Efficient Quantum Circuits for Schur and Clebsch-Gordan Transforms

Dave Bacon

Department of Computer Science and Engineering, University of Washington, Seattle, Washington 98195, USA
Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA
Santa Fe Institute, Santa Fe, New Mexico 87501, USA

Isaac L. Chuang

Center for Bits and Atoms, Research Laboratory for Electronics, Department of Physics, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, USA

Aram W. Harrow

Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
Department of Computer Science, University of Bristol, Bristol, BS8 1UB, United Kingdom

The Schur basis on n d-dimensional quantum systems is a generalization of the total angular momentum basis that is useful for exploiting symmetry under permutations or collective unitary rotations. We present efficient {size poly$[n, d, \log(1/\epsilon)]$ for accuracy ϵ} quantum circuits for the *Schur transform*, which is the change of basis between the computational and the Schur bases. Our circuits provide explicit efficient methods for solving such diverse problems as estimating the spectrum of a density operator, quantum hypothesis testing, and communicating without a shared reference frame. We thus render tractable a large series of methods for extracting resources from quantum systems and for numerous quantum information protocols.

A key component of quantum algorithms is their ability to reveal information stored in nonlocal degrees of freedom. In particular, one of the most important building blocks known is the quantum Fourier transform (QFT) [1], an efficient circuit construction for conversion between discrete position and momentum bases. The QFT converts a vector of 2^n amplitudes in $O(n^2)$ steps, in contrast to the $O(n2^n)$ steps required classically.

Another elementary basis change important in quantum physics is between independent local states and those of definite total generalized angular momentum. When two identical spin-1/2 particles interact with a global excitation, due to their permutation symmetry they appear as a singlet or a triplet to the external interaction. When this basis is generalized to n d-dimensional systems (n "qudits"), we call it the *Schur basis* and call the unitary transformation between local and Schur bases the *Schur transform*.

The Schur transform is central to a plethora of quantum information protocols and to many optimal physical methods for extracting information or resources from a quantum system. These include methods to estimate the spectrum of a density operator [2], perform quantum hypothesis testing [3], perform universal quantum source coding [4], concentrate entanglement noiselessly [5], create states immune to collective decoherence [6], and communicate without a shared reference frame [7]. For all of these tasks (and others), inefficient protocols also exist that work in local bases; however, only the protocols using the Schur basis are optimal. This suggests that the Schur basis is a natural

way to treat quantum states based on independent and identically distributed random variables, i.e., to experiments in which many copies of a single quantum state are given. However, unlike the QFT, no efficient algorithm for the Schur transform has been found, rendering protocols which use it nonconstructive. If we wish to implement the Schur transform in the lab to solve any of the problems listed above, an explicit efficient circuit construction for the Schur transform is needed.

Here, we resolve this problem by giving an efficient construction of the Schur transform on n qudits, for arbitrary n and d. This is achieved using a quantum circuit of size poly$[n, d, \log(1/\epsilon)]$ for accuracy ϵ. We believe that this basis change is important not only for quantum information and useful for extracting information about physical systems, but also as a new building block for future quantum algorithms.

The Schur transform.—Consider a system of n qudits, each with a standard local ("computational") basis $|i\rangle$, $i = 1 \ldots d$. The Schur transform relates transforms on the system performed by local d-dimensional unitary operations to those performed by permutation of the qudits. Recall that the symmetric group $\mathcal{S}_n$ is the group of all permutations of n objects. This group is naturally represented in our system by

$$\mathbf{P}(\pi)|i_1 i_2 \cdots i_n\rangle = |i_{\pi^{-1}(1)} i_{\pi^{-1}(2)} \cdots i_{\pi^{-1}(n)}\rangle, \quad (1)$$

where $\pi \in \mathcal{S}_n$ is a permutation and $|i_1 i_2 \ldots\rangle$ is shorthand for $|i_1\rangle \otimes |i_2\rangle \otimes \ldots$. Let $\mathcal{U}_d$ denote the group of $d \times d$ unitary operators. This group is naturally represented in

our system by

$$\mathbf{Q}(U)|i_1 i_2 \cdots i_n\rangle = U|i_1\rangle \otimes U|i_2\rangle \otimes \cdots \otimes U|i_n\rangle, \quad (2)$$

where $U \in \mathcal{U}_d$.

