Complexity of full counting statistics of free quantum particles in product states

Dmitri A. Ivanov [1] and Leonid Gurvits[2]
[1]*Institute for Theoretical Physics, ETH Zürich, 8093 Zürich, Switzerland*
[2]*Department of Computer Science, The City College of New York, New York, New York 10031, USA*

(Received 27 April 2019; published 7 January 2020)

We study the computational complexity of quantum-mechanical expectation values of single-particle operators in bosonic and fermionic multiparticle product states. Such expectation values appear, in particular, in full-counting-statistics problems. Depending on the initial multiparticle product state, the expectation values may be either easy to compute (the required number of operations scales polynomially with the particle number) or hard to compute (at least as hard as a permanent of a matrix). However, if we only consider full counting statistics in a finite number of final single-particle states, then the full-counting-statistics generating function becomes easy to compute in all the analyzed cases. We prove the latter statement for the general case of the fermionic product state and for the single-boson product state (the same as used in the boson-sampling proposal). This result may be relevant for using multiparticle product states as a resource for quantum computing.

DOI: 10.1103/PhysRevA.101.012303

I. INTRODUCTION

Future quantum computers are predicted to efficiently solve certain problems difficult for classical ones [1]. One indication of this "quantum supremacy" is the computational complexity of quantum amplitudes: computationally simple quantum states and operators may generate expectation values of higher complexity. This consideration lead to a quantum-computing proposal named "boson sampling" [2], where bosonic multiparticle amplitudes are given by (presumably) computationally difficult permanents [3]. In the boson-sampling proposal, the origin of the computational complexity of the corresponding noninteracting multiparticle amplitudes may be traced down to the quantum nature of the initial single-boson state. A similar construction with fermions would require suitably entangled fermionic states, in order to generate scattering amplitudes of the same complexity level [4,5].

Those examples suggest that we may benefit from a more systematic study of the complexity of expectation values for various classes of quantum states and operators. To some extent, this approach was already developed in the context of quantum optics [6], but we find it instructive to discuss the bosonic and fermionic cases on equal grounds. Specifically, we restrict our study to the computational complexity of matrix elements $\langle\Phi_1|\hat{U}|\Phi_2\rangle$, where $|\Phi_1\rangle$ and $|\Phi_2\rangle$ are multiparticle bosonic or fermionic states constructed as direct products of $N \gg 1$ identical states and $\hat{U}$ is a noninteracting multiparticle operator (e.g., a noninteracting evolution operator or a similar operator without the unitarity condition; we use a hat for noninteracting multiparticle operators, while the same letters without a hat denote the corresponding single-particle operators, as explained in Sec. II B); see Fig. 1(a). In this formulation, the states $|\Phi_i\rangle$ only require a finite number of parameters for their description (which is automatic in the fermionic case and implies an extra assumption for bosons) and the operator $\hat{U}$ is defined by the underlying single-particle operator U and is thus parametrized by $O(N^2)$ parameters. We are interested in a criterion for the matrix element $\langle\Phi_1|\hat{U}|\Phi_2\rangle$ to be computable in a polynomial in N time. In this paper, we only consider the problem of an exact computation and do not discuss the issue of approximations (the latter may be relevant for practical quantum-computing applications [6]).

We do not have a full answer to this question, but in this paper we collect a few known examples: some of them where a polynomial in N algorithm exists and others (specifically, the boson-sampling and entangled-fermion examples) that are (at least) as complex as a permanent (and therefore are believed to belong to a higher complexity class noncomputable in polynomial time).

After this overview of the known results for the general $\hat{U}$, we consider a variation of the problem where $\hat{U}$ is generated by a single-particle operator $U = 1 + V$ with V having a small rank. For the bosonic version of the problem with the single-product boson state (as in the boson-sampling construction), we find a polynomial algorithm thus proving Lemma B.5 [7] of Ref. [2] presented there without proof [about the polynomial computability of the permanent $\mathrm{Per}(1 + V)$]. A similar statement for the fermionic case is also formulated and proven. We also refine the original formulation of the lemma by proving an estimate for the degree of the polynomial: the number of the required operations is bounded by $O(N^{2k+1})$ in the bosonic case and $O(N^{2k})$ in the fermionic case.

The motivation for the above formulations comes partly from the full-counting-statistics (FCS) problems, where the generating function for the probability distribution of noninteracting particles has the described structure [8]. In particular, the results for the operators $\widehat{1 + V}$ with a finite-rank matrix V correspond to the computational complexity of the "marginal" FCS in a finite number of states (tracing over the remaining states). We elaborate on this interpretation in the corresponding section of the paper.

