Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise

John A. Smolin,* Jay M. Gambetta, and Graeme Smith†
IBM T.J. Watson Research Center, Yorktown Heights, New York 10598, USA
(Received 29 June 2011; revised manuscript received 27 September 2011; published 17 February 2012)

We provide an efficient method for computing the maximum-likelihood mixed quantum state (with density matrix ρ) given a set of measurement outcomes in a complete orthonormal operator basis subject to Gaussian noise. Our method works by first changing basis yielding a candidate density matrix μ which may have nonphysical (negative) eigenvalues, and then finding the nearest physical state under the 2-norm. Our algorithm takes at worst $O(d^4)$ for the basis change plus $O(d^3)$ for finding ρ where d is the dimension of the quantum state. In the special case where the measurement basis is strings of Pauli operators, the basis change takes only $O(d^3)$ as well. The workhorse of the algorithm is a new linear-time method for finding the closest probability distribution (in Euclidean distance) to a set of real numbers summing to one.

DOI: 10.1103/PhysRevLett.108.070502 PACS numbers: 03.67.Ac, 03.65.Wj, 05.40.Ca

As scientists, we are faced again and again with the problem of determining from imperfect data what ''really happened'' in an experiment. Generally, the more data one has the better the reconstruction of the true events. Even so, the view remains imperfect, so one typically tries only to determine what was the event most likely to have led to the observed data. When quantum mechanics is considered, the situation is harder on the experimentalist, since even the results of perfectly performed measurements may have probabilistic results.

Nevertheless, one can still determine the quantum state by performing many experiments on identically prepared systems and building up good statistics on the outcomes. If the set of experiments is *informationally complete* then the mixed state density matrix describing the system can be determined. This is called quantum state tomography [1–3]. A complete determination of the state would require an infinite number of perfect measurements, so instead we concentrate on finding the maximum-likelihood state consistent with the available data (cf. [4,5]).

We consider an informationally complete set of measurements, each performed many times on identically prepared systems. From the experimental outcomes, we would like to determine the quantum state that gives the observed results with highest probability. This can be a computationally intensive task. For the two qubit experiments of [6], conventional maximum-likelihood (ML) solving took more time than the experiments themselves. And it has been reported in [7] that the ML reconstruction for 8 qubits in [8] took *weeks* of computation. Our main result is a fast algorithm for reconstructing this state when the noise is Gaussian. For 8 qubits our algorithm runs in *seconds*.

The rest of the Letter is organized as follows: First, we show that the ML state reconstruction problem with Gaussian noise is equivalent to a least-squares minimization problem on quantum states. Next, we prove that the minimum takes a particularly simple form. Finally, we give a fast algorithm that finds this minimum explicitly and benchmark it against several other methods.

Reduction to density matrix minimization.—Observables in quantum mechanics are Hermitian operators with the expectation value of a Hermitian operator σ applied to a mixed state ρ given by $\mathrm{Tr}(\sigma\rho)$. We can represent the result of an imperfect measurement of such an expectation subject to additive Gaussian noise with variance v as a probability density function

$$p(m|\rho) = \frac{1}{\sqrt{2\pi v}} e^{-[m-\mathrm{Tr}(\sigma\rho)]^2/(2v)}. \tag{1}$$

At this point we have diverged from a fully general treatment of maximum-likelihood quantum state reconstruction starting from any positive operator valued measurement in favor of a special case. In exchange, we are able find a very fast solution method with a simple intuitive explanation. The case we consider is of significant practical importance. For example, such expectation values were measured in the superconducting qubit experiments described in [6]. In other cases, due to the central limit theorem, our method will quickly converge to the correct result as a function of the number of measurement results. And importantly, our treatment corresponds to how many physicists view quantum measurements.

Given an orthonormal Hermitian operator basis $\{\sigma_i\}_{i=1}^{d^2}$ on $d \times d$ matrices (with $\mathrm{Tr}[\sigma_i\sigma_j] = d\delta_{ij}$), and a particular set of measured values m_{ij} corresponding to the jth measurement result of the expectation value of σ_i on the ''true state'' ρ_0, we want to find the mixed state ρ, a trace 1 Hermitian matrix with only non-negative eigenvalues, maximizing the likelihood function

$$\mathcal{L} = \prod_{ij} p^{(i)}(m_{ij}|\rho) = \prod_{ij} \frac{1}{\sqrt{2\pi v_i}} e^{-[m_{ij}-\mathrm{Tr}(\sigma_i\rho)]^2/(2v_i)}$$

or

$$\mathcal{L} = \prod_i \left(\frac{1}{\sqrt{2\pi v_i}}\right)^{n_i} e^{-[m_i - \mathrm{Tr}(\sigma_i \rho)]^2/(2v)}. \quad (2)$$

