

The Pure-State Consistency of Local Density Matrices Problem: In PSPACE and Complete for a Class Between QMA and QMA(2)

Jonas Kamminga  

Department of Computer Science and Institute for Photonic Quantum Systems (PhoQS),
Paderborn University, Germany

Dorian Rudolph  

Department of Computer Science and Institute for Photonic Quantum Systems (PhoQS),
Paderborn University, Germany

Abstract

In this work we investigate the computational complexity of the pure consistency of local density matrices (PureCLDM) and pure N -representability (Pure- N -Representability; analog of PureCLDM for bosonic or fermionic systems) problems. In these problems the input is a set of reduced density matrices and the task is to determine whether there exists a global *pure* state consistent with these reduced density matrices. While mixed CLDM, i.e. where the global state can be mixed, was proven to be QMA-complete by Broadbent and Grilo [JoC 2022], almost nothing was known about the complexity of the pure version. Before our work the best upper and lower bounds were QMA(2) and QMA. Our contribution to the understanding of these problems is twofold.

Firstly, we define a pure state analogue of the complexity class QMA⁺ of Aharonov and Regev [FOCS 2003], which we call PureSuperQMA. We prove that both pure- N -Representability and PureCLDM are complete for this new class. Along the way we supplement Broadbent and Grilo by proving hardness for 2-qubit reduced density matrices and showing that mixed N -Representability is QMA-complete.

Secondly, we improve the upper bound on PureCLDM. Using methods from algebraic geometry, we prove that PureSuperQMA \subseteq PSPACE. Our methods, and the PSPACE upper bound, are also valid for PureCLDM with exponential or even perfect precision, hence precisePureCLDM is not preciseQMA(2) = NEXP-complete, unless PSPACE = NEXP. We view this as evidence for a negative answer to the longstanding open question whether PureCLDM is QMA(2)-complete.

The techniques we develop for our PSPACE upper bound are quite general. We are able to use them for various applications: from proving PSPACE upper bounds on other quantum problems to giving an efficient parallel (NC) algorithm for (non-convex) quadratically constrained quadratic programs with few constraints.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum complexity theory

Keywords and phrases Quantum Complexity Theory, PSPACE, QMA(2), Consistency of Local Density Matrices, Polynomial Optimization

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.83

Related Version *Full Version*: <https://arxiv.org/abs/2411.03096> [18]

Funding *Jonas Kamminga*: supported by the DFG under grant number 450041824.

Dorian Rudolph: supported by the DFG under grant number 432788384 and 563388236.

Acknowledgements The authors would like to thank Dima Grigoryev for suggesting that the GP algorithm could be parallelized, pointing to helpful references, and giving us advice on how this could be proven.



© Jonas Kamminga and Dorian Rudolph;

licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 83; pp. 83:1–83:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

“Are these local density matrices consistent with some global state?” This quantum variant of the NP-complete *marginal problem* is known as the *consistency of local density matrices problem* (CLDM) or *quantum marginal problem* and as the *N-representability problem* (*N-Representability*) when dealing with indistinguishable particles. It is of fundamental interest to quantum physics. In fact, it was already recognized as an important question in the sixties [10]. At that time, the hope was that the ground state energy of quantum systems could be found using reduced density matrices. This hope was supported by the fact that Hamiltonians showing up in nature are all *local*. One requirement would be that it is possible to check that alleged reduced density matrices are indeed consistent with a valid global quantum state, hence the interest in the CLDM problem.

Over the years it became apparent that this hope would not materialize, especially when Kitaev proved that computing ground state energies of local Hamiltonians is QMA-hard [20]. Also the CLDM problem itself was proven to be hard. First by Liu, who proved that the (mixed state) CLDM problem is contained in QMA and QMA-complete under Turing reductions [22] and later, together with Christandl and Verstraete, proved that the *N-representability problem* is also QMA-complete under Turing reductions [23]. This was improved by Broadbent and Grilo who showed that the mixed CLDM is also QMA-hard under Karp reductions, establishing it as QMA-complete [9].

In this work we focus our attention on a variant of CLDM that is of interest to physics and complexity theory. This variant, the *pure consistency of local density matrices problem*, or PureCLDM asks not whether there exists any state consistent with the reduced density matrices, but demands that this consistent state is pure. Physically, this restriction to pure states is well motivated as isolated systems will always be in a pure state. If one has some reduced density matrices of some molecule it might be more interesting to know whether these are consistent with some global state of the molecule in isolation, i.e. a pure state, rather than some larger state where the molecule is entangled with its environment.

From a complexity point of view PureCLDM is interesting because of its connections with the class QMA(2). While the complexity of mixed CLDM is well understood to be QMA, the best known upper bound to PureCLDM is the class QMA(2) and it has been a longstanding open question whether PureCLDM is complete for QMA(2) [23]. If it would be complete that would be interesting for several reasons. Firstly, it would be a highly physically motivated complete problem for QMA(2), a class that has attracted a lot of attention but generally lacks physical complete problems. Secondly, there has been a large interest in the relation between purity testing and QMA(2). In many ways, the power of QMA(2) seems to come from the ability to do SWAP tests [16]. On the other hand, a correspondence between protocols over separable states and protocols with a purity constraint has been used in several works (e.g. [16, 32]). Recently, it has been suggested that purity testing could already be enough to give QMA(2) its power [4]. Establishing PureCLDM as a complete problem for QMA(2) would suggest that indeed SWAP tests are not necessary to capture the power of QMA(2) and indeed a purity constraint is already enough.