The paper consists of the three main parts. The first part introduces the multiparticle complexity of product states and

reviews, in this context, previously known results with only minor reformulations. This part includes Sec. II with definitions and notation and Sec. III with examples for the case of the general noninteracting operator $\hat{U}$. The second part of the paper is Sec. IV, where we prove new results for the more restrictive case $\hat{U} = \widehat{1+V}$. The third part of the paper, Sec. V, addresses motivation and interpretation of our constructions in terms of full counting statistics. Finally, in Sec. VI we summarize our results and propose questions for further studies.

II. DEFINITIONS AND NOTATION

A. Fermionic and bosonic states

We consider the fermionic and bosonic multiparticle spaces (Fock spaces) generated by a large number of single-particle levels, and we are interested in the computational complexity of matrix elements of a certain class of operators as a function of this number (whether it is polynomial or higher, e.g., exponential). More specifically, we restrict our analysis to product states: tensor products of states built on a small number (one or a few) of single-particle levels. For simplicity, in our discussion we consider all the states in these products to be identical, even though many of our results may also be extended to the case of products of different states. The following states will appear in our examples.

Single-boson product state:

$$|\mathrm{BN{=}1}\rangle^N = |\mathrm{BN{=}1}\rangle \otimes \cdots \otimes |\mathrm{BN{=}1}\rangle \quad (N \text{ times}), \tag{1}$$

where $|\mathrm{BN{=}1}\rangle = b^\dagger|\star\rangle_{B,1}$ is a single-boson state ($b^\dagger$ here and below denotes the boson creation operator and $|\star\rangle_{B,1}$ is the bosonic vacuum with one empty single-particle level). This is the state used in the boson-sampling proposal [2].

Coherent-boson product state:

$$|\mathrm{BC}{=}\alpha\rangle^N = |\mathrm{BC}{=}\alpha\rangle \otimes \cdots \otimes |\mathrm{BC}{=}\alpha\rangle \quad (N \text{ times}), \tag{2}$$

where $|\mathrm{BC}{=}\alpha\rangle = \exp(\alpha b^\dagger - \alpha^2/2)|\star\rangle_{B,1}$ is a coherent boson state.

Fermi-sea product state: a class of states constructed as

$$|\mathrm{FS}\rangle^N = |\mathrm{FS}\rangle \otimes \cdots \otimes |\mathrm{FS}\rangle \quad (N \text{ times}), \tag{3}$$

with $|\mathrm{FS}\rangle = \psi_1^\dagger \ldots \psi_k^\dagger|\star\rangle_{F,n}$, where $|\star\rangle_{F,n}$ is a fermionic vacuum with n single-particle states and $\psi_i^\dagger$ are creation operators for some $k \leqslant n$ (mutually orthogonal, for the sake of normalization) linear combinations of those states. n and k are fixed small numbers (unrelated to N). The product state $|\mathrm{FS}\rangle^N$ then belongs to the multiparticle space (Fock space) generated by Nn single-particle levels. Two particular cases of such a state are the vacuum state ($k = 0$) and the fully occupied state ($k = n$).

Entangled-quadruplet product state:

$$|\Psi_4\rangle^N = |\Psi_4\rangle \otimes \cdots \otimes |\Psi_4\rangle \quad (N \text{ times}), \tag{4}$$

where $|\Psi_4\rangle = (1/\sqrt{2})(f_1^\dagger f_2^\dagger + f_3^\dagger f_4^\dagger)|\star\rangle_{F,4}$. This state was used in Ref. [5]. It involves $2N$ fermions in $4N$ single-particle states.

B. Noninteracting operators

Every single-particle operator U generates a "multiplicative" multiparticle operator $\hat{U}$ in the multiparticle Fock space. A "physical definition" of this construction is sometimes written as

$$\hat{U} = \exp\left(\sum_{ij} a_i^\dagger (\ln U)_{ij} a_j\right), \tag{5}$$

where $a_i^\dagger$ and a_j are either fermionic or bosonic creation and annihilation operators. However, this definition formally fails when U has zero eigenvalues (noninvertible). For our purpose, we extend this definition to noninvertible matrices U, which can be done either by continuity or with a more explicit alternative definition

$$\hat{U} a_{j_1}^\dagger \ldots a_{j_k}^\dagger |\star\rangle = \sum_{i_1 \ldots i_k} U_{i_1,j_1} \ldots U_{i_k,j_k} a_{i_1}^\dagger \ldots a_{i_k}^\dagger |\star\rangle, \tag{6}$$

which describes the action of $\hat{U}$ on each of the basis vectors of the Fock space.

