Matrix-Product Operators and States: NP-Hardness and Undecidability

M. Kliesch,[1] D. Gross,[2] and J. Eisert[1]

[1]*QMIO Group, Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*Physikalisches Institut and Freiburg Center for Data Analysis and Modeling, Universität Freiburg, 79104 Freiburg, Germany*

Tensor network states constitute an important variational set of quantum states for numerical studies of strongly correlated systems in condensed-matter physics, as well as in mathematical physics. This is specifically true for finitely correlated states or matrix-product operators, designed to capture mixed states of one-dimensional quantum systems. It is a well-known open problem to find an efficient algorithm that decides whether a given matrix-product operator actually represents a physical state that in particular has no negative eigenvalues. We address and answer this question by showing that the problem is provably undecidable in the thermodynamic limit and that the bounded version of the problem is NP-hard (nondeterministic-polynomial-time hard) in the system size. Furthermore, we discuss numerous connections between tensor network methods and (seemingly) different concepts treated before in the literature, such as hidden Markov models and tensor trains.

Computational quantum many-body physics is marred by the fact that standard computational descriptions of states require exponentially many parameters. Fortunately, for many physically relevant problems, one does not need to consider all those parameters to capture natural properties extremely accurately. One of the pillars on which computational many-body approaches rest is the framework of tensor network methods. Here, the relevant degrees of freedom are parametrized by very few numbers, which are organized in terms of tensor networks that are contracted in order to compute expectation values [1–8]. Notably, the density-matrix renormalization group approach, the most successful method to numerically determine ground state properties of strongly correlated one-dimensional models, can be cast into such a form [1,2]. In this language, the problem of minimizing the energy can be phrased as a variational principle over matrix-product (or purely generated C^*-finitely correlated) states [9]. The natural analogue that also encompasses mixed quantum states is matrix-product operators. Again, they feature strongly in numerical algorithms [10,11], for example when investigating stationary states of local Liouvillians modeling open quantum systems [12,13] or Gibbs states [14,15].

However, general matrix-product operators are not guaranteed to represent physical states, which is the source of considerable conceptual and computational difficulties. It would thus be highly desirable to design an efficient algorithm capable of checking whether a given matrix-product representation defines a positive operator. To decide if such an efficient "local test for positivity" exists is a fundamental problem in the field, implicit already in its early formulations (see the Appendix of Ref. [9]).

Here, we address and answer this question: determining whether a matrix-product operator defines a physical state

in the thermodynamic limit is a provably undecidable problem. We also show that the bounded version of the problem is nondeterministic-polynomial-time (NP) hard in the number of tensors, burying hopes that one could find an efficient algorithm testing for positivity exactly. This is proven for quantum spin chains with local dimension $d = 2$ by a polynomial reduction from the Post correspondence problem and a bounded variant thereof.

To give a practical example, one can approximate stationary states of local Liouvillians by iteratively applying the Liouvillian to a state described as a matrix-product operator and subsequently truncating the tensors. To avoid inconsistent results, one has to check whether the truncation step has caused the state to become "too unphysical" in that it has created eigenvalues that are more negative than some chosen tolerance threshold. We prove such a check to be unfeasible. The practical implications of our work are as follows. On the one hand, they motivate the quest for finding specific feasible instances that might exist. This quest is reminiscent of the task of finding, e.g., efficient contractions of two-dimensional planar tensor networks, even though this task has been identified to be #P-complete [16]. On the other hand, it shows that one should direct one's efforts towards finding approximate solutions.

The insight presented here adds a natural many-body problem to the list of quantum mechanical questions that have recently been identified not only as computationally hard, but as outright undecidable [17–20]. Along the way of introducing these novel results, we discuss a number of connections between concepts that have arisen in the literature, but whose relation has received surprisingly little attention (see Table I).

Tensor networks: In quantum many-body theory, tensor network methods are widely used in order to avoid

TABLE I. Concepts of tensor networks discussed here.