The Schur transform is based on Schur duality, a well-known [8] and powerful way to relate the representation theory of $\mathbf{P}(\pi)$ and $\mathbf{Q}(U)$. For example, consider the case of two qubits ($n = 2$, $d = 2$). The two-qubit Hilbert space $(\mathbb{C}^2)^{\otimes 2}$ decomposes under $\mathbf{Q}$ into a one-dimensional spin-0 singlet space spanned by $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$ and a three-dimensional spin-1 triplet space spanned by $|00\rangle$, $|11\rangle$, and $\frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$. Both of these spaces are acted upon in an irreducible manner under the action of $\mathbf{Q}(U)$, $U \in \mathcal{U}_2$, meaning that the action of $\mathbf{Q}(U)$ does not mix these two subspaces and these are the minimal such nonmixing subspaces which exist. Schur duality is related to the fact that these subspaces also happen to be irreducible representations (irreps) of $\mathcal{S}_2$. The singlet state changes sign under permutation of the two spins, and the triplet states are invariant under permutation. These correspond to the sign $\mathcal{P}_{sign}$ and the trivial $\mathcal{P}_{trivial}$ irreps of $\mathcal{S}_2$, and thus we can write $(\mathbb{C}^2)^{\otimes 2} \cong (\mathcal{Q}_1 \otimes \mathcal{P}_{trivial}) \oplus (\mathcal{Q}_0 \otimes \mathcal{P}_{sign})$, where $\mathcal{Q}_j$ is the spin-j irrep of $\mathcal{U}_2$.

This relation between the two representations exists for an arbitrary number of qudits, and in general both the $\mathcal{U}_d$ and $\mathcal{S}_n$ irreps will be nontrivial. For example, the Hilbert space of three qubits ($n = 3$, $d = 2$) decomposes into $(\mathcal{Q}_{3/2} \otimes \mathcal{P}_{trivial}) \oplus (\mathcal{Q}_{1/2} \otimes \mathcal{P}_{2,1})$, where $\mathcal{P}_{2,1}$ denotes a particular two-dimensional mixed symmetry irrep of S_3. In terms of the original (local) basis the $\mathcal{Q}_{1/2} \otimes \mathcal{P}_{2,1}$ space contains two spin-1/2 objects, one spanned by $|110\rangle + \omega|011\rangle + \omega^*|101\rangle$ (suppressing normalization) and $|001\rangle + \omega|100\rangle + \omega^*|010\rangle$, and the other obtained by replacing $\omega = e^{2\pi i/3}$ with ω^*. These two spaces correspond to the two dimensions of $\mathcal{P}_{2,1}$.

The general theorem of Schur duality states that for any (integer) d and n,

$$(\mathbb{C}^d)^{\otimes n} \cong \bigoplus_{\lambda \in \mathrm{Part}[n,d]} \mathcal{Q}_\lambda \otimes \mathcal{P}_\lambda, \quad (3)$$

where λ is chosen from the set of possible partitions of n into $\leq d$ parts, and simultaneously labels the $\mathcal{U}_d$-irrep $\mathcal{Q}_\lambda$ and the $\mathcal{S}_n$-irrep $\mathcal{P}_\lambda$. This goes beyond simultaneously diagonalizing the commuting representations $\mathbf{P}$ and $\mathbf{Q}$ because $\mathcal{P}_\lambda$ depends only on n (through λ) and not d. Schur duality means that there exists a basis for $(\mathbb{C}^d)^{\otimes n}$ with states $|\lambda, q_\lambda, p_\lambda\rangle_{Sch}$, where λ labels the subspaces $\mathcal{Q}_\lambda \otimes \mathcal{P}_\lambda$ and $|q_\lambda\rangle \in \mathcal{Q}_\lambda$ and $|p_\lambda\rangle \in \mathcal{P}_\lambda$ label bases for $\mathcal{Q}_\lambda$ and $\mathcal{P}_\lambda$, respectively.