With this definition, we have a set of the "noninteracting operators" $\hat{U}$ defined as those obtainable from single-particle matrices U. This is a representation of the monoid of matrices U with respect to multiplication (i.e., $\widehat{U_1 U_2} = \hat{U}_1 \hat{U}_2$). In particular, this set is closed with respect to multiplication. An example of such an operator is a quantum evolution operator for a noninteracting system of particles [given by (5) with $\ln U$ playing the role of the Hamiltonian]. Another example motivated by full-counting-statistics problems is presented in Sec. V below.

C. Multiparticle complexity of a quantum state (or of a pair of states)

Now we are ready to define the main object of our study. We define the *multiparticle complexity* of a pair of states $|\Phi_1\rangle$, $|\Phi_2\rangle$ (or of a single state $|\Phi_0\rangle$) as the maximal computational complexity of the matrix element

$$\langle\Phi_1|\hat{U}|\Phi_2\rangle \quad (\text{or } \langle\Phi_0|\hat{U}|\Phi_0\rangle, \text{ respectively}) \tag{7}$$

(with the maximum taken over all noninteracting operators $\hat{U}$); see Fig. 1(a). The computational complexity is understood as scaling of the required number of operations as a function of N (see more explanations in Sec. II D below). The operator $\hat{U}$ is parametrized by its single-particle counterpart U, which requires N^2 parameters. The quantum states $|\Phi_i\rangle$, in their full generality, use an exponential number of amplitudes; therefore, the definition is only meaningful if we restrict it to a subclass of states parametrized by at most a polynomial in N set of parameters. One possible restriction of this sort is to consider product states (as defined in Sec. II A), where each of the factors involves only a finite number of parameters (we do not formalize this restriction further). Note that the operators U generally act across all the factors in the product state, which may make the matrix element (7) computationally demanding.

In Sec. IV below, we also consider a modification of this definition where $\hat{U}$ is further restricted to be generated by a matrix $U = 1 + V$, where V is a matrix of a finite rank. We will call this a *finite-rank complexity* of a state (or of a pair of states).

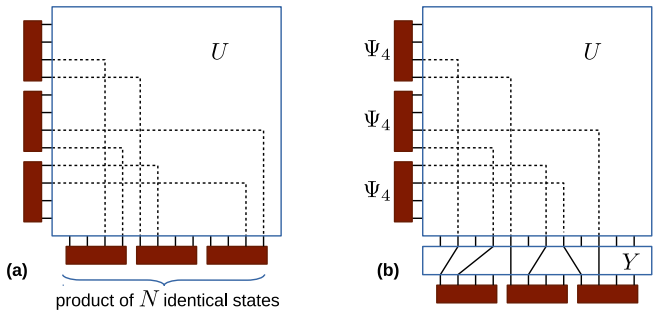

FIG. 1. (a) Schematic representation of the matrix element $\langle\Phi_0|\hat{U}|\Phi_0\rangle$ (or $\langle\Phi_1|\hat{U}|\Phi_2\rangle$, if the two states are different). The solid orange rectangles represent the factors of the product states $|\Phi_0\rangle$ (or $|\Phi_1\rangle$ and $|\Phi_2\rangle$). The big square represents the single-particle matrix U, with the dashed lines representing the matrix elements of U [there are $O(N^2)$ of such matrix elements]. (b) The construction used in Sec. III D for proving the equivalence of the hardness of the quantum amplitudes from Ref. [5] to the hardness in the definition of the present paper. Here, all the orange rectangles denote specifically the state Ψ_4. The lower rectangle represents the auxiliary operator Y and the solid lines across this rectangle denote the matrix elements of Y equal to 1 (with the rest of the matrix elements being zero).

D. Computational complexity for real and complex functions

Defining computational complexity for functions with continuous variables is sometimes a subtle issue [9], and we do not want to go deeply into this topic here. Instead, since the expectation values of interest are all polynomials of the matrix elements of U and of the wave-function components, we define the computational complexity as the scaling of the number of required arithmetic operations with N (with the exception of the coherent-boson case, which involves the exponentiation operator; see more details in Sec. III B). To simplify our notation, we only distinguish two levels of complexity: "easy" (computable in a polynomial in N number of operations) and "hard" (at least as difficult as computing a matrix permanent).