States	Classical	Quantum
Pure	Deterministic finite [55]	Matrix-product states [3], purely generated C^*-finitely correlated states [9], tensor trains [26]
Mixed, inherently positive	Hidden Markov models [21], probabilistic finite automaton [46]	C^*-finitely correlated states [9], local purification [29], quantum Markov chains [27]
Mixed, not inherently positive	Quasirealizations [21]	Finitely correlated states [9], matrix-product density operators [10,11]

intractability problems. The idea is to resort to variational classes of states, where the attention is restricted to low-dimensional manifolds of states that seem to capture well the relevant physics of the model under study. It is less widely appreciated in the physics community that similar structures are ubiquitous in classical probability theory: the hidden Markov model (HMM) is a generalization of a Markov chain, where the observable process $(Y_t)_{t \in \mathbb{N}}$ does not need to be Markovian but there is a stochastic process $(X_t)_{t \in \mathbb{N}}$ carrying additional information that renders the combined process (X_t, Y_t) Markovian. We only consider the case where X_t and Y_t have finitely many outcomes and call the number of outcomes of X_t the bond dimension D. With transition probabilities $M_{i,j}^{(\alpha)} = \Pr[(X_{t+1}, Y_{t+1}) = (j, \alpha)|X_t = i]$, boundary condition $p_j = \Pr[X_1 = j]$, and $\underline{1} := (1, 1, ..., 1)^T \in \mathbb{R}^D$, the probabilities of outcome sequences of the process Y_t are given by the matrix product

$$\Pr[Y_1 = \alpha_1, ..., Y_n = \alpha_n] = p M^{(\alpha_1)} \cdots M^{(\alpha_n)} \underline{1}. \quad (1)$$

In order for a HMM to describe a stationary process, p is usually taken to be a stationary distribution, i.e., $\sum_{\alpha=1}^d M^{(\alpha)} p = p$. The description complexity of the HMM is independent of n, or, if we allow the Markov kernels M to vary as a function of t, linear. Non-negativity of the probabilities in Eq. (1) is guaranteed because they arise as the contraction over elementwise non-negative tensors.

From Eq. (1), it follows that the matrices $F^{(k,n)}$ defined by

$$F_{(\alpha_1,...,\alpha_k),(\alpha_{k+1},...,\alpha_{k+n})}^{(k,n)} := \Pr[Y_1 = \alpha_1, ..., Y_{k+n} = \alpha_{k+n}]$$

have rank at most D, which upper bounds the so-called Hankel rank [21].

A natural question is whether the rank condition alone characterizes those distributions that allow for a HMM with bond dimension D. It has been known since the 1960 s that this is not the case: there are distributions where $\text{rank}(F^{(k,n)}) \leq D$ for all k, n, yet no HMM with finite bond dimension exists [22,23]. However, a relatively straightforward argument (based on sequential "rank-revealing decompositions," e.g., singular value decompositions) shows that every distribution with rank bounded by D allows for a representation as in Eq. (1) where the tensors M

are not necessarily positive. This fact seems to have been discovered independently in different contexts, e.g., Refs. [7,9,24–26]. The resulting form is known as a quasirealization, offering the same concise description of the distribution as a HMM. These are, however, more difficult to work with computationally, as any variation of the local tensors can destroy global positivity. An important question thus is as follows: are the conditions on the tensor M that guarantee global positivity computationally efficiently verifiable? As we will see, the answer is no.