Just as in the examples above, the Schur basis states $|\lambda, q_\lambda, p_\lambda\rangle_{Sch}$ are superpositions of the n qudit computational basis states $|i_1 i_2 \ldots i_n\rangle$,

$$|\lambda, q_\lambda, p_\lambda\rangle_{Sch} = \sum_{i_1, \ldots, i_n} [\mathbf{U}_{Sch}]^{\lambda, q_\lambda, p_\lambda}_{i_1, i_2, \ldots, i_n} |i_1 i_2 \cdots i_n\rangle. \quad (4)$$

By the isomorphism of Eq. (3), this defines a unitary transformation $\mathbf{U}_{Sch}$ (with matrix elements as given), the Schur transform we desire. If we think of $\mathbf{U}_{Sch}$ as a quantum circuit, it maps the state $|\lambda, q_\lambda, p_\lambda\rangle_{Sch}$ into the computational basis state $|\lambda, q_\lambda, p_\lambda\rangle$, with λ, q_λ, and p_λ expressed as bit strings. Since $\dim(\mathcal{Q}_\lambda)$ and $\dim(\mathcal{P}_\lambda)$ vary with λ we need to pad the $|q\rangle$ and $|p\rangle$ registers; this requires only constant spatial overhead. We know of no efficient classical algorithms to calculate even a single matrix element of $\mathbf{U}_{Sch}$, the best known results being recursive definitions of these matrix elements which require exponential time to evaluate [9]. The main purpose of this Letter is to show how the entire transformation can be performed on a quantum computer in poly(n, d) steps {implying as a corollary a classical algorithm for Schur transforming a vector of length d^n in time $O[d^n \mathrm{poly}(n, d)]$}.

The defining property of $\mathbf{U}_{Sch}$ is that it reduces the action of $\mathbf{Q}$ and $\mathbf{P}$ into irreps. For any $\pi \in S_n$ and any $U \in \mathcal{U}_d$, $\mathbf{P}(\pi)$ and $\mathbf{Q}(U)$ commute, so we can express both reductions at once as

$$\mathbf{U}_{Sch}\mathbf{Q}(U)\mathbf{P}(\pi)\mathbf{U}^\dagger_{Sch} = \sum_{\lambda \in \mathrm{Part}(d,n)} |\lambda\rangle\langle\lambda| \otimes \mathbf{q}_\lambda(U) \otimes \mathbf{p}_\lambda(\pi),$$
$$(5)$$

where $\mathbf{q}_\lambda$ and $\mathbf{p}_\lambda$ are irreps of $\mathcal{U}_d$ and $\mathcal{S}_n$, respectively.

Example of the Schur transform.—Consider the case of two qubits ($n = 2$, $d = 2$). Here the Schur transform is the transform between the standard computational basis $|i_1, i_2\rangle$ and a basis describing the singlet and triplet states. Explicitly the matrix of elements for the Schur transform, as in Eq. (4), are given by

$$
\begin{array}{c}
\begin{array}{cccc} |00\rangle & |01\rangle & |10\rangle & |11\rangle \end{array} \\
\begin{array}{l}
|\lambda = (1,1), q_\lambda = 0, p_\lambda = 0\rangle_{Sch} \\
|\lambda = (2,0), q_\lambda = +1, p_\lambda = 0\rangle_{Sch} \\
|\lambda = (2,0), q_\lambda = 0, p_\lambda = 0\rangle_{Sch} \\
|\lambda = (2,0), q_\lambda = -1, p_\lambda = 0\rangle_{Sch}
\end{array}
\begin{bmatrix}
0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 \\
1 & 0 & 0 & 0 \\
0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\
0 & 0 & 0 & 1
\end{bmatrix}.
\end{array}
$$
$$(6)$$

Here $\lambda = (1, 1)$ labels the singlet and $\lambda = (2, 0)$ labels the triplet. In this simple case, the permutation irreps are both one dimensional. Further as noted above, when we implement this we must express the label λ, p_λ, q_λ in terms of bit strings from some computational basis. For example, we could label λ by a single qubit and q_λ by two qubits (no qubits are required for p_λ in this example).

Applications of the Schur transform.—The numerous applications of the Schur transform mentioned in the introduction [2–7] solve a variety of problems which are relevant to quantum information theory as well as to experiments designed to acquire information or resources from a quantum system. Applying the Schur transform extracts λ, q, and p values for a given state, allowing the values be manipulated like any other quantum data. Here we briefly review a few of these applications, focusing on the ones most relevant to physics.

One application of the Schur transform is spectrum estimation [2]. In spectrum estimation, we are given access to n copies of a density operator $\rho = \sum_i p_i |i\rangle\langle i|$. Suppose the experimentalist wishes to estimate the values of the eigenvalues p_i but does not know the basis $|i\rangle$. Using the Schur transform is the optimal method for estimating this spectrum for any value of n. In particular, if we are given $\rho^{\otimes n}$, then performing the Schur transform on this state followed by measuring the irrep label λ provides an estimate of the spectrum by taking the partition $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_d)$ and dividing each λ_i by n: $p_i \approx (\frac{\lambda_1}{n}, \frac{\lambda_2}{n}, \ldots, \frac{\lambda_d}{n})$. In the limit of large n this estimate is optimal [2].