1.1 Our results

In this work we investigate the complexity of PureCLDM and Pure-*N-Representability* and give evidence towards a negative answer to the longstanding open question whether they are QMA(2)-complete.

► **Definition 1.1** (*k*-PureCLDM, informal). We are given pairs $(\rho_1, C_1), \dots, (\rho_m, C_m)$, where the ρ_i are reduced density matrices and $C_i \subseteq [n]$ with $|C_i| \leq k$ for all i . Additionally, we are given parameters α and β with $\beta - \alpha \geq 1/\text{poly}(n)$. Decide whether:

- YES: there exists a consistent pure state, that is, a state $|\psi\rangle$ such that $\|\text{Tr}_{\overline{C_i}}(|\psi\rangle\langle\psi|) - \rho_i\| \leq \alpha$ for all $i \in [m]$.
- NO: all pure states are “far from” consistent, that is, for all $|\psi\rangle$, there is an $i \in [m]$ with $\|\text{Tr}_{\overline{C_i}}(|\psi\rangle\langle\psi|) - \rho_i\| \geq \beta$.

1.1.1 PureCLDM is complete for PureSuperQMA

Our first main result on the complexity of PureCLDM is that it is complete for the class of (promise) problems decided by a “super-verifier” (c.f. [2]) that is given a pure state proof. Following [2] we call this class PureSuperQMA.¹

► **Definition 1.2** (PureSuperQMA, informal). A promise problem $A = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ is in PureSuperQMA(m, ε, δ) if there exist m constraints $\mathcal{V} = \{(V_{x,i}, r_{x,i}, s_{x,i})\}_{i \in [m]}$ such that:

- $\forall x \in A_{\text{yes}} \exists |\psi\rangle: \Pr_i(|p(V_{x,i}, \psi) - r_{x,i}| \leq s_{x,i}) = 1$, that is, the $V_{x,i}$ accept $|\psi\rangle$ with acceptance probability at most $s_{x,i}$ away from $r_{x,i}$.
- $\forall x \in A_{\text{no}} \forall |\psi\rangle: \Pr_i(|p(V_{x,i}, \psi) - r_{x,i}| \leq s_{x,i} + \varepsilon) \leq 1 - \delta$, that is, at least a δ fraction of the $V_{x,i}$ accept $|\psi\rangle$ with probability more than $s_{x,i} + \varepsilon$ away from $r_{x,i}$.

Here $p(V, \psi)$ denotes the acceptance probability of the circuit V when ran on input $|\psi\rangle$.

In the following we refer to m as the number of checks or constraints and to ε and δ as precision parameters. We denote the union of PureSuperQMA(m, ε, δ) where m is polynomial and ε and δ are inverse polynomial as $\text{PureSuperQMA}(\text{poly}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}}) =: \text{PureSuperQMA}$.

We show that the definition of PureSuperQMA can be significantly simplified: we can set all $r_{x,i} = 1/2$ and all $s_{x,i} = 0$ without changing the complexity. Our proof of PureSuperQMA-completeness holds already for k -PureCLDM₁, which is k -PureCLDM with “exact consistency” as well as for Bosonic/Fermionic Pure- N -Representability.

► **Theorem 1.3.** k -PureCLDM₁ is PureSuperQMA-complete for all $k \geq 2$.

► **Corollary 1.4.** Fermionic/bosonic Pure- N -Representability₁ is PureSuperQMA-complete.²

Note that PureSuperQMA as the complete complexity class is quite natural, since Liu’s proof for $\text{CLDM} \in \text{QMA}$ also goes via $\text{SuperQMA} = \text{QMA}$. Since a pure state analogue to this latter statement is not known, we have to remain the super-verifier regime. Along the way to Theorem 1.3, we show that the k -CLDM problem is already QMA-hard for $k \geq 2$. This improves upon the results by Broadbent and Grilo [9] who show hardness for $k \geq 5$. We also resolve one of their open questions by showing that the (mixed) fermionic and bosonic N -representability problems are QMA-hard, already for 2-particle reduced density matrices.

1.1.2 PureSuperQMA as a “light” version of QMA(2)

Next, we prove results suggesting PureSuperQMA can be viewed as a “light” version of QMA(2). It lies between QMA and QMA(2) but contains NP already with log-size proofs and becomes NEXP when the number of checks and the precision are exponential. This mirrors

¹ [2] use the name QMA+, but recently QMA+ has been used to refer to QMA with proofs with nonnegative amplitudes [17, 5]. As QMA+ is sometimes referred to as QMA with a super-verifier, we use SuperQMA.

² Technical details of Pure- N -Representability are in the full version [18].

QMA(2) that also contains NP with log-size proofs [6] and becomes equal to NEXP when made exponentially precise [26]. On the other hand, QMA with log-size proofs is equal to BQP [24], which is believed not to contain NP, and PreciseQMA = PSPACE [11] is believed to be different from NEXP. Figure 1 depicts the relationships between the classes discussed here.

► **Proposition 1.5.** $\text{QMA} \subseteq \text{PureSuperQMA} \subseteq \text{PureSuperQMA}(\text{exp}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}}) \subseteq \text{QMA}(2)$.

► **Proposition 1.6.** NP can be decided in PureSuperQMA already with log-size proofs.