Here n_i is the number of measurement results for the expectation of σ_i, $m_i = \sum_{j=1}^{n} m_{ij}/n_i$ is the average value of those results and we have chosen n_i so that $v = v_i/n_i$ is independent of i. Then, the same ρ that maximizes $\mathcal{L}$ will minimize the log likelihood function

$$\mathcal{L}_{\log} = \sum_i [m_i - \mathrm{Tr}(\sigma_i \rho)]^2. \quad (3)$$

Working in the operator basis of the $\{\sigma_i\}$'s is not convenient, but fortunately the distance is just the Hilbert-Schmidt, or 2-norm, which is basis independent. We show this here for completeness:

Lemma.—$\sum_i (m_i - r_i)^2 = \|\mu - \rho\|_2^2/d$ where $m_i = \mathrm{Tr}[\sigma_i \mu]$ and $r_i = \mathrm{Tr}[\sigma_i \rho]$.

Proof.—

$$\begin{aligned}\|\mu - \rho\|_2^2 = \mathrm{Tr}(\mu - \rho)^2 &= \mathrm{Tr}\left[\left(\sum_i (m_i - \mathrm{Tr}\rho\sigma_i)\sigma_i\right)^2\right] \\ &= \sum_{ij} \mathrm{Tr}[(m_i - \mathrm{Tr}\rho\sigma_i)\sigma_i (m_j - \mathrm{Tr}\rho\sigma_j)\sigma_j] \\ &= \sum_{ij} (m_i - \mathrm{Tr}\rho\sigma_i)(m_j - \mathrm{Tr}\rho\sigma_j)\mathrm{Tr}[\sigma_i \sigma_j] \\ &= d\sum_i (m_i - \mathrm{Tr}\rho\sigma_i)^2 \quad (4)\end{aligned}$$

■

The matrix $\mu = (1/d)\sum_i m_i \sigma_i$ can be thought of as the experimentally noisy view of the density matrix ρ_0. Note that it is trace one by construction, but may have negative eigenvalues. Calculation of μ from the m_i's is a change of operator basis, and in general requires time $O(d^4)$ (there are d^2 values of i and each σ_i is a $d \times d$ matrix). This will actually be the limiting step in our overall algorithm, as all other steps will be $O(d^3)$ or better. In many cases of interest, however, the operator basis change can be done more quickly. This will be true whenever the matrices representing the σ_i's in the canonical basis are sparse. In particular, if the σ_i's are tensor products of the Pauli matrices

$$\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\} = \left\{\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}, \begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}, \begin{pmatrix}0 & -i\\ i & 0\end{pmatrix}, \begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix}\right\}$$

on n qubits, so that $d = 2^n$, each σ_i has only d nonzero elements and the change of basis can be carried out in $O(d^3)$ steps.

Now, after the change of basis from the m_i's to μ, our original maximum-likelihood problem has been transformed into the following:

Subproblem 1: Given a trace-one Hermitian matrix μ, find the closest density matrix ρ (a trace-one Hermitian matrix with only non-negative eigenvalues) under the 2-norm:

$$\|\mu - \rho\|_2^2 = \mathrm{Tr}[(\mu - \rho)^2] = \sum_{ij} |\mu_{ij} - \rho_{ij}|^2. \quad (5)$$

This is immediately familiar as a least-squares minimization problem, for which standard minimizer packages are well suited. Indeed, this is how the problem is often solved in practice. Unfortunately, finding the solution can be computationally intensive. In standard state reconstruction algorithms, this is the most expensive step by far (see Fig. 1).

Simple form for the minimum.—To improve matters, since the 2-norm is basis independent we can work in the eigenbasis of μ. The optimum ρ is diagonal in this basis. This is immediate from the form of (5), where any off diagonal terms can only contribute positive amounts to the sum. Thus, the problem reduces to finding the eigenvectors and eigenvalues of μ and picking the d non-negative eigenvalues for ρ minimizing (5). Eigensystem solvers are $O(d^3)$ and good packages exist [9].

