Complexity of the Quantum Separability Problem and Its Variants

Charles Xu

December 12, 2014

Abstract

We review a family of tightly related complexity-theoretic results introduced in two recent papers of Gutoski, Hayden, Milner, and Wilde. These pertain to variants of the quantum separability problem: roughly, deciding whether the output of a quantum circuit is separable or entangled. Fully specifying each problem means determining distance measures for completeness and soundness, whether to accept pure or mixed outputs, and whether to test for separable or product states. Most such variants have been shown to be complete for one of the interactive proof classes BQP, QMA, QMA(2), QSZK, and QIP; one is a candidate QIP(2)-complete problem. A wide variety of proof techniques come into play, some more than once, many relying on notions of k-extendibility and tailored separability and product testing protocols; we conclude by giving a flavor of some of the more important ones.

1 Introduction and Statement of the Problem

Entanglement is one of the most distinctive features of quantum mechanics, generally considered necessary for the full computational speedup it promises over classical models. It is also notoriously difficult to detect operationally in the presence of classical probabilistic correlations such as those in general mixed states, for which it is important to distinguish completely uncorrelated *product states* $\rho_{AB} = \rho_A \otimes \rho_B$ as a special subclass of the convex set of *separable states*

$$\rho_{AB} = \sum_i p_i \rho_A^{(i)} \otimes \rho_B^{(i)} \tag{1}$$

(Here and in the remainder of this survey, we restrict attention to the bipartite case of a state shared by Alice and Bob.) Indeed, Gurvits showed in [1] that it is NP-hard to decide whether a bipartite state ρ_{AB}, given as an $N \times N$ density matrix, is separable across the A/B cut or ϵ-far in trace distance from the set $\mathcal{S}$ of separable states (for ϵ scaling as $1/e^{O(N)}$), promised one of these is the case. Gharibian [2] later strengthened the separation ϵ to an inverse polynomial in N, thus making it robust to Cook polynomial-time reductions.

The overall structure of the general quantum separability problem, as a promise problem with a $1/\text{poly}(N)$ gap between the distance bounds, has been preserved in all subsequent variants addressed in the literature. However, in one major respect the formulation of Gurvits and Gharibian is unsatisfactory from both an operational and a complexity-theoretic standpoint: ρ_{AB} is given explicitly in terms of its matrix entries. It is well-known that fully reconstructing an unknown density matrix is itself an inefficient, poorly scaling, and computationally intensive process; moreover, N itself scales as 2^n where the number of qubits n is the conventional metric for input size in complexity theory. But any preparation of an arbitrary n-qubit input state ρ involves an underlying unitary process, suggesting it is more "natural" in some sense to characterize ρ as the output of a generalized circuit. This alternative scheme is also advantageous from a complexity-theoretic perspective because the circuit generating ρ can be directly embedded in that of the separability-testing algorithm. We can additionally differentiate several classes of input states by further specifying the structure of their generating circuits—which, combined with refinements on the desired output form or the distance measure used, could result in separability problem formulations of widely varying complexity.

In [3] and [4], Gutoski, Hayden, Milner, and Wilde fully exploit the potential of such "custom" variants of the separability problem, all of which take the following general form:

Separability problem schema: *We are given as input a description of a quantum circuit (isometry/channel) and promised that its output ρ is one of the following:*

> ***Yes:*** *α-close in trace distance to a product/separable state (for some pure/mixed input ρ_0 to the isometry/channel),*
>
> $$\min_{\rho} \min_{\sigma \in \mathcal{S}(A|B)} ||\rho - \sigma||_1 \leq \alpha, \tag{2}$$
>
> ***No:*** *β-far in some distance norm $||X||_N \leq ||X||_1$ from all product/separable states (for all pure/mixed inputs ρ_0 to the isometry/channel),*
>
> $$\min_{\rho} \min_{\sigma \in \mathcal{S}(A|B)} ||\rho - \sigma||_N \geq \beta, \tag{3}$$

for some α and β such that $\beta - \alpha \geq 1/poly(n)$, where $\mathcal{S}(A|B)$ is the set of pure/mixed states that are of product/separable form across the $A|B$ cut. Decide which is the case.