► **Proposition 1.7.** $\text{PureSuperQMA}(\text{exp}, \frac{1}{\text{exp}}, \frac{1}{\text{exp}}) = \text{NEXP}$

1.1.3 PureSuperQMA is upper bounded by PSPACE

Proving a non-trivial upper bound on the complexity of QMA(2) is a major open question in the field. As such, one might wonder whether our “light” version of QMA(2) does admit a non-trivial upper bound. Our second main result is that it does indeed, namely PSPACE. This sharpens the upper bound on the complexity of PureCLDM from $\text{QMA}(2) \subseteq \text{NEXP}$ to PSPACE.

► **Theorem 1.8.** $\text{PureSuperQMA} \subseteq \text{PureSuperQMA}(\text{poly}, 0, \frac{1}{\text{poly}}) \subseteq \text{PSPACE}$.

Our proof relies on methods from algebraic geometry and also works for PureSuperQMA with exponential or even perfect precision *as long as the number of constraints is polynomial*.

What does this mean for QMA(2)? Of course showing that PureCLDM is QMA(2)-hard would imply $\text{QMA}(2) \subseteq \text{PSPACE}$ but there is a catch. To prove Theorem 1.8 we use methods from algebraic geometry that also work for PureSuperQMA with exponential precision and even with perfect precision. This allows us to obtain the following corollary.

► **Corollary 1.9.** *If Precise-PureCLDM is PreciseQMA(2)-hard, then PSPACE = NEXP.*

This corollary can be interpreted either as evidence that PureCLDM is not QMA(2)-hard, or as a barrier to a hardness proof: any such proof must fail in the precise case, assuming that $\text{PSPACE} \neq \text{NEXP}$.

Furthermore, by combining Proposition 1.5 and Theorem 1.3 we conclude that PureCLDM can only be QMA(2)-complete if a polynomial number of constraints give the same power as exponentially many constraints.


► **Corollary 1.10.** *If $\text{PureSuperQMA}(\text{poly}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}}) \subsetneq \text{PureSuperQMA}(\text{exp}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}})$, then PureCLDM is not QMA(2)-hard.*

Lastly, the fact that QMA(2) contains NP already with log-size proofs is sometimes taken as evidence that QMA(2) could be equal to NEXP. We argue that our results imply that this is not evidence at all. Indeed, PureSuperQMA also contains NP with log-size proofs but is contained in PSPACE.

1.1.4 Solving systems of quadratic equations efficiently in parallel

To prove Theorem 1.8 we rely on techniques by Grigoriev and Pasechnik [14] for solving certain polynomials. We call such polynomials *GP systems*.

► **Definition 1.11** (GP system (informal)). *A GP system is a polynomial of the form $p(Q(X))$, where $p: \mathbb{R}^k \rightarrow \mathbb{R}$ is a degree d polynomial in “few” variables, and $Q: \mathbb{R}^N \rightarrow \mathbb{R}^k$ is quadratic in “many” variables. We assume the coefficients of p and Q are integers of bit size at most L . A solution to the system is a point $X^* \in \mathbb{R}^N$ such that $p(Q(X^*)) = 0$.*



■ **Figure 1** Diagram of the complexity classes discussed in this paper. Theorem 1.8 shows $\text{PureSuperQMA} \subseteq \text{PSPACE}$. BellPureSymQMA is discussed in the full version [18]. Proposition 1.5 shows $\text{PureSuperQMA}(\text{exp}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}}) \subseteq \text{QMA}(2)$. $\text{PSPACE} = \text{PreciseQMA}$ is shown in [11]. The third level of the quantum polynomial hierarchy, $\text{Q}\Sigma_3$, is between $\text{QMA}(2)$ and NEXP [12]. $\text{NEXP} = \text{PreciseQMA}(2)$ is shown in [6, 26]. The relationship between PSPACE and $\text{QMA}(2)$ remains a major open problem.

Grigoriev and Pasechnik give an algorithm that can compute a set of univariate representation of solutions of $p(Q(X))$ that intersects all connected components of the solution set [14, Theorem 1.2]. Their algorithm requires $(dN)^{O(k)}$ arithmetic operations.

We show how to write a PureSuperQMA verifier as a GP system with $k = \text{poly}(n)$ and $N, L = 2^{\text{poly}(n)}$ such that the system has a solution iff the verifier accepts some pure proof. The results by Grigoriev and Pasechnik then immediately imply $\text{PureSuperQMA} \subseteq \text{EXP}$. To prove containment in PSPACE we modify the algorithm to work efficiently in parallel. The techniques we develop turn out to be much more general. They also show that solving quadratic systems with few constraints is in NC , or Nick’s Class, which captures efficient parallel computation.

► **Definition 1.12** (Nick’s Class [3]). *The i -th level of NC , NC^i , consists of all languages decidable in time $O(\log^i(n))$ using $\text{poly}(n)$ parallel processors. NC is the union of NC^i over all $i \in \mathbb{N}$.³*

The class $\text{NC}(\text{poly})$ of all decision problems decidable in $\text{poly}(n)$ time on $\text{exp}(n)$ processors is equal to PSPACE [7]. Our modification of Grigoriev and Pasechnik’s algorithm yields:

► **Theorem 1.13.** *Let p, Q be a GP system. There is a parallel algorithm for deciding whether $p(Q(X))$ has a root. The algorithm uses $L^2(k^2Nd)^{O(k)}$ parallel processors and needs $\text{poly}(\log N, \log d, k, \log L)$ time. In particular, if $\log N, \log d, \log L, k \leq \text{poly}(n)$ the computation can be done in $\text{NC}(\text{poly}) = \text{PSPACE}$. If $N, d, L \leq \text{poly}(n)$ and $k = O(1)$ then the algorithm is in NC .*

We further build upon [14] and consider the following optimization problem which includes the task of quantifying how inconsistent local density matrices are.