The above constructions generalize to the quantum setting: a C^*-finitely correlated state [9] (also known as quantum Markov chain [27,28]) ρ is obtained by replacing the elements of Eq. (1) by their quantum counterparts. We substitute p by a $D \times D$ density matrix σ, the stochastic map M by a completely positive map Φ that maps states on $\mathbb{C}^D$ to those on $\mathbb{C}^D \otimes \mathbb{C}^d$, and $\underline{1}$ by the partial trace (cf. Fig. 1). This immediately yields a local purification [29]: one can write Φ in Kraus representation by choosing operators $K_i : \mathbb{C}^D \to \mathbb{C}^D \otimes \mathbb{C}^d$ satisfying $\sum_{i=1}^E K_i K_i^\dagger = \mathbb{1}$ and $\Phi(\cdot) = \sum_{i=1}^E K_i^\dagger \cdot K_i$. Here, E can be assumed to be smaller than or equal to dD. Then the n-fold application of K to a purification $|\sqrt{\sigma}\rangle$ of σ yields the state vector $|\psi\rangle = K \circ K \circ ... \circ K |\sqrt{\sigma}\rangle$ in $\mathbb{C}^D \otimes (\mathbb{C}^d \otimes \mathbb{C}^E)^{\otimes n} \otimes \mathbb{C}^D$ that is a local purification of ρ, see Fig. 1. For quantum states one can, once more, define "quasirepresentations." Here Φ can be a general linear map and σ is some operator (no positivity constraints). This results in what is known as a matrix-product density operator (MPDO) or finitely correlated state (FCS) (not C^*-FCS). A discussion of different notions of positivity is provided in the Supplemental Material [30]. More concretely, with $[d] := \{1, 2, ..., d\}$,
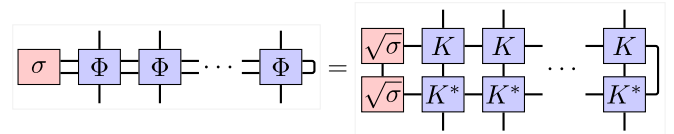

FIG. 1 (color online). A C^*-FCS as a tensor network. The channel Φ can be written in terms of Kraus operators, $\Phi(\rho) = \sum_i K_i^\dagger \rho K_i$. The vertically contracted indices between K and K^* correspond to the sum over i. The tensor network to the right is referred to as a local purification.

a MPDO is a density matrix written in the following form.

Definition 1 [matrix-product operator (MPO)].—An instance of MPO tensors is given by $M = (M_{i,j}^{(\alpha,\beta)})_{\alpha,\beta\in[d],i,j\in[D]} \in \mathbb{C}^{d\times d\times D\times D}$ and boundary vectors $|L\rangle, |R\rangle \in \mathbb{C}^D$. The dimension d is called the physical dimension and D the (MPO)-bond dimension. The generated translation invariant MPO for system size n is

$$\rho(L,M,R,n) := \sum_{j\in[D]^{n+1}} L_{j_1} M_{j_1,j_2} \otimes \ldots \otimes M_{j_n,j_{n+1}} R_{j_{n+1}}.$$

Main results: In order to precisely define the problems that are shown to be computationally unfeasible, we employ the standard language of theoretical computer science: the task of identifying objects with a certain property (e.g., those MPOs that are positive) is a decision problem. A specific case (e.g., given by a concrete tensor and boundary vectors) is an instance. A decision problem is NP-hard if it is at least as hard (in a precise sense) as all other problems from the complexity class NP. It is deemed highly unlikely that any NP-hard problem can be solved efficiently on either a classical or a quantum computer. A problem is (Turing) undecidable if no computer, even if endowed with unbounded resources, is capable of correctly solving all instances. In the statements of the various problems below, MPO tensors are specified by rational numbers. These have finite descriptions and can thus serve as inputs to computer programs. Allowing for more general numbers (e.g., complex rationals) would make the problem only harder.

In the precise statement of the problem, we allow for a threshold λ, which bounds the "degree of negativity" that is deemed acceptable. Moreover, we call positive semi-definite operators more concisely just positive.

Problem 2.—[bounded MPO threshold problem (BTP)] Instance: MPO tensors $M \in \mathbb{Q}^{d\times d\times D\times D}$, $|L\rangle, |R\rangle \in \mathbb{Q}^D$, threshold $\lambda \in \mathbb{Q}$, and system size n. Question: is the MPO $\rho(L,M,R,n) + \lambda\mathbb{1}$ positive?

Problem 3.—[MPO threshold problem (TP)] The TP is defined in the same way as the the BTP except that there is no restriction on the system size and the question is as follows: is there an $n \in \mathbb{Z}^+$ such that $\rho(L,M,R,n) + \lambda\mathbb{1}$ is not positive?