We are left with a minimization over d variables, effectively the square root of the difficulty of the original problem. Call the eigenvalues of μ, ρ μ_i, λ_i and arrange them such that $\mu_i \geq \mu_{i+1}$. We now want to minimize $\sum_i (\lambda_i - \mu_i)^2$ such that $\sum_i \mu_i = \sum_i \lambda_i = 1$ and $\lambda_i \geq 0$. Using the method of Lagrange multipliers to impose this

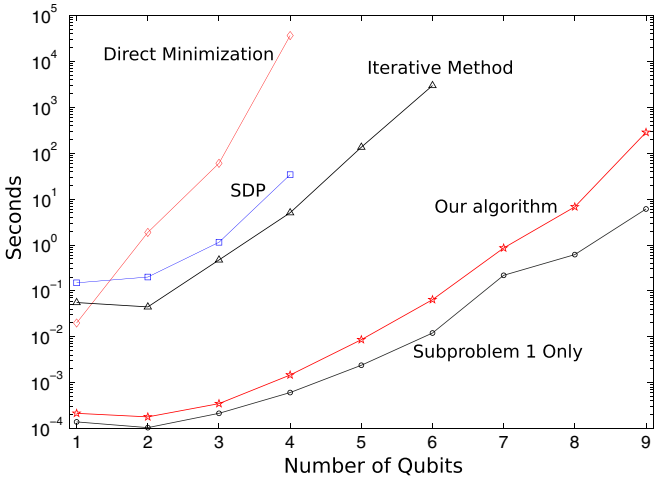

FIG. 1 (color online). Run time for maximum-likelihood reconstruction of random n-qubit pure states mixed with the identity and subjected to Gaussian noise on Pauli measurements. We compare four techniques: The diamond points are MATLAB's fminsearch minimizing $\mathrm{Tr}[(\mu - \rho)^2]$ directly. Squares points are timings for a semidefinite programming method (SeDuMi) [14]. Our implementation of the iterative method of [15] is shown with triangles. Circles indicate our algorithm for Subproblem 1. The stars denote our complete algorithm, including Subproblem 1 and the basis change from measurement outcomes to μ. All timings were performed in MATLAB on a single core of a 3 GHz Intel E8400 CPU.

constraint, and writing $\lambda_i = x_i^2$ to enforce non-negativity of λ_i we write the objective function

$$\Lambda = \sum_i (x_i^2 - \mu_i)^2 - L\left(\sum_i x_i^2 - 1\right). \tag{6}$$

Differentiating with respect to x_i we have

$$\frac{\partial \Lambda}{\partial x_i} = 4(x_i^2 - \mu_i)x_i - 2Lx_i = 0. \tag{7}$$

This equation has two solutions, either $x_i = 0$ or

$$x_i^2 = L/2 + \mu_i. \tag{8}$$

Note that L does not depend on i so each $\lambda_i = x_i^2$ is either set to zero or given by μ_i plus the very same number. To evaluate L we must pick a set $s = \{i \text{ such that } x_i \neq 0\}$. Then, summing (8) over i we have $1 = \sum_i x_i^2 = |s|L/2 + \sum_{i\in s}\mu_i$ or

$$L/2 = \frac{1}{|s|}\sum_{i\notin s}\mu_i \tag{9}$$

and

$$\Lambda = \frac{1}{|s|}\left(\sum_{i\notin s}\mu_i\right)^2 + \sum_{i\notin s}\mu_i^2. \tag{10}$$

Lemma.—Consider an i and j with $\mu_i > \mu_j$ with $i \in s$ and $j \notin s$. Then Λ_i, the distance function for this case, is always less than or equal to Λ_j, the distance function for a new set $\mathbf{s}' = s + \{j\} - \{i\}$.

Proof.—We can write

$$\Lambda_i = |s|\left(\frac{L}{2}\right)^2 + \sum_{k\notin s}\mu_k^2 \tag{11}$$

and for the case with $j \in s$ and $i \notin s$

$$\Lambda_j = |s|\left(\frac{L}{2} + \frac{\mu_i - \mu_j}{|s|}\right)^2 - \mu_j^2 + \mu_i^2 + \sum_{k\notin s}\mu_k^2. \tag{12}$$

If we had $\Lambda_j < \Lambda_i$ this would imply

$$L < \frac{\mu_j - \mu_i - |s|(\mu_i + \mu_j)}{|s|}. \tag{13}$$

Then we would have

$$\lambda_j = \mu_j + \frac{L}{2} + \frac{\mu_i - \mu_j}{|s|} < \frac{(|s| - 1)(\mu_j - \mu_i)}{2|s|} \leq 0 \tag{14}$$

because $\mu_i \geq \mu_j$. But λ_j must be non-negative, therefore Λ is never decreased by moving i into and j out of s. ■

The lemma tells us that all the λ_i's that are zero are together at the end, matching up with the smallest μ_i's. Thus, rather than the 2^d possible choices for s we need only decide where to put the break between zero and nonzero λ_i's, for which there are only d choices.