³ Equivalently, NC^i can be defined as the class of languages decidable by an $O(\log^i n)$ -depth circuit of size $\text{poly}(n)$.

► **Definition 1.14** (OptGP). *An OptGP instance consists of a degree $\leq d$ polynomial $r: \mathbb{R}^k \rightarrow \mathbb{R}$ and a GP system $p(Q(X))$ with bounded solution set $Z = \{X \in \mathbb{R}^N: p(Q(X)) = 0\}$. The task is to compute or approximate $\theta = \min_{X \in Z} r(Q(X))$.*

We give a parallel algorithm for approximating the optimum of a OptGP system as well as approximating a witness to this optimum.

► **Theorem 1.15.** *Let r, p, Q be an OptGP instance. There is a parallel algorithm that computes a δ -approximation to $\theta = \min_{X \in Z} r(Q(X))$ and to some X^* with $r(Q(X^*)) = \theta$ and $p(Q(X^*)) = 0$. The algorithm uses $\text{poly}(L, |\log \delta|, (kNd)^{O(k^2)})$ parallel processors and needs $\text{poly}(\log N, \log d, k, \log L, \log |\log \delta|)$ time.*

Grigoriev and Pasechnik [14, Theorem 1.5] also state a theorem with a sequential algorithm for the optimization problem that also applies to unbounded Z , but their proof is not yet available to the best of our knowledge.

Finally, we showcase the applicability of the developed algorithms. The formal treatment of these applications are deferred to the full version [18] but we already state the results.

First, we improve upon a result by Shi and Wu. In [30] they give a PSPACE algorithm for optimizing the energy of “decomposable” Hamiltonians over separable states. Using our framework we are able to reprove this fact, and even get a better runtime dependence on the error.

Second, we show that deciding if there exists a *unique*⁴ pure state that is consistent with given local density matrices is also in PSPACE. In other words, we can decide in PSPACE whether the local density matrices *fully* describe the physics of the system.

Third, we show how to decide a variant of PureCLDM, where the input only specifies the spectrum of the local density matrices. This version is sometimes referred to as the quantum marginal problem, although others use that name for our PureCLDM.

Fourth, we consider a pure-state variant of BellSymQMA(poly). This class (see [1, 8]) captures QMA protocols where Merlin promises to send poly many copies of a proof in tensor product, but Arthur has measure each copy independently and then decide to accept or reject based on the classical measurement outcomes. It turns out that BellSymQMA(poly) is actually equal to QMA [8]. The variant we consider, BellPureSymQMA(poly) adds the additional promise that the proof is pure and does not obviously equal QMA. It is easily seen that PureSuperQMA \subseteq BellPureSymQMA(poly) as Arthur can approximate the acceptance probability of the checks by doing many independent measurements. We slightly strengthen Theorem 1.8 by proving that also BellPureSymQMA(poly) \subseteq PSPACE.

Lastly, we consider an entirely classical problem: quadratic programming with few constraints. A quadratically constrained quadratic program (QCQP) is an optimization problem of the form

$$\min_{x \in \mathbb{R}^n} x^\top A_0 x + a_0^\top x \tag{1a}$$

$$\text{subject to } x^\top A_i x + a_i^\top x \leq b_i \quad \forall i \in [m]. \tag{1b}$$

As a final application of the GP framework, we show that QCQPs can be solved in NC *if the number of constraints m is constant*. Note that we do not require any assumptions on the A_i such as positivity, symmetry or bounds on the condition number.

⁴ We say a state $|\phi\rangle$ is the unique state consistent with some local density matrices if any state that is orthogonal to $|\phi\rangle$ is far from consistent.

1.2 Proof techniques

Sketches of our proof can be found in Sections 4–6. The formal proofs can be found in the full version [18].

1.3 Related work

The computational complexity of (mixed) CLDM and N -Representability has previously been studied by Liu, Broadbent and Grilo, as mentioned before. Liu [22] proves that (mixed) CLDM is contained in QMA and hard under Turing reductions. Similar results for N -Representability are proven in [23]. This was improved by Broadbent and Grilo who proved (among other results regarding zero-knowledge proof systems) that (mixed) CLDM is also QMA-hard under Karp reductions, thereby fully resolving its complexity [9]. Both Liu, and Broadbent and Grilo do not intensively study PureCLDM, although [23] does show containment of fermionic Pure- N -Representability in QMA(2), leaving hardness as an open question. A similar containment for bosonic Pure- N -Representability was shown in [31].

That does not mean that PureCLDM and Pure- N -Representability have not been studied before. There is a large body of work focussing on finding necessary and/or sufficient conditions for reduced density matrices to be consistent with a global state. Among these works is [21], which focuses on the case where the reduced density matrices are non-overlapping. The paper establishes conditions that are necessary and sufficient for the existence of a consistent pure state in this case. Mazziotti [25] derives necessary conditions for a two-fermion density matrix to have a consist global N -fermion pure state.