We obtain the following results, where the latter one adapts ideas from Ref. [46] to the quantum setting.

Theorem 4.—(NP-hardness of the bounded MPO threshold problem) For any $\lambda \in \mathbb{Q}$ and physical dimension $d \geq 2$, the BTP is NP-hard.

Theorem 5.—(undecidability of the MPO threshold problem) For each threshold $\lambda \in \mathbb{Q}$ the TP is undecidable. In particular, this holds for the case where the physical dimension is $d = 2$, the bond dimension is $D = 42$, and the matrices $M_{i,j}$ are diagonal for all $i, j = 1, \ldots, D$.

Outlook: An important question is whether there are physically relevant instances for which positivity is
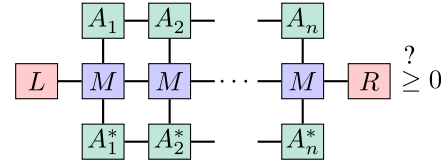

FIG. 2 (color online). Contracting MPOs with MPS can detect negativity for some instances. Hence, this provides a hierarchy of efficient tests labeled by the MPS-bond dimension, a strategy practically accessible by density-matrix renormalization group (DMRG) approaches.

efficiently decidable and how this can be exploited best in numerical algorithms. Sometimes one can, e.g., efficiently detect negativity locally by calculating expectation values with respect to matrix-product states (MPS) of small bond dimension, see Fig. 2.

In Ref. [29] local purifications of positive MPOs in terms of matrix-product states are investigated and it is shown that the arising MPS-bond dimension can in general not be bounded independently of the system size. This already suggests that such purifying MPS would require high bond dimensions when used instead of MPOs in numerical simulations. However, two constructive purification methods are suggested that are efficient when the rank of the MPO is polynomially bounded but in general necessarily inefficient [29]. From our Theorem 4 it also follows that this is no coincidence. To be more precise, a local purification method is an algorithm that receives a MPO instance $M, |L\rangle, |R\rangle$ and a system size n with $\rho(L,M,R,n) \geq 0$ as input and outputs a local purification of $\rho(L,M,R,n)$.

Corollary 6.—Local purification methods are inefficient in the system size.

In the BTP one is asked to exactly delineate the MPOs with smallest eigenvalues above $-\lambda$ from those with smallest eigenvalues below $-\lambda$. In practice, it would be acceptable if an algorithm reliably recognizes whether a state ρ is either sufficiently positive, i.e., $\rho \geq -\lambda$, or violates a threshold by at least $\epsilon \geq 0$, i.e., $\rho \ngeq -(\lambda + \epsilon)$. Such an approximate version is allowed to give unspecified results on the narrow band between the two cases. In order to make this precise, we state the BTP as a weak membership problem: for $\epsilon > 0$ one is only required to decide whether a MPO instance (L, M, R) with $\mathrm{Tr}(\rho(L,M,R,n)) = 1$ satisfies either $\rho(L,M,R,n) \geq -\lambda$ or $\rho(L,M,R,n) \ngeq -(\lambda + \epsilon)$. The MPO provided in the proof of Theorem 4 has a trace that is exponentially bounded from above. Hence, as a corollary, one obtains that the BTP remains NP-hard as a weak membership problem if ϵ is exponentially small in n. This statement remains true for algebraic and not necessarily rational inputs. Weak membership formulations seem to be natural for a variety of problems in quantum information. For instance, NP-hardness of testing separability of quantum states as a weak membership problem was established first [47] for an exponentially small "error" ϵ and, much later [48], for a polynomially

small ϵ, in fact, using key methods of the previous approach [47]. Hence, our work is an invitation to explore whether the BTP as a weak membership problem is also NP-hard for only polynomially bounded ϵ or, instead, to actually find an algorithm that efficiently solves that problem.