We briefly elaborate now on each of the choices needed to specify a particular variant, indicated above by slashes and parenthetical additions. We have already addressed the distinction between product and separable states, which applies only to the mixed case. As for the form of the process that generates ρ, we allow both for *circuits* in which the input to the unitary U is fixed as the all-zeros state $|0\rangle$, and for *isometries* or *channels* that take an auxiliary input ρ_0 ($|\psi_0\rangle$ in the pure case) in tensor product with $|0\rangle$. The output is also allowed to be just Alice and Bob's shared state ρ, or some ρ_{RAB} distributed over three registers R, A, and B, the first of which is traced over to obtain $\rho = \mathrm{Tr}_R(\rho_{RAB})$. "Isometries" refer to the first case and "channels" to the second; obviously it only makes sense to work with mixed states if ρ is generated by tracing over part of the output, but otherwise this formulation admits pure- and mixed-state variants for each configuration.

The last area of specification deserves special attention: our choice of norm for the No case's lower bound β on the distance to the set of separable/product states $\mathcal{S}(A|B)$. The definition of the trace distance $||\rho - \sigma||_1$ involves an optimization over all unitaries acting on the *full* Hilbert space of ρ and σ, or equivalently over all projectors on that space. But we are considering ρ and σ as bipartite states shared between two (possibly well-separated) parties Alice and Bob who can act only on their respective parts, and implementing a general unitary on their joint Hilbert space may require extensive coordination and arbitrarily many rounds of communication and joint local operations. Therefore it makes operational sense to define another metric, the *one-way LOCC norm*, that requires only one:

$$||\rho - \sigma||_{1-LOCC} \equiv \max_{\Lambda_{B\to M}} ||(I_A \otimes \Lambda_{B\to M})(\rho - \sigma)||_1 \tag{4}$$

where we maximize over quantum-to-classical channels $\Lambda_{B\to M} = \sum_m \mathrm{Tr}(\Lambda_m \rho)|m\rangle\langle m|$ with $\{|m\rangle\}$ an orthonormal basis as $\{\Lambda_m\}$ a general POVM. Essentially this restricts to the set of operations in which Bob locally measures his part of the state and writes the result to the register M, which he can then (classically) communicate to Alice. Owing to this restriction we obviously have $||X||_{1-LOCC} \leq ||X||_1$ for all matrices X, meaning that replacing the trace distance with the one-way LOCC distance in the "No" criterion results in a stronger promised separation between the two cases—and thus a reduction in computational complexity, if there is any change at all.

2 Main Results

We earlier mentioned that one advantage of taking a (generalized) circuit description as input, in this formalism for quantum separability problems, is our ability to embed it—including possible ancillary inputs ρ_0 and output registers R—into the very circuit that decides the separability/product-ness of its output ρ. This approach is particularly suited to the setting of *quantum interactive proofs*, in which the computationally unbounded prover "Merlin" and the polynomially-limited verifier "Arthur" hold their respective private qubit registers P and V, as well as a shared register M whose contents constitute their messages to each other. In particular, for variants in which part of the output is traced over to yield ρ, the ancillary output register R can be treated as a quantum message from Arthur (who holds ρ in his private register) to Merlin. On the other hand, constructions that involve optimization on an auxiliary input ρ_0 lend themselves naturally to interactive-proof protocols in which Merlin (or multiple Merlins) supply the quantum message ρ_0 in order to convince Arthur of α-closeness to a separable/product state for Yes instances.

Of course, we also have the freedom to add as many further rounds of communication and computation as we like; however, in complexity-theoretic terms these collapse into a very limited number of distinct classes. This follows from Kitaev and Watrous's result [5] that three messages suffice for any quantum interactive proof, i.e. QIP=QIP(3) where QIP(m) is the class of promise problems efficiently decidable in m rounds of communication. (Note that QIP(0)=BQP and QIP(1)=QMA in these terms.) Another direction for generalization is to allow multiple provers: when restricting to the case of a single message from Merlin to Arthur, this gives the family of classes QMA(k) where k is the number of Merlins. Harrow and Montanaro have shown [6] that QMA(k)=QMA(2) so this hierarchy collapses dramatically

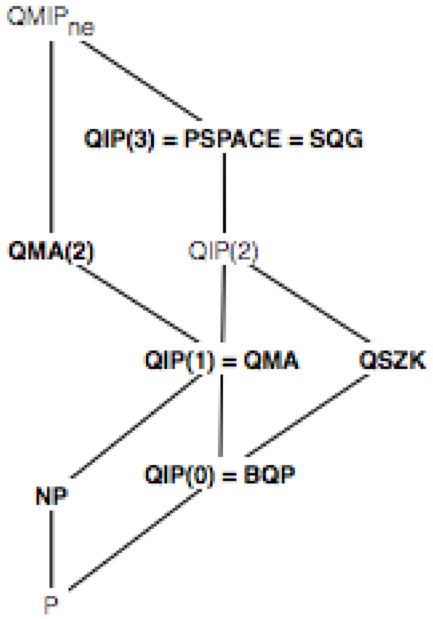

Figure 1: The quantum interactive proof hierarchy, with bold face denoting classes for which some variant of separability testing is a complete problem. Image credit [4].

as well. We can further impose the *statistical zero knowledge* condition that any message from Merlin certifies one yes-instance and (in a rigorous sense) reveals nothing more about the structure of the problem, giving the class QSZK.