[32] rewrite PureCLDM as an optimization problem over separable state. They then apply the method of symmetric extensions to this notoriously hard problem to describe PureCLDM as a hierarchy of SDP's. That is, they describe SDP's depending on a parameter N such that any “No” instance will be discovered by the SDP for sufficiently large N . They do not, however, prove any upper bounds on the size of N required.

In [4], the authors consider QMA with an *internally separable* proof. They prove that when this proof is mixed, the class is contained in EXP, whereas it is equal to NEXP if the proof is pure. This provides the, to our knowledge first, instance where pure proofs are provably stronger than mixed proofs, modulo standard complexity theoretic assumptions.

An algorithm for solving polynomial systems more general than those considered in Theorem 1.13 is given in [13]. It shows that a system of k polynomials of degree d in n variables can be solved in time $\text{poly}(n^{d^{3k}})$. One downside to this algorithm is that it finds a solution over the complex numbers instead of the reals. This makes it hard to constrain the norm of variables, as the complex conjugate is not a polynomial.

1.4 Discussion and open questions

Our work sheds some more light on the complexity of Pure- N -Representability and PureCLDM. However, the story is far from complete as the relation between PureSuperQMA and QMA or PSPACE remains poorly understood. We conjecture

► **Conjecture 1.16.** $\text{QMA} \subsetneq \text{PureSuperQMA} \subsetneq \text{QMA}(2)$.

We give some evidence that PureSuperQMA differs from QMA(2). Indeed, we prove that their precise versions are equal only if $\text{PSPACE} = \text{NEXP}$. However, this does not necessarily carry over from the precise setting to the “standard” setting. It would therefore be nice to see more evidence that $\text{PureSuperQMA} \subsetneq \text{QMA}(2)$, such as an oracle separation. Of course,

separating PureSuperQMA from QMA(2) relative to an oracle is at least as hard as separating QMA from QMA(2) in this way, something that has been eluding researchers to this date. Perhaps, however, the new perspective offered by PureSuperQMA can lead to new insights.

Recently, it has been suggested that purity testing is at the heart of QMA(2)'s power [4]. While we provide evidence that PureCLDM is not QMA(2)-hard, that does not mean the end for this suggestion. One way to formalize the idea that QMA(2)'s power derives from purity would be to prove that $\text{QMA}(2) = \text{PureSuperQMA}(\text{exp}, \frac{1}{\text{poly}}, \frac{1}{\text{poly}})$. Note that our results do not provide evidence against this equality, as the PSPACE upper bound crucially relies on there being only polynomially many constraints.

The relationship between QMA and PureSuperQMA with a constant number of checks could also be interesting. It can be shown that with a single check $\text{PureSuperQMA}(1, \frac{1}{\text{poly}}, 0) = \text{QMA}$ as there is no single check passed by a mixed state but not by a pure state⁵. We have been unable to find a set of 2 constraints with a consistent mixed state but no consistent pure state but leave $\text{PureSuperQMA}(2, \frac{1}{\text{poly}}, 0) \stackrel{?}{=} \text{QMA}$ as an open question.

Lastly, it would be nice to see if the GP system framework used for our PSPACE upper bound can find other uses. An approach one could take here is to try to use it for a PSPACE or EXP upper bound on QMA(2). There are two main obstacles here. Firstly, any such approach needs to make essential use of the promise gap in order to work for QMA(2) but not for $\text{QMA}(2)_{\text{exp}}$ (assuming $\text{PSPACE} \neq \text{NEXP}$). Secondly, naively converting a QMA(2) instance into polynomials yields degree 3, for which the techniques from [14] no longer work.

2 Preliminaries

The main object of study in this paper is the k -PureCLDM problem on qubits.

► **Definition 2.1** (k -PureCLDM). *We are given a set of reduced density matrices ρ_1, \dots, ρ_m with $\text{poly}(n)$ bits of precision, where each ρ_i acts on qubits $C_i \subseteq [n]$ with $|C_i| \leq k$, as well as thresholds α, β with $\beta - \alpha \geq 1/\text{poly}(n)$. Decide:*

- *YES: There exists a state $|\psi\rangle \in \mathbb{C}^{2^n}$ such that $\|\text{Tr}_{\overline{C_i}}(|\psi\rangle\langle\psi|) - \rho_i\| \leq \alpha$ for all $i \in [m]$.*
- *NO: For all states $|\psi\rangle \in \mathbb{C}^{2^n}$, there exists an $i \in [m]$ such that $\|\text{Tr}_{\overline{C_i}}(|\psi\rangle\langle\psi|) - \rho_i\| \geq \beta$. For $\alpha = 0$, we denote the problem by k -PureCLDM₁. Define k -CLDM and k -CLDM₁ analogously, but allowing a mixed state in place of $|\psi\rangle$.*

Note that PureCLDM and (mixed) CLDM are indeed different. The following example shows that a consistent mixed state may exist even if no consistent pure state exists.