Details: For any finite set Σ (alphabet) we denote by Σ^n the set of sequences w (words) of $n = |w|$ elements (letters) from Σ and by $\Sigma^* := \bigcup_{n \in \mathbb{N}} \Sigma^n$ the set of words. For $w \in [d]^*$ we denote by $|w\rangle = \otimes_{j=1}^{|w|} |w_j\rangle$ the tensor product of the corresponding canonical basis states $|1\rangle, \ldots, |d\rangle$. Given two monoids W and W' we call a map $U : W \to W'$ a morphism if U maps the identity element of W to the identity element of W' and $U(w_1 w_2) = U(w_1) U(w_2)$ for all $w_1, w_2 \in W$. The monoids we encounter here are either given by words over an alphabet with concatenation or by matrices with matrix multiplication. Next, we introduce the famous Post correspondence problem [49] and a bounded variant thereof.

Problem 7.—[bounded Post correspondence problem (BPCP)] Instance: pairs of words $(u_\alpha, v_\alpha)_{\alpha \in [d]}$ over a finite alphabet Σ and length n in unary notation [50]. Question: does there exist a nonempty word $w \in [d]^n$ of length n such that $u_{w_1} u_{w_2} \cdots u_{w_n} = v_{w_1} v_{w_2} \cdots v_{w_n}$?

Problem 8.—[Post correspondence problem (PCP)] The PCP is defined in the same way as the BPCP except that there is no restriction on the word length.

The two sets of words $(u_\alpha, v_\alpha)_{\alpha \in [d]}$, referred to as dominos, define two morphisms $U, V : [d]^* \to \Sigma^*$ given by $U(w) = u_{w_1} u_{w_2} \cdots u_{w_{|w|}}$ and similar for V.

Theorem 9.—(PCP is undecidable, see Ref. [51]) For every $d \geq 7$ the PCP with $\Sigma = \{0, 1\}$ is undecidable.

In fact, by noting a simpler proof [52] of a variant of this theorem with larger d, one can make the following computation theoretic statement.

Observation 10.—(BPCP is NP complete) There is a polynomial p such that for any nondeterministic Turing machine $\mathcal{M}$ and input x there is a reduction to an instance U, V of the BPCP such that $\mathcal{M}$ accepts x in n steps iff there is a solution of U, V of length $p(n)$.

In the usual textbook proof of the undecidability of the PCP (see, e.g., Ref. [52]) the halting problem is reduced to the PCP. The idea of the proof is to encode the computation history of a given Turing machine into the two morphisms of the PCP in two different ways such that there is a solution iff the Turing machine halts. Specifically, there is a polynomial p such that exactly when a Turing machine accepts an input after n steps then there is a solution of the corresponding PCP instance of length $p(n)$. The encoding works in the same way when the Turing machine is replaced by a nondeterministic Turing machine because having a transition relation instead of a transition function allows us to define the PCP instances in the same way (possibly with more "dominos" (u_α, v_α)). This shows that if one could solve the BPCP in polynomial time, then one could also solve NP problems in polynomial time. As

solutions to the BPCP can be verified by a Turing machine in polynomial time the BPCP is in NP.

For the proofs of Theorems 4 and 5 the two following polynomial reductions are needed (see the Supplemental Material [30] for proofs building on Refs. [46,53,54]).

*Lemma 11.—*Let $(u_j, v_j)_{j \in [d]}$ be any instance of the PCP and $\lambda \in \mathbb{Q}$ be a threshold. Then there exist boundary vectors $|L\rangle, |R\rangle$, and matrices $A^{(1)}, \ldots, A^{(d)} \in \mathbb{N}^{6 \times 6}$ that define a morphism $A : [d]^* \to \mathbb{N}^{6 \times 6}$ such that for all $w \in [d]^*$

$$\langle L|A^{(w)}|R\rangle = -(\lambda + 1) \quad \text{if} \quad U(w) = V(w),$$
$$\langle L|A^{(w)}|R\rangle \geq -\lambda \qquad \text{if} \quad U(w) \neq V(w). \quad (2)$$