There is a general belief that computing a permanent requires a higher than polynomial number of operations, which implies $P \neq NP$ [3]. We also need this assumption in order for our classification to be meaningful. However, otherwise we never make use of it.

III. COMPLEXITY IN CASE OF GENERAL $\hat{U}$

We do not have a general criterion for product states to be "easy" or "hard", but we can give a few examples of states of each of them as follows.

(i) Single-boson product state is "hard".
(ii) Coherent-boson product state is "easy".
(iii) Fermi-sea product state is "easy".
(iv) Entangled-quadruplet product state is "hard".

A. Single-boson product state is "hard"

The corresponding expectation value is a permanent,

$$\langle \mathrm{BN}{=}1|^N \hat{U} |\mathrm{BN}{=}1\rangle^N = \mathrm{Per}\, U, \tag{8}$$

so it is "hard" by definition. This high complexity was used in Ref. [2] to conjecture the "quantum supremacy" of boson sampling.

B. Coherent-boson product state is "easy"

Since noninteracting operators $\hat{U}$ act within the space of coherent states (and this action can be written in single-particle terms), one can easily calculate the matrix element of $\hat{U}$ between any two coherent states. In particular,

$$\langle \mathrm{BC}{=}\alpha|^N \hat{U} |\mathrm{BC}{=}\alpha\rangle^N = \exp\left[\alpha^2\left(\sum_{ij} U_{ij} - N\right)\right]. \tag{9}$$

In this example, unlike in all the others, we use a sloppy definition of complexity: instead of the wave-function components (there are infinitely many of them), we use the parameter α of the coherent state and are allowed one exponentiation at the end of the calculation.

C. Fermi-sea product state is "easy"

The product of Fermi seas (3) is also a Fermi sea with Nk fermions. For this large Fermi sea, one easily finds

$$\langle \mathrm{FS}|^N \hat{U} |\mathrm{FS}\rangle^N = \det_{i,j} \langle\psi_i|U|\psi_j\rangle, \tag{10}$$

where the determinant is of the Nk-dimensional matrix of the single-particle matrix elements between the states generating the large Fermi sea. This proves that this matrix element is computable in polynomial time. Note that this argument equally applies to products of nonidentical Fermi seas.

D. Entangled-quadruplet product state is "hard"

This was shown in Ref. [5] (it also follows from the results on mixed discriminants in Refs. [10,11]). Strictly speaking, in that work, the hardness of $\langle x|\hat{U}|\Psi_4\rangle^N$ was proven, where $|x\rangle$ is a Fermi sea with arbitrary $2N$ states (orthogonal, for simplicity), $|x\rangle = \psi_1^\dagger \ldots \psi_{2N}^\dagger |\star\rangle_{F,4N}$. However, we can easily convert this statement into one for the expectation value in the state $|\Psi_4\rangle^N$. Namely, consider a single-particle operator Y transforming the states $\psi_1, \ldots, \psi_{2N}$ into the basis states $f_1, f_2, f_5, f_6, \ldots, f_{4N-3}, f_{4N-2}$ (in arbitrary order) and zeroing out the orthogonal complement of $\psi_1, \ldots, \psi_{2N}$. Then [see Fig. 1(b)]

$$\begin{aligned}\langle\Psi_4|^N \hat{Y}\hat{U}|\Psi_4\rangle^N &= \langle\Psi_4|^N \hat{Y}|x\rangle\langle x|\hat{U}|\Psi_4\rangle^N \\ &= 2^{-N/2}\langle x|\hat{U}|\Psi_4\rangle^N,\end{aligned} \tag{11}$$

which proves the hardness of the left-hand side of the above equation.

IV. FINITE-RANK COMPLEXITY

In this section, we consider the finite-rank complexity: a modified version of the complexity definition (Sec. II C), where the operators $\hat{U}$ are restricted to those generated by

$$U = 1 + V, \tag{12}$$

where V is a matrix of a finite rank:

$$V_{ij} = \sum_{s=1}^{k} u_i^{(s)} v_j^{(s)}. \tag{13}$$

Obviously, the finite-rank complexity cannot be higher than the complexity for the general $\hat{U}$. In particular, for all the examples considered above, the finite-rank complexity is "easy" (polynomial). Moreover, we can prove that the finite-rank complexity is polynomial for a general product state in the fermionic case. Specifically, we prove the following two statements below.

(i) The finite-rank complexity of the single-boson product state is "easy". We can further prove that the number of required operations scales as $O(N^{2k+1})$.

(ii) The finite-rank complexity of *any* fermionic product state is "easy". The number of required operations is also limited as $O(N^{2k})$.