► **Example 2.2.** Let $\rho = \frac{1}{n} \sum_{i=1}^n |\psi_i\rangle\langle\psi_i|$, where $|\psi_i\rangle = |0^{i-1}10^{n-i}\rangle \in \mathbb{C}^{2^n}$. Then all 2-local reduced density matrices of ρ are $\rho_{ij} = \frac{n-2}{n}|00\rangle\langle 00| + \frac{1}{n}|01\rangle\langle 01| + \frac{1}{n}|10\rangle\langle 10|$. Assume there exists a pure state $\sigma = |\phi\rangle\langle\phi|$ such that $\sigma_{ij} = \rho_{ij}$ for all $i, j \in [n]$. Then $|\phi\rangle \in \text{Span}\{|0^n\rangle, |\psi_1\rangle, \dots, |\psi_n\rangle\}$ since ρ has no overlap with any string of Hamming weight ≥ 2 . Hence, $|\phi\rangle$ must be of the form $|\phi\rangle = \sum_{i=1}^n a_i |\psi_i\rangle$ with $|a_i| = \sqrt{1/n}$. Then $\sigma_{12} = \frac{2}{n} |\eta\rangle\langle\eta| + \frac{n-2}{n} |00\rangle\langle 00|$, where $|\eta\rangle = \sqrt{n/2}(a_1|10\rangle + a_2|01\rangle)$. However, then $\sigma_{12} \neq \rho_{12}$ since their rank differs, which contradicts the choice of $|\phi\rangle$.

⁵ If there exists a mixed state passing the check then there exist pure states $|\psi\rangle, |\phi\rangle$ such that $|\phi\rangle$ is accepted with probability $\leq 1/2$ and $|\psi\rangle$ is accepted with probability $\geq 1/2$. By interpolating between these states and the Intermediate Value Theorem it follows that there exists some pure state accepted with probability exactly $1/2$.

3 PureSuperQMA

In this section we define the complexity class PureSuperQMA. Following the definition we prove that PureSuperQMA contains PureCLDM and some other properties of the class, suggesting it can be thought of as a “light” version of QMA(2). We end the section by establishing a canonical form for PureSuperQMA, cleaning up the definition. This canonical form will be important in the next section, where we use it for the hardness proof.

► **Definition 3.1** (PureSuperQMA, c.f. [2]). *A promise problem A is in $\text{PureSuperQMA}(m, \varepsilon, \delta)$ if there exists a uniformly generated super-verifier⁶ $\mathcal{V} = \{(V_{x,i}, r_{x,i}, s_{x,i})\}_{i \in [m]}$ on $n_1(n) \in n^{O(1)}$ proof qubits and $n_2(n) \in n^{O(1)}$ ancilla qubits ($n = |x|$) such that*

- $\forall x \in A_{\text{yes}} \exists |\psi\rangle \in \mathcal{P}: \Pr_i(|p(V_{x,i}, \psi) - r_{x,i}| \leq s_{x,i}) = 1,$
- $\forall x \in A_{\text{no}} \forall |\psi\rangle \in \mathcal{P}: \Pr_i(|p(V_{x,i}, \psi) - r_{x,i}| \leq s_{x,i} + \varepsilon) \leq 1 - \delta,$

where \mathcal{P} is the set of unit vectors on $n_1(n)$ qubits, $i \in [m]$ is drawn uniformly at random, and $p(V, \psi) := \text{Tr}(\Pi_{\text{acc}} V |\psi\rangle \langle \psi| 0^{n_2} \langle \psi| 0^{n_2} |V^\dagger)$ denotes the acceptance probability of V on input $|\psi\rangle$.⁷ We call each triple $(V_{x,i}, r_{x,i}, s_{x,i})$ a constraint or a check. Let

$$\text{PureSuperQMA} = \bigcup_{m \in n^{O(1)}, \varepsilon, \delta \in n^{-O(1)}} \text{PureSuperQMA}(m, \varepsilon, \delta). \quad (2)$$

► **Lemma 3.2.** k -PureCLDM \in PureSuperQMA for $k \leq O(\log n)$.

The proof is analogous to the containment of Pure- N -Representability in QMA(2) [23] and can be found in the full version [18].

► **Proposition 1.5.** $\text{QMA} \subseteq \text{PureSuperQMA} \subseteq \text{PureSuperQMA}(\exp, \frac{1}{\text{poly}}, \frac{1}{\text{poly}}) \subseteq \text{QMA}(2)$.

Proof sketch. First note that $\text{QMA} \subseteq \text{PureSuperQMA}$ follows by setting $r = 1$ and $s = \frac{1}{\exp(n)}$ in the definition of PureSuperQMA. For the other inclusion we use the fact that $\text{QMA}(2) = \text{QMA}(\text{poly})$ [16]. With probability 1/2 each, the verifier performs one of the following tests: (i) Run swap tests between random disjoint pairs of the registers to ensure the input state is close to a state of the form $|\psi\rangle^{\otimes k}$. (ii) Pick $i \in [m]$ uniformly at random and run $V_{x,i}$ on all k proofs, recording the outcomes as $y_1, \dots, y_k \in \{0, 1\}$, and let $\mu = \frac{1}{k} \sum_{i=1}^k y_k$. Accept if $|\mu - r_{x,i}| \leq s_{x,i} + \varepsilon/2$. For sufficiently large $k \in n^{O(1)}$, we can use Hoeffding’s inequality to prove completeness and soundness. ◀

Note that we still get containment in QMA(2), even with an exponential number of constraints, as long as in the NO-case the acceptance probability deviates from $r_{x,i}$ for a significant fraction of constraints. However, we do have some reason to believe that k -PureCLDM is only hard for PureSuperQMA with a polynomial number of constraints since the hardness proof works for any precision parameter and k -PureCLDM \in PSPACE (see Theorem 1.13) even for exponentially small precision. In contrast, if we have both an exponential number of constraints and exponential precision, then PureSuperQMA contains NEXP.