Lemma 12.—(see Ref. [54]) Let $d, D \geq 2$, $A^{(1)}, \ldots, A^{(d)} \in \mathbb{Q}^{D \times D}$ be matrices that define a morphism $A : [d]^* \to \mathbb{Q}^{D \times D}$, and $|L\rangle, |R\rangle \in \mathbb{Q}^D$ be boundary vectors. Then there exist two matrices $B^{(1)}$ and $B^{(2)}$ that define a morphism $B : [2]^* \to \mathbb{Q}^{Dd \times Dd}$ together with an injective morphism $X : [d]^* \to [2]^*$ satisfying $|X(w)| = d|w|$ such that $\langle \tilde{L}|B^{(X(w))}|\tilde{R}\rangle = \langle L|A^{(w)}|R\rangle$ for all $w \in [d]^*$, where $|\tilde{X}\rangle := (|X\rangle, 0, \ldots, 0)^T$.

*Proof of Theorems 4 and 5.—*We prove the theorem by using the encoding $A : [d]^* \to \mathbb{Q}^{D \times D}$ of the PCP with d dominos into the matrices from Lemma 11. Using Lemma 12 we reduce the physical dimension d to 2 at the expense of having a larger bond dimension dD and an increase of the system size n to dn. This results in an encoding $C : [2]^* \to \mathbb{Q}^{dD \times dD}$ with boundary vectors $|\tilde{L}\rangle$ and $|\tilde{R}\rangle$. Now we define a MPO tensor M by $M^{(\alpha, \beta)} = \sum_{\gamma=1}^d \delta_{\alpha, \gamma} \delta_{\beta, \gamma} C^\gamma$. Then $(|L\rangle, M, |R\rangle)$ is an encoding of the PCP to the TP. All successively used encodings are polynomial reductions. In particular, an instance of the BPCP with word length n can be written as an instance of the BTP with system size dn. Hence, Theorem 9 and Observation 10 finish the proof.

Conclusions: In this work, we have shown that a problem naturally occurring in the context of tensor network states is NP-hard and in the thermodynamic limit even undecidable. The findings point to the challenge for reliable numerical methods for, e.g., finding Gibbs and stationary states of quantum many-body systems: truncations in the bond dimension—a common step in existing numerical algorithms—can introduce inconsistencies that cannot be found computationally. This insight provides an interesting twist to numerical methods to capture mixed quantum many-body systems as well as to notions of Hamiltonian complexity [56,57]. Future research should follow a dual aim: first, identify instances and approximations where (near) positivity can be guaranteed; second, search for further problems in the context of tensor network states that are not decidable algorithmically.