Fig. 1 shows the hierarchy of quantum interactive proof classes and displays the extraordinary results of Gutoski, Harrow, Milner, and Wilde: some variant of the quantum separability problem is complete for *each* of the classes in bold. With the exception of NP, these variants are all encompassed by the circuit-based schema from 1, and the only class in the hierarchy for which such a complete problem is as yet unknown is QIP(2)—for which the one-way LOCC version of (mixed-state) separability testing is a strong candidate for completeness. Table 1 summarizes these results in full detail.

3 Proof Techniques

Unsurprisingly, an impressive array of proof techniques must be marshaled in establishing all the results in Table 1, of which we can give only a flavor within the scope of this review. Recall that each completeness statement includes both membership in the class and hardness for it, and that the former generally requires yes- and no-instances to be addressed separately for some fixed relation between the bounds α and β. Nonetheless, many of the respective proofs of membership and of hardness are essentially slight variants of each other, building on the same basic circuit with additional rounds of communication from Merlin(s) to Arthur at the start, or from Arthur to Merlin(s) at the end, appended according to the nature of the specific problem.

Problem	Input	Output	Product?	Pure?	"No" norm	Complexity
PURE PRODUCT STATE	$\|0\rangle$	AB	—	pure	$\|\|X\|\|_1$	BQP-complete
PURE PRODUCT STATE (1-LOCC)	$\|0\rangle$	AB	—	pure	$\|\|X\|\|_{1-LOCC}$	BQP-complete
SEPARABLE ISOMETRY OUTPUT (1-LOCC)	$\rho_0 \otimes \|0\rangle\langle 0\|$	AB	separable	mixed	$\|\|X\|\|_{1-LOCC}$	QMA-complete
PURE PRODUCT ISOMETRY OUTPUT	$\|\psi_0\rangle \otimes \|0\rangle$	AB	—	pure	$\|\|X\|\|_1$	QMA(2)-complete
PRODUCT ISOMETRY OUTPUT	$\rho_0 \otimes \|0\rangle\langle 0\|$	AB	product	mixed	$\|\|X\|\|_1$	QMA(2)-complete
SEPARABLE ISOMETRY OUTPUT	$\rho_0 \otimes \|0\rangle\langle 0\|$	AB	separable	mixed	$\|\|X\|\|_1$	QMA(2)-complete
PRODUCT STATE	$\|0\rangle$	RAB	product	mixed	$\|\|X\|\|_1$	QSZK-complete
SEPARABLE STATE (1-LOCC)	$\|0\rangle$	RAB	separable	mixed	$\|\|X\|\|_{1-LOCC}$	in QIP(2), QSZK-hard
SEPARABLE CHANNEL OUTPUT (1-LOCC)	$\rho_0 \otimes \|0\rangle\langle 0\|$	RAB	separable	mixed	$\|\|X\|\|_{1-LOCC}$	QIP-complete

Table 1: The nine variants on separability testing addressed in [3] and [4], their circuit, output, and norm specifications, and their complexity properties.

An archetypal problem for both membership and hardness proof techniques is SEPARABLE STATE, as discussed in detail in [3]. The interactive proof system that establishes SEPARABLE STATE $\in$ QIP(2) proceeds as follows: first, Arthur runs the circuit that produces ρ and sends the ancillary output $\mathrm{Tr}_{AB}(\rho_{RAB})$ in the R register to Merlin. Merlin then appends ancilla qubits in the $|0\rangle$ state so as to purify this to some output on registers $R', B_2, ..., B_k$ where all the B_i have the same size as B; he sends the systems $B_1, ..., B_k$ to Arthur. Finally, Arthur performs the *permutation test* on Merlin's $k-1$ systems and his own system B:

1. Prepare a $k!$-dimensional register W in a uniform superposition over all basis states by applying a quantum Fourier transform to $|0\rangle$.
2. Apply a controlled unitary that permutes $B, B_2, ..., B_k$ according to the index in W.
3. Invert the QFT on W and measure in the computational basis; accept iff the result is $|0\rangle$.