A. Finite-rank complexity of the single-boson product state is "easy"

The matrix element is given by the permanent

$$\langle \mathrm{BN{=}1}|^N \hat{U} |\mathrm{BN{=}1}\rangle^N = \mathrm{Per}\, U = \mathrm{Per}(1+V). \tag{14}$$

Below we show that, if V has a finite rank k, the permanent (14) may be expressed in terms of the coefficients of an auxiliary polynomial of degree $2N$ in $2k$ variables, which, in turn, requires only a polynomial in N number of operations.

A simple combinatorial argument expresses the permanent (14) in terms of the vectors $u^{(s)}$ and $v^{(s)}$ from Eq. (13):

$$\mathrm{Per}(1+V) = \sum_{X\subseteq\{1,\ldots,N\}} \sum_{\substack{s_u(X),\, s_v(X) \\ [s_u(X)]=[s_v(X)]}} \prod_{x\in X} u_x^{(s_u(x))} v_x^{(s_v(x))} \prod_{r=1}^{k} n_r!, \tag{15}$$

where the first sum is taken over all subsets X of the set of indices $\{1,\ldots,N\}$, the second sum is over the label sets s_u and s_v (ranging from 1 to the rank k) for elements of X such that they form identical multisets (sets with repetitions) $[s_u(X)] = [s_v(X)]$, but possibly permuted with respect to each other. Finally, in the last product n_r denotes the multiplicity of r in the multiset $[s_u(X)]$ (or, equivalently, $[s_v(X)]$). This expression may, in turn, be computed with the help of the auxiliary polynomial of $2k$ formal variables

$$\begin{aligned} &F\left(a_u^{(1)},\ldots a_u^{(k)}, a_v^{(1)},\ldots,a_v^{(k)}\right) \\ &= \prod_{x=1}^{N}\left[1+\sum_{s=1}^{k}\sum_{s'=1}^{k} a_u^{(s)} a_v^{(s')} u_x^{(s)} v_x^{(s')}\right] \\ &= \sum_{\{n_r\},\{n'_r\}} F_{n_1,\ldots,n_k,n'_1,\ldots,n'_k} \\ &\quad\times \left(a_u^{(1)}\right)^{n_1}\ldots\left(a_u^{(k)}\right)^{n_k}\left(a_v^{(1)}\right)^{n'_1}\ldots\left(a_v^{(k)}\right)^{n'_k}, \end{aligned} \tag{16}$$

where the first equality is the definition of the polynomial $F(a_u^{(1)},\ldots a_u^{(k)}, a_v^{(1)},\ldots,a_v^{(k)})$ and the second equality is its expansion in powers of $a_u^{(s)}$ and $a_v^{(s)}$ defining its coefficients. On inspection, the "diagonal" coefficients of this polynomial reproduce the terms in the sum (15), up to combinatorial coefficients, and one finds

$$\mathrm{Per}(1+V) = \sum_{\{n_r\}} F_{n_1,\ldots,n_k,n_1,\ldots,n_k} \prod_{r=1}^{k} n_r!. \tag{17}$$

There are altogether $O(N^{2k})$ coefficients $F_{n_1,\ldots,n_k,n'_1,\ldots,n'_k}$, including $O(N^k)$ diagonal coefficients (with $n_r = n'_r$). Their calculation involves multiplying out N terms in Eq. (16), where at each multiplication the $O(N^{2k})$ coefficients need to be updated. Therefore, the calculation of $\mathrm{Per}(1+V)$ using Eqs. (16) and (17) can be done in $O(N^{2k+1})$ operations, as claimed. This proves Lemma B.5 [7] of Ref. [2] and the "finite-rank easiness" of the single-boson product state.

B. Finite-rank complexity of any fermionic product state is "easy"

The idea of the proof is that $\hat{U}$, in the finite-rank construction (12)–(13), acts nontrivially only in a small subspace spanned by a small number of fermionic states and therefore may be written in terms of a small number of fermionic operators. Specifically, $\hat{U}$ may be written in terms of the creation and annihilation operators defined as

$$\hat{u}_s^\dagger = \sum_i u_i^{(s)} f_i^\dagger, \quad \hat{v}_s = \sum_j v_j^{(s)} f_j, \tag{18}$$

where $f_i^\dagger$ and f_j are the fermionic creation and annihilation operators in the original basis. Using the definition (6), one can verify that $\hat{U}$ may be expressed as the polynomial in those operators,

$$\hat{U} = \sum_{\{s_i\}} \hat{u}_{s_1}^\dagger \ldots \hat{u}_{s_r}^\dagger \hat{v}_{s_r} \ldots \hat{v}_{s_1}, \tag{19}$$

where the sum is taken over all subsets $\{s_i\}$ of indices (including the empty subset, which contributes the unity operator) and r is the number of elements in the subset. The polynomial (19) has 2^k terms with its degree limited by $r \leqslant k$.