Next, we show that the choice of s should be the largest set satisfying the constraint that all the λ_i's are non-negative. Starting from Eq. (11), we imagine removing some element j from s. Then

$$\Lambda' = \frac{1}{|s| - 1}\left(|\mathbf{s}|\frac{L}{2} + \mu_j\right)^2 + \mu_j^2 + \sum_{ks}\mu_k^2 \tag{15}$$

and

$$\begin{aligned}\Lambda' - \Lambda_i &= \frac{1}{|s| - 1}\left(|s|\frac{L}{2} + \mu_j\right)^2 + \mu_j^2 - |\mathbf{s}|\left(\frac{L}{2}\right)^2 \\ &= \frac{|s|}{|s| - 1}\left[\left(\frac{L}{2}\right)^2 + 2\mu_j\frac{L}{2} + \mu_j^2\right] \\ &= \frac{|s|}{|s| - 1}\left(\frac{L}{2} + \mu_j\right)^2 \geq 0.\end{aligned} \tag{16}$$

In other words, setting any more of the λ_i's to zero than necessary increases the distance function. We are now ready to give an algorithm for Subproblem 1.

Fast algorithm for Subproblem 1.—(1) Calculate the eigenvalues and eigenvectors of μ. Arrange the eigenvalues in order from largest to smallest. Call these μ_i, $|\mu_i\rangle$, $1 \leq i \leq d$. (2) Let $i = d$ and set an accumulator $a = 0$. (3) If $\mu_i + a/i$ is non-negative, go on to step 4. Otherwise, set $\lambda_i = 0$ and add μ_i to a. Reduce i by 1 and repeat step 3. (4) Set $\lambda_j = \mu_j + a/i$ for all $j \leq i$. (5) Construct $\rho = \sum_i \lambda_i |\lambda_i\rangle\langle\lambda_i|$.

Figure 2 works through an example of this algorithm.

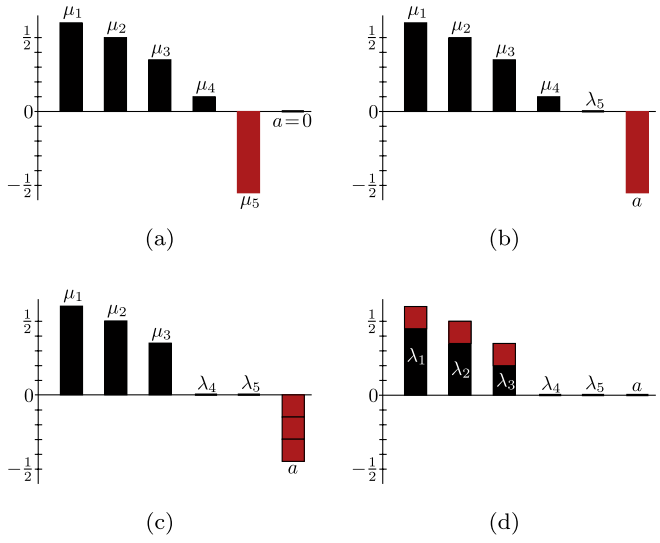

FIG. 2 (color online). Example of our algorithm for Subproblem 1: (a) We start with $\mu_1 = 3/5$, $\mu_2 = 1/2$, $\mu_3 = 7/20$, $\mu_4 = 1/10$, $\mu_5 = -11/20$, and accumulator $a = 0$. (b) Since $\mu_5 + a/5$ is negative, λ_5 is set to 0 and a to $-11/20$. (c) Since $\mu_4 + a/4$ is negative, λ_4 is set to 0 and $a = -9/20$. (d) Finally, λ_1, λ_2, and λ_3 each have $a/3 = -3/20$ added to the corresponding μ. The final result is $\lambda_1 = 9/20$, $\lambda_2 = 7/20$, $\lambda_3 = 1/5$, and $\lambda_4 = \lambda_5 = 0$.

Efficiency of the algorithm.—The slowest step is step 1, solving the eigensystem, which is $O(d^3)$ (standard libraries such as LAPACK are limited by the cost of reducing a Hermitian matrix to tridiagonal form using the Householder method which is $O(d^3)$ [10]). Step 2 is obviously constant time, and step 3 and 4 together are easily seen to be $O(d)$. Step 5 involves a choice of whether one wants the answer in the eigenbasis of μ, in which case it has already been computed, or in some other basis, in which case it is $O(d^3)$ or better [11]. Thus, the overall complexity is $O(d^3)$. The actual run time of an implementation will depend primarily on the eigensystem solver used. Including the basis transformation step, which is $O(d^4)$, renders this somewhat moot. Figure 1 compares the run time of our algorithm with that of a traditional ML optimization, with a semidefinite programming method [14], and with the iterative algorithm of [15].