Consider, for example, spectrum estimation for the $n = 2$, $d = 2$ example given above. Let $\rho = p|0\rangle\langle 0| + (1 - p)|1\rangle\langle 1|$ be a fixed state and assume the experimentalist does not know the basis $|0\rangle$, $|1\rangle$. In order to estimate this spectrum if we are given two copies of ρ, we perform the Schur transform, Eq. (6), on these two qubits and measure the λ register. If we get $\lambda = (2, 0)$ we estimate that the spectrum is that of a pure state $p_1 = 1$, $p_2 = 0$ and if we get $\lambda = (1, 1)$ we estimate that the spectrum is that of the fully mixed state $p_1 = \frac{1}{2}$, $p_2 = \frac{1}{2}$. For a given value of p, the $\lambda = (2, 0)$ case occurs with probability $1 - p(1 - p)$, and the $\lambda = (1, 1)$ case occurs with probability $p(1 - p)$. Note that only if $p = 0$ or 1 does this estimate exactly reproduce the spectrum. Hence we learn a little about the spectrum with two copies of ρ; note that what we have learned is independent of the basis $|0\rangle$, $|1\rangle$. In the limit of a large number of copies, $n \gg 1$, the Schur transform provides the optimal estimate of the spectrum.

Another application of the Schur transform is to encode quantum information into noiseless subsystems which arise due to collective decoherence [6]. Here we run the Schur transform backwards (this can be done for the circuit by applying the inverse of every gate and reversing the order of the gates). If we input into the inverse Schur transform a fixed label $|\lambda\rangle$, some arbitrary information in the $|q_\lambda\rangle$ register, and the information we wish to encode in a noiseless manner into the $|p_\lambda\rangle$ basis, then the n qudit states output from this transform are encoded in a noiseless manner. In particular, the effect of decoherence which couples identically to each of the n qudit states acts trivially on the encoded information. Noiseless subsystems have already been implemented in ion trap quantum computers [10] and our transform makes feasible their use for larger systems.

As an example of the Schur transform in quantum information theory, consider the situation where Alice and Bob share n copies of a partially entangled state $|\psi\rangle_{AB} = \sum_i \sqrt{p_i} |i\rangle_A |i\rangle_B$ and they wish to extract the maximal number of maximally entangled states, $\frac{1}{\sqrt{d_E}} \times \sum_{k=1}^{d_E} |i\rangle_A |i\rangle_B$, from these n copies. Alice and Bob's local density matrices are invariant under permutations of their n copies, so if they perform the Schur transform and measure the $|\lambda\rangle$ basis, this leaves their $|p_\lambda\rangle$ registers in a maximally

entangled state. If $|\psi\rangle$ is unknown and no classical communication is allowed, then this is an optimal distortion-free entanglement protocol [5]. Note that in order to make this protocol computationally tractable, we need to describe how the $|p_\lambda\rangle$ basis states are labeled in a way that can be efficiently and reversibly mapped to the integers $\{1, \ldots, \dim(\mathcal{P}_\lambda)\}$ [11].

Quantum circuit for the Schur transform.—We construct a quantum circuit [12] for $\mathbf{U}_{\text{Sch}}$ in two stages, first for $d = 2$, then generalizing to $d > 2$. Each of these constructions follows an iterative structure, in which the Schur transform on n qudits is realized using n elementary steps, each of which adds a single qudit to an existing Schur state of the form $|\lambda, q, p\rangle$.

For $d = 2$, this elementary step corresponds to the addition of angular momentum, and the matrix elements of the unitary transform are known as Clebsch-Gordan (CG) coefficients [13]. In this case, λ and q can be conveniently denoted by half integers j and m (with $|m| \leq j \leq n/2$) which give the total angular momentum and the z-component of angular momentum, respectively. And in terms of j, the CG transform takes as input $|j, m\rangle$ and a single spin $|s = \pm 1/2\rangle$, and outputs a linear combination of the states $|j' = j \pm 1/2, m' = m + s\rangle$. The amplitudes of the linear combination are readily computed using the usual ladder operators for raising and lowering angular momenta [13]. In addition, however, we must distinguish between multiple distinct pathways which add up to give the same total j, as demonstrated by the three qubit example above. In fact, it is the permutation symmetry of these pathways which gives rise to $\mathcal{P}_j$, and thus we track the pathway with another output label $p = j' - j$.