Now consider the expectation value of each term of the polynomial (19) in any product state

$$|\Phi_0\rangle = |\Psi_{(1)}\rangle \otimes \cdots |\Psi_{(N)}\rangle, \tag{20}$$

where each of the states $|\Psi_{(i)}\rangle$ belongs to a Fock space generated by a "small" (not growing with N) number of single-particle states (for our argument, we do not even need these states to be identical). Our states (3) and (4) are particular cases of this construction. Without loss of generality, we may take the states $|\Psi_{(i)}\rangle$ to be normalized.

To calculate the expectation value of a term of degree r in the polynomial (19) in the state $|\Phi_0\rangle$, we decompose each of the operators $\hat{u}_s^\dagger$ and $\hat{v}_s$ into N components:

$$\hat{u}_s^\dagger = (\hat{u}_s^\dagger)_1 \oplus \cdots \oplus (\hat{u}_s^\dagger)_N, \tag{21}$$

where $(\hat{u}_s^\dagger)_i$ acts in the ith space (hosting the state $|\Psi_{(i)}\rangle$). The same decomposition is done for the operators $\hat{v}_s$. Now we expand the product of $2r$ operators in Eq. (19) to obtain a sum N^{2r} terms. Each of these terms is itself a product of N factors by the number of subspaces in Eq. (20). These factors have the form $\langle\Psi_{(i)}|(\hat{u}_{j_1}^\dagger)_i \ldots (\hat{u}_{j_m}^\dagger)_i (\hat{v}_{j'_1})_i \ldots (\hat{v}_{j'_{m'}})_i |\Psi_{(i)}\rangle$: they

represent an expectation value in a small (not growing with N) space and therefore can be computable in a small number of operations. Moreover, at most $2k$ of those expectation values are nontrivial, and the rest are equal to one, since we have taken $|\Psi_{(i)}\rangle$ to be normalized. Therefore, the product of the N factors can actually be computed in a small (not growing in N) number of operations. Since there are N^{2r} such products for each term of degree $r \leqslant k$ in the polynomial (19), we can compute the expectation value $\langle\Phi_0|\hat{U}|\Phi_0\rangle$ in $O(N^{2k})$ operations. This proves our statement.

V. IMPLICATIONS FOR FULL COUNTING STATISTICS

A. Generating function for the particle-number probability distribution

The above discussion of the complexity of the expectation values may be interpreted in the language of so-called full counting statistics (FCS): a class of problems addressing the probability distribution of a quantum observable [8]. Namely, our results may be reformulated in terms of complexity of FCS generating functions for noninteracting particles initially prepared in a certain state $|\Phi_0\rangle$. Indeed, consider an initial state $|\Phi_0\rangle$ that is subject to a noninteracting evolution $\hat{U}_0$. We may further define the generating function

$$\chi(\lambda_1, \ldots, \lambda_N) = \sum_{\{n_i\}} e^{i\sum_{i=1}^N \lambda_i n_i} P(n_1, \ldots, n_N), \tag{22}$$

where

$$P(n_1, \ldots, n_N) = |\langle n_1, \ldots, n_N|\hat{U}_0|\Phi_0\rangle|^2 \tag{23}$$

is the probability to observe the counts n_i in the single-particle states i after the evolution $\hat{U}_0$.

In fact, there are three commonly used formulations for the full-counting-statistics problem: one may be interested either in computing the probabilities (23) or in the generating function (22) or in sampling the probability distribution with a randomized algorithm. The translation between these three formulations may turn out to be computationally intensive in the case of large N. For the purpose of this paper, we only consider the problem of calculating the generating function (22) and do not discuss its connections to the other two formulations (except for a short remark at the end of Sec. V B).

The generating function (22) has the required structure $\chi(\lambda_1, \ldots, \lambda_N) = \langle\Phi_0|\hat{U}|\Phi_0\rangle$, where

$$U = U_0^{-1} e^{i\sum_{i=1}^N \lambda_i n_i} U_0 \tag{24}$$

and n_i is the single-particle projector on the state i (see Fig. 2).