► **Proposition 1.6.** NP can be decided in PureSuperQMA already with log-size proofs.

⁶ Our definition of a super verifier here might seem slightly different from Aharonov and Regev’s definition. Note that their definition can be recovered from our by repeating the same constraint as necessary.

⁷ Here and in the following, we use ψ to denote the density operator $|\psi\rangle\langle\psi|$.

83:10 On the Pure-State Consistency of Local Density Matrices Problem

Following [6]’s proof of containment of NP in QMA(2) with log-size proofs, the idea behind the proof is to verify a 3-coloring instance by having the prover send a state of the form $|\psi\rangle = \frac{1}{\sqrt{n}} \sum_{v \in V} |v\rangle_V |C(v)\rangle_C$. To force the prover to send a proof of this form we design constraints to force every vertex to show up in the superposition and to have exactly one color associated to it. For the technical details, see the full version [18].

A similar proof using succinct 3-coloring rather than normal 3-coloring shows the following result.

► **Proposition 1.7.** $\text{PureSuperQMA}(\text{exp}, \frac{1}{\text{exp}}, \frac{1}{\text{exp}}) = \text{NEXP}$

Since $\text{PureSuperQMA}(\text{exp}, \frac{1}{\text{exp}}, \frac{1}{\text{exp}}) \subseteq \text{QMA}(2)_{\text{exp}}$ (i.e. QMA(2) with exponentially small promise gap) by the same argument as Proposition 1.5, we also recover the following known result:

► **Corollary 3.3** ([6, 26]). $\text{QMA}(2)_{\text{exp}} = \text{NEXP}$, where $\text{QMA}(2)_{\text{exp}}$ denotes QMA(2) with exponentially small promise gap.

To prove that 2-PureCLDM is PureSuperQMA-complete, we will use a simplified but equivalent definition:

► **Definition 3.4** (PSQMA). A promise problem A is in $\text{PSQMA}(m, \varepsilon)$ if there exists a uniformly generated super-verifier $\mathcal{V} = \{V_{x,i}\}_{i \in [m]}$ on $n_1(n) \in n^{O(1)}$ proof qubits and $n_2(n) \in n^{O(1)}$ ancilla qubits ($n = |x|$) such that

- $\forall x \in A_{\text{yes}} \exists |\psi\rangle \in \mathcal{P} \forall i \in [m]: \left| \text{Tr}(\Pi_{\text{acc}} V_{x,i} |\psi, 0^{n_2}\rangle \langle \psi, 0^{n_2} | V_{x,i}^\dagger) - \frac{1}{2} \right| = 0,$
- $\forall x \in A_{\text{no}} \forall |\psi\rangle \in \mathcal{P} \exists i \in [m]: \left| \text{Tr}(\Pi_{\text{acc}} V_{x,i} |\psi, 0^{n_2}\rangle \langle \psi, 0^{n_2} | V_{x,i}^\dagger) - \frac{1}{2} \right| \geq \varepsilon,$

where \mathcal{P} is the set of unit vectors on $n_1(n)$ qubits.

Let $\text{PSQMA} = \bigcup_{m \in n^{O(1)}, \varepsilon \in n^{-O(1)}} \text{PSQMA}(m, \varepsilon)$.

In other words, PSQMA is PureSuperQMA where all the r ’s are set to $\frac{1}{2}$ and all the s ’s to 0.

► **Lemma 3.5.** $\text{PSQMA} = \text{PureSuperQMA}$.

4 2-PureCLDM is PureSuperQMA-complete

► **Theorem 1.3.** k -PureCLDM₁ is PureSuperQMA-complete for all $k \geq 2$.

We give a proof sketch only. The complete proof can be found in the full version [18].

Proof sketch. Containment holds by Lemma 3.2. It remains to prove that k -PureCLDM₁ is PureSuperQMA = PSQMA-hard (Lemma 3.5). For simplicity, we only prove hardness of 2-PureCLDM₁, but $k > 2$ is analogous since simulatability works for any k .

Our proof closely follows the proof for the QMA-completeness of the (mixed) CLDM problem by Broadbent and Grilo [9] based on *locally simulatable codes*. The broad idea behind their proof is as follows: starting with a QMA verifier V we can use Kitaev’s circuit-to-Hamiltonian construction to build a Hamiltonian H that has low-energy state if and only if V would accept. This low-energy state, the history state, encodes the behavior of the verifier. Broadbent and Grilo’s aim is to construct local density matrices that are consistent with the history state if V accepts, and are inconsistent otherwise.

The main difficulty is that the behavior of the verifier, and hence the history state, depend on the proof given to the verifier. This is solved by considering a different verifier V' that mirrors the actions of V but acts on *encoded* data. By choosing an s -simulatable code for this encoding this has as crucial advantage that the s -local density matrices at any point during the computation can be computed efficiently.⁸

One technical detail that will be especially important to us is that the new verifier V' needs to check whether the proof it received is properly encoded. To ensure that the local density matrices can still be computed during this checking procedure, Broadbent and Grilo have the prover send the original proof encoded under a one-time pad and the one-time pad keys. It is this one-time pad constructions that causes the Broadbent Grilo proof to fail for PureCLDM as a cheating prover could send different pure states for each key, making the full state behave like a mixed state.

We will now sketch our construction. Highlighting especially where we deviate from Broadbent and Grilo. For the technical details and proofs we refer to the full version [18].