Putting this together, we can define an elementary Clebsch-Gordan transform step $\mathbf{U}_{\text{CG}}$ as a rotation between two specific basis states,

$$\begin{bmatrix} |j'_-, m', & p = -\frac{1}{2}\rangle \\ |j'_+, m', & p = +\frac{1}{2}\rangle \end{bmatrix} = \begin{bmatrix} \cos\theta_{j,m'} & -\sin\theta_{j,m'} \\ \sin\theta_{j,m'} & \cos\theta_{j,m'} \end{bmatrix}$$
$$\times \begin{bmatrix} |j, m_+\rangle|s = -\frac{1}{2}\rangle \\ |j, m_-\rangle|s = +\frac{1}{2}\rangle \end{bmatrix}, \quad (7)$$

where $j'_\pm = j \pm 1/2$, $m_\pm = m' \pm 1/2$, and $\cos\theta_{j,m'} = \sqrt{\frac{j+m'+1/2}{2j+1}}$. $\mathbf{U}_{\text{CG}}$ can be realized with three gates in a quantum circuit, as shown in Fig. 1, using as one gate a controlled rotation about $\hat{y}$ by angle $\theta_{j,m'}$. This angle is computed using usual quantum and reversible circuit techniques [12] with error ϵ, using $\text{poly}[\log(1/\epsilon)]$ standard circuit elements.

The full Schur transform is implemented by cascading $\mathbf{U}_{\text{CG}}$ as shown in Fig. 2. The complexity of this circuit is thus $\mathcal{O}(n\text{poly}\log(1/\epsilon))$. We now claim that $|p_1, \ldots, p_n\rangle := |p\rangle$ labels a basis for $\mathcal{P}_j$. This follows from Eq. (3) and the fact that the $p_k = j_k - j_{k-1}$, $k = 1, \ldots, n$ are invariant under $\mathbf{Q}$, while j and m are invariant under $\mathbf{P}$. In fact, since j_k describes the action of $\mathcal{S}_k$ on the
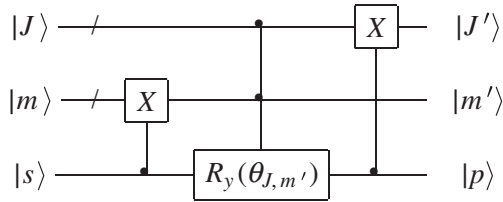
FIG. 1. Quantum circuit implementing $\mathbf{U}_{CG}$ to convert between the $|j, m\rangle|s\rangle$ and $|j', m', p\rangle$ bases, for the $d = 2$ (qubit) case. Following standard conventions [12], time goes from left to right, the $|j\rangle$ and $|m\rangle$ wires hold multiple qubits, and $|s\rangle$ is one qubit. The controlled X operation $\mathbf{C}_X$ adds the control to the target qubits, i.e., $\mathbf{C}_X|s\rangle|m\rangle = |s\rangle|m + s\rangle$. The doubly controlled $R_y(\theta_{j,m'})$ gate implements the rotation given by Eq. (7) using the j and m' qubits.

FIG. 2. Quantum circuit for the Schur transformation $\mathbf{U}_{Sch}$, transforming between $|i_1 i_2 \cdots i_n\rangle$ and $|j, m, p\rangle$. The fact that the $|p_1, p_2, \ldots, p_n\rangle$ is a full basis is intimately related to Schur duality.

first k qubits, $|p_1, \ldots, p_n\rangle$ is a subgroup-adapted basis for the chain $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \ldots \subset \mathcal{S}_n$, also known as Young's orthogonal basis [14]. This basis is also used in the only known fast quantum Fourier transform over $\mathcal{S}_n$ [14,15].