Some of the parameters λ_i used for counting particles in different single-particle states may be set to zero: in this case, the corresponding particle numbers are simply ignored (a trace is taken over all the particle numbers). Alternatively, it is also possible to directly take the limit $\lambda_i \to i\infty$ (or, equivalently, $e^{i\lambda_i} \to 0$, which corresponds to projecting onto the states with zero occupancy of the corresponding single-particle state [in this case, the operator (24) is no longer unitary].

At the same time, the state $|\Phi_0\rangle$ may be allowed to span only a subset of the available single-particle input states. Together with the possibility to exclude some output states

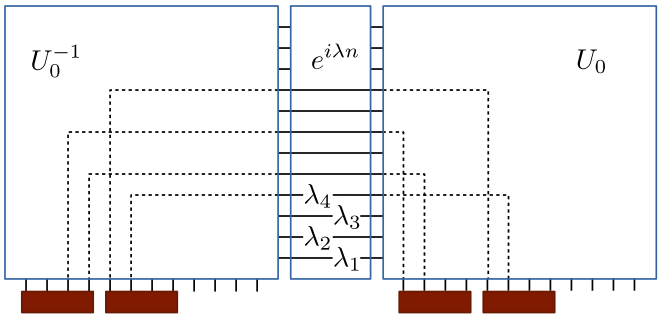

FIG. 2. Schematic illustration of the full-counting-statistics generating function $\chi(\lambda_1, \ldots, \lambda_N) = \langle\Phi_0|\hat{U}|\Phi_0\rangle$, where U is given by Eq. (24). The solid orange rectangles represent the factors of the product states $|\Phi_0\rangle$. The big squares are the factors U_0^{-1} and U_0 in Eq. (24). The rectangle in the middle is the factor $e^{i\sum\lambda_i n_i}$. The horizontal lines marked with λ_i correspond to general values of the parameters λ_i, solid lines to $\lambda_i = 0$, and missing lines to $\lambda_i \to i\infty$.

with the $\lambda_i \to i\infty$ limit, this provides a lot of flexibility for constructing the operator U, without the constraint of unitarity. Therefore, we conjecture that our results from Sec. III literally translate into the computational complexity of the generating function (22) for the general choice of the noninteracting evolution U_0 and of the complex variables λ_i.

B. Full counting statistics in a small subset of states

The discussion above applies to the case of the general choice of the parameters λ_i for a large (of order N) number of states. We may, however, consider a simpler problem with only a small number of nonzero λ_i (while keeping the total number of single-particle states and particles large, of order N). Then the matrix (24) has the form $1 + V$, where V has a finite rank. Indeed, we can rewrite (24) as

$$U = 1 + U_0^{-1}(1 - e^{i\sum_{i=1}^N \lambda_i n_i})U_0, \tag{25}$$

and the second term has a finite rank for a finite number of nonzero λ_i.

Therefore, our results of Sec. IV fully translate into the statement about full counting statistics. Namely, in our setup, for a finite subset of states, the FCS generating function is computable in polynomial time for the states $|\Phi_0\rangle$ considered in Sec. IV (single-boson product state and any fermionic product states).

Note that, for a small subset of states, there is no difference between the three different formulations of the FCS, since the generating function (22) and probabilities (23) are related by a Fourier transform computable in a polynomial (in N) time in this case. Furthermore, the sampling can also be performed in a polynomial time in our examples, since there is only a polynomial in N number of probabilities (23).

VI. SUMMARY AND DISCUSSION

The purpose of this paper is twofold. First, we introduce the notion of the "multiparticle complexity" of product states. This definition naturally leads to the question of formulating a criterion for a product state to be "hard". From examples, one may conjecture that most of such states are actually

"hard" except for a few special cases. One such special case is the so-called Gaussian states, where the Wick theorem applies (see, e.g., Ref. [12] for a definition in the bosonic case). Our examples of coherent-boson and Fermi-sea product states belong to this class of Gaussian states. One finds more Gaussian states among the mixed states described by a density matrix, but in this paper we restrict our discussion to pure states only. We do not know if there is any non-Gaussian state that would generate "easy" product states.

Another question in connection with this multiparticle complexity concept is its possible implications for quantum computing. Reference [2] suggests that this setup (specifically, the example of boson sampling) is insufficient for universal quantum computing, but, to our knowledge, without solid justification. In any case, it would be an interesting problem to characterize the class of problems solvable in polynomial time with this "full counting statistics" setup: an initial preparation of a certain product state (e.g., one of our "hard" examples), then evolution with a single-particle operator (which encodes the "quantum algorithm"), and finally a measurement of a certain generating function (or of a set of generating functions) (22). It seems plausible that all "hard" quantum states are equivalent for this quantum-computing setup, and therefore all of them would be equivalent to boson sampling.

