform of $D^-(\cdot)$. Let's speak in terms of the decimal representation $x = dec(x^n)$ of the points $x^n \in Bool^n$. Call X the minimum of $E(x)$. Assume $n = 5$ for concreteness. The domain of the function D^- is $\{0, 1, \dots, 31\}$. Calculate $D(X_0)$ with $X_0 = 15$. If $D(15)$ is much smaller than 1, then that means that the peak X is in $\{16, 17, \dots, 31\}$ so set $X_1 = 23$, the midpoint of $\{16, 17, \dots, 31\}$. Otherwise, if $D(15)$ is close to 1, then that means that the peak X is in $\{0, 1, \dots, 15\}$ so set $X_1 = 7$, the midpoint of $\{0, 1, \dots, 15\}$. Repeating this procedure, one gets a finite sequence X_0, X_1, X_2, \dots that converges to the peak X . We are simply performing a binary search for X . Of course, for large n , this technique for finding minima is only useful if $|\psi_D^-\rangle$ (where $\sqrt{D^-(x^n)} = \langle x^n | \psi_D^- \rangle$) can be compiled into a SEO of poly(n) length.

Many papers (see, for example, **Ref.HeTi** and references therein) have proposed methods for discovering from data the graph (aka structure) of a classical Bayesian network. **Ref.HeTi** explains how some of those methods require that we perform sums $\sum_{pa_j \subset Mom}$ over all pa_j such that $pa_j \subset Mom$, where pa_j is the set of nodes which are parents to node j in a classical Bayesian network with n nodes, and where Mom is a “mother” set of nodes. As is well known (see **Ref.Tuc-qMobius** for an explanation if it's not obvious to the reader), the statement $x^{-n} \leq x^n$ and the statement $pa_j \subset Mom$ can be identified. Hence, sums over all pa_j such that $pa_j \subset Mom$ are examples of Mobius transforms which can be calculated using this invention.

For convenience, the quantum circuits generated by an embodiment of this invention may include gates that act on more than 3 qubits at a time. Such “fat” gates might be judged by some not to be elementary gates as defined earlier in this patent. However, such fat gates should be allowed inside the SEO's covered by this invention

for cases in which they are trivially expandable (TE) fat gates. By TE fat gates we mean, fat gates for which there are well known, expanding methods for replacing them by a sequence of gates that are strictly elementary, in the sense that they act on just one or two qubits at a time. Multi-controlled rotations and multiplexors are examples of TE fat gates. In fact, see the Java classes `MultiCRotExpander` and `MultiplexorExpander` and related classes included in the code listing appendix to this patent. These classes automate such expanding methods for multi-controlled rotations and multiplexors.

So far, we have described some exemplary preferred embodiments of this invention. Those skilled in the art will be able to come up with many modifications to the given embodiments without departing from the present invention. Thus, the inventor wishes that the scope of this invention be determined by the appended claims and their legal equivalents, rather than by the given embodiments.