Construction of the Schur transform for $d > 2$ follows the same ideas as for $d = 2$, but is complicated by the challenge of showing that the elementary $\mathbf{U}_{CG}$ steps for $d > 2$ can be computed in poly(d) steps [a direct construction along the lines of Eq. (7) would require $n^{O(d)}$ steps]. $\mathbf{U}_{Sch}$ is constructed as a cascade of $O(n)$ $\mathbf{U}_{CG}$ transforms, just as for $d = 2$. Each $\mathbf{U}_{CG}$ combines a state $|\lambda, q_\lambda\rangle$ [with $\lambda \in \text{Part}(d, k - 1)$ and $|q_\lambda\rangle \in \mathcal{Q}_\lambda$] with a single qudit state $|i_k\rangle$, to obtain a superposition of states $|\lambda', q'_{\lambda'}\rangle$ [with $\lambda' \in \text{Part}(d, k)$ and $|q'_{\lambda'}\rangle \in \mathcal{Q}_{\lambda'}$]. Simultaneously, the permutation labels $|p\rangle$ are constructed; equivalently, we could save the values of λ that we generate in each step, just as $p_1, \ldots, p_n$ are equivalent to $j_1, \ldots, j_n$ for $d = 2$. $\mathbf{U}_{CG}$ can be computed efficiently because of a recursive relationship between $\mathbf{U}_{CG}$ for $\mathcal{U}_d \times \mathcal{U}_d$ and that of $\mathcal{U}_{d-1} \times \mathcal{U}_{d-1}$ in terms of reduced Wigner coefficients [16]. Crucially, there is an efficient classical algorithm for the computation of the reduced Wigner coefficients [9] needed for $\mathbf{U}_{CG}$. Specific details of this calculation are given in detail elsewhere [11]. The complexity of the full Schur transform is thus found to be polynomial in n, d, and $\log(\epsilon^{-1})$.

Conclusion.—We have shown how to efficiently perform the Schur transform. Without efficient implementations of the Schur transform, the various physical and quantum information tasks we have discussed [2–7] are not practical in the lab. As a final note, we comment on the Schur transform as it relates to the search for new quantum algorithms. An important open problem here is to find a black-box problem for which the Schur transform offers a speedup over classical algorithms. In this respect, there are few unitary transforms which have both an efficient quantum circuit and interpretations which might allow these transforms to be useful in an algorithm. We are hopeful that our circu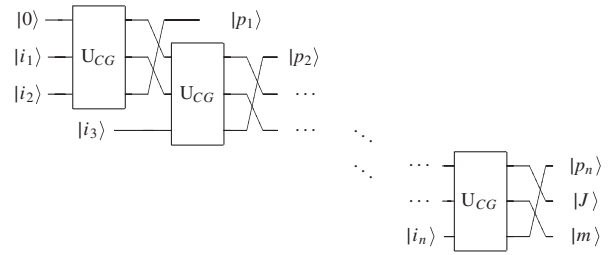its will be useful for quantum algorithms exactly because they have such clear group representation theory interpretations.

[1] P. W. Shor, in *Proceedings of the 37th Symposium on the Foundations of Computer Science* (IEEE, New York, 1996), pp. 56–65.
[2] M. Keyl and R. F. Werner, Phys. Rev. A **64**, 052311 (2001).
[3] M. Hayashi, J. Phys. A **35**, 10759 (2002).
[4] M. Hayashi and K. Matsumoto, Phys. Rev. A **66**, 022311 (2002); M. Hayashi and K. Matsumoto, Quantum Inf. Comput. **2**, 519 (2002).
[5] M. Hayashi and K. Matsumoto, quant-ph/0209030.
[6] E. Knill, R. Laflamme, and L. Viola, Phys. Rev. Lett. **84**, 2525 (2000); J. Kempe, D. Bacon, D. A. Lidar, and K. B. Whaley, Phys. Rev. A **63**, 042307 (2001).
[7] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, Phys. Rev. Lett. **91**, 027901 (2003).
[8] R. Goodman and N. Wallach, *Representations and Invariants of the Classical Groups* (Cambridge University Press, Cambridge, England, 1998).
[9] L. C. Biedenharn and J. D. Louck, Commun. Math. Phys. **8**, 89 (1968).
[10] D. Kielpinski *et al.*, Science **291**, 1013 (2001).
[11] D. Bacon, I. Chuang, and A. Harrow, quant-ph/0601001.
[12] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
[13] J. J. Sakurai, *Modern Quantum Mechanics* (Addison-Wesley, Reading, MA, 1994), Chap. 3, p. 214.
[14] R. Beals, in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (ACM Press, New York, 1997), pp. 48–53.
[15] C. Moore, D. Rockmore, and A. Russell, in *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms* (SIAM, Philadelphia, 2004), pp. 778–787.
[16] J. D. Louck, Am. J. Phys. **38**, 3 (1970).