In this context, we would like to comment on the relation of our multiparticle complexity examples to earlier studies on quantum computing with free bosons and fermions. In Ref. [13], computations with noninteracting operators on fermionic systems were shown to be "easy". This is consistent with our examples and with our conjecture above, since Ref. [13] only considered initial states in the form of "bitstrings", which fall in the category of Gaussian states (our Fermi-sea example in Sec. III C). It is the choice of the initial state that allows computationally hard expectation values in our examples. Note that Ref. [13] considered a more general form of quadratic operators possibly including pair creation and annihilation terms, while we only restrict our discussion to operators conserving the particle number. In Ref. [14], the authors have shown that free bosons may be used for efficient quantum computing. This is again consistent with our examples, since Ref. [14] uses single-boson states, which provide the key ingredient for creating quantum amplitudes of high computational complexity.

We would like to remind the reader that the "hardness" of a matrix element $\langle\Phi_0|\hat{U}|\Phi_0\rangle$ does not imply a possibility to actually compute this quantity with a quantum system. There are two reasons for this. First, the quantum measurement implies sampling, and achieving a good precision in a typically exponentially small expectation value would require exponentially many repeated measurements. Second, in this paper we only address the question of an exact computation, while for experimental implications it may be more relevant to study approximations. Computational complexity of approximate computations of permanents, mixed discriminants, and other related functions is addressed in many recent works [6,10,11,15,16].

The second goal of the paper is to report two results related to the "finite-rank" full counting statistics. For the cases we managed to prove (any fermionic product states and the single-boson product state), we have shown that counting particles in a *finite* number of final states is an "easy" task (computable in polynomial time). It seems plausible that this statement might be extended to a wider class of bosonic states (e.g., to any bosonic product states based on states with a finite number of particles). We leave this extension for future studies.

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).

[2] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, Theory Comput. **9**, 143 (2013).

[3] L. G. Valiant, The complexity of computing the permanent, Theor. Comput. Sci. **8**, 189 (1979).

[4] V. S. Shchesnovich, Boson-sampling with noninteracting fermions, Int. J. Quantum Inf. **13**, 1550013 (2015).

[5] D. A. Ivanov, Computational complexity of exterior products and multiparticle amplitudes of noninteracting fermions in entangled states, Phys. Rev. A **96**, 012322 (2017).

[6] S. Aaronson, A linear-optical proof that the permanent is #P-hard, Proc. Roy. Soc. A **467**, 3393 (2011).

[7] The same lemma has number 67 in the arXiv version of Ref. [2] and number 68 in the manuscript version on S. Aaronson's website.

[8] L. S. Levitov, H.-W. Lee, and G. B. Lesovik, Electron counting statistics and coherent states of electric current, J. Math. Phys. **37**, 4845 (1996).

[9] K. Weihrauch, *Computable Analysis, An Introduction* (Springer, Berlin, 2000).

[10] L. Gurvits, On the complexity of mixed discriminants and related problems, in *Mathematical Foundations of Computer Science* (Springer, Berlin, 2005), Vol. 3618, p. 447.

[11] L. Gurvits, On complexity of the mixed volume of parallelograms, in *Proceedings of the 25th European Workshop on Computational Geometry, Brussels, Belgium*, 2009, p. 337 (also available as report LA-UR-09-00709, Los Alamos National Laboratory, New Mexico, 2009, https://www.osti.gov/biblio/956384).

[12] X.-B. Wang, T. Hiroshima, A. Tomita, and M. Hayashi, Quantum information with Gaussian states, Phys. Rep. **448**, 1 (2007).

[13] B. M. Terhal and D. P. DiVincenzo, Classical simulation of noninteracting-fermion quantum circuits, Phys. Rev. A **65**, 032325 (2002).

[14] E. Knill, R. Laflamme, and G. J. Milburn, A scheme for efficient quantum computation with linear optics, Nature (London) **409**, 46 (2001).

[15] S. Aaronson and T. Hance, Generalizing and derandomizing Gurvits' approximation algorithm for the permanent, Quantum Inf. Comput. **14**, 541 (2014).

[16] N. Anari, L. Gurvits, S. Oveis Gharan, and A. Saberi, Simply exponential approximation of the permanent of positive semidefinite matrices, in *Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Berkeley, CA, 2017), p. 914.