[1] S. Rommer and S. Ostlund, Phys. Rev. B **55**, 2164 (1997).
[2] U. Schollwoeck, Ann. Phys. (N.Y.) **326**, 96 (2011).
[3] F. Verstraete, J. I. Cirac, and V. Murg, Adv. Phys. **57**, 143 (2008).
[4] N. Schuch, arXiv:1306.5551.
[5] R. Orus, Ann. Phys. (N.Y.) **349**, 117 (2014).
[6] J. Eisert, Model. Simul. **3**, 520 (2013).
[7] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Quantum Inf. Comput. **7**, 401 (2007).
[8] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).
[9] M. Fannes, B. Nachtergaele, and R. F. Werner, Commun. Math. Phys. **144**, 443 (1992).
[10] F. Verstraete, J. J. Garcia-Ripoll, and J. I. Cirac, Phys. Rev. Lett. **93**, 207204 (2004).
[11] M. Zwolak and G. Vidal, Phys. Rev. Lett. **93**, 207205 (2004).
[12] M. C. Banuls, "Tensor Network Techniques for the Study of Dissipative Dynamics," 2013 (unpublished).
[13] Z. Cai and T. Barthel, Phys. Rev. Lett. **111**, 150403 (2013).
[14] T. Barthel, U. Schollwöck, and S. R. White, Phys. Rev. B **79**, 245101 (2009).
[15] E. M. Stoudenmire and S. R. White, New J. Phys. **12**, 055026 (2010).
[16] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **98**, 140506 (2007).
[17] J. Eisert, M. P. Muller, and C. Gogolin, Phys. Rev. Lett. **108**, 260501 (2012).
[18] M. M. Wolf, T. S. Cubitt, and D. Perez-Garcia, arXiv:1111.5425.
[19] J. Morton and J. Biamonte, Phys. Rev. A **86**, 030301(R) (2012).
[20] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, in Proceedings of the XVII Conference on Quantum Information Processing, Barcelona, 2014 (unpublished).
[21] M. Vidyasagar, Math. Control Signals Syst. **23**, 1 (2011).
[22] M. Fox and H. Rubin, Ann. Math. Stat. **39**, 938 (1968).
[23] S. W. Dharmadhikari and M. G. Nadkarni, Ann. Math. Stat. **41**, 207 (1970).
[24] R. V. Erickson, Ann. Math. Stat. **41**, 843 (1970).
[25] G. Vidal, Phys. Rev. Lett. **93**, 040502 (2004).
[26] I. V. Oseledets, SIAM J. Sci. Comput. **33**, 2295 (2011); I. V. Oseledets and E. E. Tyrtyshnikov, SIAM J. Sci. Comput. **31**, 3744 (2009).
[27] M. Guta and J. Kiukas, arXiv:1402.3535.
[28] O. E. Barndorff-Nielsen, U. Franz, R. Gohm, B. Kümmerer, and S. Thorbjørnsen, *Quantum Independent Increment Processes II* (Springer, Heidelberg, 2006).
[29] G. De las Cuevas, N. Schuch, D. Pérez-García, and J. I. Cirac, New J. Phys. **15**, 123021 (2013).
[30] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.113.160503, which includes Refs. [31–45], for a discussion of different notions of positivity and proofs of Lemmas 11 and 12.
[31] S. A. Vavasis, SIAM J. Optim. **20**, 1364 (2010).
[32] D. D. Lee and H. S. Seung, Nature (London) **401**, 788 (1999).
[33] D. D. Lee and H. S. Seung, Adv. Neural Inf. Process. Syst. **13**, 556 (2000).
[34] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf, Proc. ACM STOC **44**, 95 (2012).
[35] M. Yannakakis, Proc. ACM STOC **20**, 223 (1988).
[36] J. Gouveia, P. A. Parrilo, and R. Thomas, Math. Oper. Res. **38**, 248 (2013).
[37] R. de Wolf, SIAM J. Comput. **32**, 681 (2003).
[38] C. Stark, arXiv:1210.1105.
[39] C. D. Aliprantis and R. Tourky, *Cones and Duality* (American Mathematical Society, Providence, 2007).
[40] S. Fiorini, S. Massar, M. K. Patra, and H. Raj Tiwary, arXiv:1310.4125.
[41] S. W. Dharmadhikari, Ann. Math. Stat. **34**, 705 (1963).
[42] J. Barrett, Phys. Rev. A **75**, 032304 (2007).
[43] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, arXiv:quant-ph/0611295.
[44] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North Holland, Amsterdam, 1982).
[45] D. Gross, M. Muller, R. Colbeck, and O. C. O. Dahlsten, Phys. Rev. Lett. **104**, 080402 (2010).
[46] V. D. Blondel and V. Catarini, Theory Comput. Syst. **36**, 231 (2003).
[47] L. Gurvits, Proc. ACM STOC **35**, 10 (2003).
[48] S. Gharibian, Quantum Inf. Comput. **10**, 343 (2010).
[49] E. Post, Bull. Am. Math. Soc. **52**, 264 (1946).
[50] "Unary notation" effectively means that the problem size of the input n equals n.
[51] Y. Matiyasevicha and G. Sénizergues, Theor. Comput. Sci. **330**, 145 (2005).
[52] M. Sipser, "*Introduction to the Theory of Computation*", 2nd ed. (Thomson Course Technology, Boston, 2005).
[53] M. Hirvensalo, Lect. Notes Comput. Sci. **4362**, 309 (2007).
[54] V. D. Blondel and J. N. Tsitsiklis, Inf. Proc. Lett. **63**, 283 (1997).
[55] J. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation* (Addison Wesley, Reading, MA, 2007).
[56] T. J. Osborne, arXiv:1106.5875.
[57] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, arXiv:1011.3445.