This is a generalization of the well-known swap test. Establishing membership for both yes- and no-instances relies on a number of frequently recurring results concerning trace distance and a property known as k-extendibility which we state without proof:

Fuchs-van-de-Graaf inequalities: $1 - \sqrt{F(\rho,\sigma)} \leq \frac{1}{2}||\rho - \sigma||_1 \leq \sqrt{1 - F(\rho,\sigma)}$ *where the fidelity* $F(\rho,\sigma)$ *is the maximum overlap* $|\langle\phi_\rho|\phi_\sigma\rangle|^2$ *between purifications of* ρ, σ.

k**-extendibility:** *A bipartite state* ρ_{AB} *is called* k*-extendible if there exists a state* $\sigma_{AB_1\cdots B_k}$ *invariant under permutation of the* B_i *such that* $\rho_{AB} = \mathrm{Tr}_{B_2\cdots B_k}(\sigma_{AB_1\cdots B_k})$. *Clearly, separable states are* k*-extendible for all* k*; conversely, a state that is* ϵ*-far from separable in one-way LOCC distance is also* δ*-far from* k*-extendible in trace distance for some* k *and* $\delta < 0$.

Variations on this proof suffice to show SEPARABLE ISOMETRY OUTPUT (1-LOCC) $\in$ QMA as well as SEPARABLE CHANNEL OUTPUT (1-LOCC) $\in$ QIP(3)=QIP. As for the QSZK-hardness of SEPARABLE STATE, the proof follows by reduction to the known QSZK-hard problem QUANTUM STATE DISTINGUISHABILITY (QSD), which takes the description of a mixed-state circuit generating two outputs ρ_0 and ρ_1 and promises they are either close or far in trace distance. The reduction is accomplished using a circuit that applies U to $|0\rangle$, controlled by half of an EPR pair $|\Phi^+\rangle$. Close variants of this argument prove hardness of PRODUCT STATE for QSZK and SEPARABLE CHANNEL OUTPUT (1-LOCC) for QIP, the latter by reduction to a problem similar to QSD but complete for QIP instead.

The membership and hardness proofs for PURE PRODUCT STATE in BQP are the templates for almost all of the remaining results. Membership follows from a simple argument using the *product test*, which evaluates whether a pure state $|\psi\rangle$ is of m-partite product form by taking two copies in registers $A_1, ...A_m$ and $B_1, ..., B_m$ and performing m swap tests on the respective pairs (A_i, B_i) [6]. Hardness, on the other hand, is established by reducing an arbitrary promise problem in BQP to the one-way LOCC version of PURE PRODUCT STATE using a controlled swap between half of a $2n$-qubit maximally entangled state (e.g. n EPR pairs) and an all-zeros register; the reduction makes use of a bound on the one-way LOCC distance between such a state and a $2n$-qubit separable state proved in [4]. Close variations on these arguments prove that each of the three ISOMETRY OUTPUT problems (in which trace distance is used for the no-instance bound) is both contained in QMA(2) and hard for that class.

4 Conclusions

We introduced and justified, on both operational and complexity-theoretic grounds, a general schema for promise problems that test the closeness of states generated by (generalized) circuits to the set of bipartite separable states. These problems vary in a number of "customizable" dimensions: whether they require auxiliary inputs and outputs, admit pure or mixed outputs, test for separability or product form, and use trace distance or one-way LOCC distance to define no-instances. Following recent results of Gutoski, Hayden, Milner, and Wilde, this variation of problem specifications is reflected by variation in complexity: various problems of this type are complete for each of the interactive proof classes BQP, QMA, QMA(2), QSZK, and QIP. However, while in each dimension of variation it is clear which alternative ought to be more computationally difficult—namely the existence of auxiliary inputs and outputs, mixed states, separability, and trace distance respectively—Table 1 shows that there seems to be no pattern to how the choices interact with each other. For instance, the product/separable distinction induces a separation in complexity for mixed-state circuits acting on $|0\rangle$, but not for the mixed-state isometry problems. Further resolving these relations is a promising area for future research.

References

[1] L. Gurvits, (2003), quant-ph/0303055 .

[2] S. Gharibian, Quantum Inf. and Comp. **10**, 343 (2010).

[3] P. Hayden, K. Milner, and M. M. Wilde, Quantum Inf. and Comp. **14**, 384 (2014).

[4] G. Gutoski, P. Hayden, K. Milner, and M. M. Wilde, quant-ph/1308.5788 .

[5] A. Kitaev and J. Watrous, in *Proceedings of the 32nd ACM Symposium on Theory of Computing* (2000) pp. 608–617.

[6] A. Harrow and A. Montanaro, in *Proceedings of the 51st Annual IEEE Symposium on the Foundations of Computer Science* (2010) pp. 633–642.