Discussion.—As we have pointed out, in trade for great speed our method is not fully general, but corresponds well to many cases of interest. One can see from Fig. 1 that doing a full ML reconstruction may, in any event, be unfeasible on systems larger than just a few qubits. We also caution that ML methods should not be treated as a panacea: The most likely state may not be very likely. Much work is ongoing into modified ML methods and more general Bayesian approaches to characterizing the entire space of possible state reconstructions (cf. [16–18]).

Another less-than-general requirement of our argument that ML state reconstruction reduces to finding the closest density matrix under the 2-norm to a nonphysical candidate matrix μ is the demand that the effective variance $v = v_i/n_i$ of the measurements of the expectation values of $\mathrm{Tr}(\sigma_i\rho)$ in Eq. (2) not depend on i. Though this may not be the case in practice, often an experimentalist *can* chose to perform more of the noisier measurements, thus equalizing the variances. For example, in circuit QED (superconducting qubits) quantum nondemolition measurements of Paulis are performed with $v = 1/(\Gamma\tau)$ where Γ is the measurement rate and τ is the measurement time.

Aside from state reconstruction, it is often desired to perform *process tomography* [19], that is, to determine the quantum input-output relation (a trace-preserving completely positive map) implemented by an apparatus. Because of the Choi-Jamiolkowski isomorphism [20,21] such a map can be represented by a density matrix. Thus, the generalization of our result to process tomography should be straightforward.

Finally, we note a connection to classical probability theory. If one considers the eigenvalues μ_i (some of which may be negative) as a noisy view of a *probability distribution*, then the algorithm starting from step 2 is an algorithm for finding the nearest proper probability distribution. Furthermore, if the noise is Gaussian, this finds the maximum-likelihood probability distribution.

We thank J. Chow, A. Corcoles, and M. Steffen for valuable discussions, and A. Lvovsky and S. Merkel for sharing their code. We especially appreciate S. Glancy's discovery of an error in an earlier draft. J. A. S. and J. M. G. were supported by the IARPA MQCO program under Contract No. W911NF-10-1-0324.

*smolin@us.ibm.com
†gsbsmith@gmail.com

[1] M. Raymer, M. Beck, and D. McAlister, Phys. Rev. Lett. **72**, 1137 (1994).
[2] U. Leonhardt, Phys. Rev. Lett. **74**, 4101 (1995).
[3] D. Leibfried *et al.*, Phys. Rev. Lett. **77**, 4281 (1996).
[4] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Phys. Rev. A **64**, 052312 (2001).
[5] Z. Hradil, J. Rehacek, J. Fiurasek, and M. Jezek, Lect. Notes Phys. **649**, 59 112 (2004).
[6] J. M. Chow *et al.*, Phys. Rev. Lett. **107**, 080502 (2011).
[7] D. Gross *et al.*, Phys. Rev. Lett. **105**, 150401 (2010).
[8] H. Haffner *et al.*, Nature (London) **438**, 643 (2005).
[9] E. Anderson *et al.*, *LAPACK Users' Guide* (Society for Industrial and Applied Mathematics, Philadelphia, 1999), 3rd ed..
[10] A. S. Householder, J. Am. Chem. Soc. **5**, 339342 (1958).
[11] Basis transformation on a density matrix is accomplished by conjugation by unitaries, involving matrix multiplication. Matrix multiplication is $O(d^3)$ by the most naive method, but can be done in $O(d^{2.376})$ with a method that is beneficial only for impractically large d [12]. A more practical method sometimes employed for reasonable sized matrices (though which should used with caution since it is less numerically stable than the naive method) runs in $O(d^{\log_2 7}) \approx O(d^{2.807})$ [13].
[12] D. Coppersmith and S. Winograd, J. Symb. Comput. **9**, 251 (1990).
[13] V. Strassen, Numer. Math. **13**, 354 (1969).
[14] J. F. Sturm, Optimization Methods and Software **11**, 625 (1999).
[15] A. I. Lvovsky, J. Opt. B **6**, S556 (2004).
[16] J. Rehacek, D. Mogilevtsev, and Z. Hradil, New J. Phys. **10**, 043022 (2008).
[17] R. Blume-Kohout, Phys. Rev. Lett. **105**, 200504 (2010).
[18] R. Blume-Kohout, New J. Phys. **12**, 043034 (2010).
[19] I. L. Chuang and M. A. Nielsen, J. Mod. Opt. **44**, 2455 (1997).
[20] A. Jamiolkowski, Rep. Math. Phys. **3**, 275 (1972).
[21] M.-D. Choi, Linear Algebra Appl. **10**, 285 (1975).