Consider a problem $A \in \text{PSQMA}$ that is decided by a super-verifier $\mathcal{V} = \{V_{x,i}\}_{i \in [m]}$ on n_1 proof qubits and n_2 ancilla qubits. We will now describe how \mathcal{V} is reduced to a PureCLDM instance that is consistent if the super-verifier accepts and inconsistent if it rejects.

First, we define a modified super-verifier $\mathcal{V}^{\text{otp}} = \{V_{x,i}^{\text{otp}}\}_{i \in [m']}$ that accepts a one-time padded proof. A crucial difference with Broadbent and Grilo is that we use the same key to encrypt every qubit. The expected proof will be of the form

$$|\psi^{\text{otp}}\rangle = \frac{1}{2} \sum_{a,b \in \{0,1\}} (X^a Z^b)^{\otimes n_1} |\psi_{ab}\rangle |abab\rangle, \quad (3)$$

where the $|\psi_{ab}\rangle$ are proofs accepted by \mathcal{V} that can (but do not have to) depend on a, b . Every individual qubit will still be maximally mixed, but there are only 4 different keys. By projecting the key register down to fixed values of a, b , the modified super-verifier can check that $|\psi_{ab}\rangle$ is indeed accepted by \mathcal{V} . As there are only a 4 different key values (rather than exponentially many for BG), the super-verifier can include checks for each of the possible projections.

Next, we unify all constraints in \mathcal{V}^{otp} into a single circuit that we can apply a circuit-to-Hamiltonian construction to, and, following Broadbent and Grilo, make this circuit act on encoded data using an s -simulatable code.

► **Definition 4.1** (s -simulatable code [9]⁹). *Let \mathcal{C} be an $[[N, 1, D]]$ -QECC, and \mathcal{G} a universal gateset, such that for each logical gate $G \in \mathcal{G}$ on k_G qubits, there exists a physical circuit $U_1^{(G)}, \dots, U_\ell^{(G)}$ with $\ell \in \text{poly}(N)$ that implements G with the help of an m_G -qubit magic state τ_G . We say \mathcal{C} is s -simulatable if there exists a deterministic $2^{O(N)}$ -time algorithm $\text{Sim}_{\mathcal{C}}(G, t, S)$ with $G \in \mathcal{G}$, $t \in \{0, \dots, \ell\}$, $S \subseteq [N(m_G + k_G)]$, $|S| \leq s$, and output $\rho(G, t, S)$, such that for any k_G -qubit state σ*

$$\rho(G, t, S) = \text{Tr}_{\bar{S}} \left((U_t^{(G)} \dots U_1^{(G)}) \text{Enc}(\sigma \otimes \tau_G) (U_t^{(G)} \dots U_1^{(G)})^\dagger \right), \quad (4)$$

where $\text{Enc}(\rho)$ denotes the encoding of ρ under \mathcal{C} .

► **Lemma 4.2** ([15, 9]). *For every $k > \log(s + 3)$, the k -fold concatenated Steane code is s -simulatable.*


⁸ Note that these no longer depend on the proof. Intuitively, all information contained in the proof has been moved to the correlations between larger sets of qubits.

⁹ Our definition deviates from [9, Definition 4.1] in that it does not require transversal circuits for all gates. In fact the simulatable codes in [9] also use non-transversal gates for the T -gadgets given in [9, Section 4.3].


83:12 On the Pure-State Consistency of Local Density Matrices Problem

The unified circuit $V_x^{(s)}$ expects the proof to be encoded with \mathcal{C} , such that \mathcal{C} is $(3s + 2)$ -simulatable. $V_x^{(s)}$ has a proof register B on n'_1 logical qubits, and an ancilla register on $n''_2 > n'_2$, where the additional ancilla qubits are used as resources for the T-gates. $V_x^{(s)}$ is defined in Figure 2


The way we combine the $V_{x,i}^{\text{otp}}$ might seem puzzling at first: we run the encoded version of $V_{x,i}^{\text{otp}}$, decode its output qubit, encode it again and run the encoded version of $V_{x,i}^{\text{otp}\dagger}$. While this acts as the identity, we will be able to *extract* the acceptance behavior of the $V_{x,i}^{\text{otp}}$ from the history state of $V_x^{(s)}$.



(a) $V_x^{(s)}$ first checks whether the proof is encoded, then generates resource states, and finally checks all constraints of \mathcal{V}^{otp} .



(b) ChkEnc successively encodes and decodes each qubit. The annotation “ $I/2$ ” denotes the expected reduced density matrix of that qubit. U_{dec} is a unitary such that $U_{\text{dec}} \text{Enc}(|\psi\rangle) = |\psi\rangle \otimes |0^{N-1}\rangle$ and $U_{\text{enc}} = U_{\text{dec}}^\dagger$.



(c) To verify $V_{x,i}^{\text{otp}}$, first apply the verifier in the code space, then decode the output qubit, which is maximally mixed for acceptance probability $1/2$, and finally undo these steps again.

■ **Figure 2** Super-verifier $V_x^{(s)}$. Wires represent *logical* qubits under \mathcal{C} . Note that while the Chk procedures act as identity, their outcome can be read from the history state via the Extraction Lemma.

Unlike Broadbent and Grilo, we use the 2-local circuit-to-Hamiltonian construction from [19] to obtain a Hamiltonian H .¹⁰ Let us now briefly state an important property of the 2-local construction and the Extraction Lemma, which we use to extract the acceptance probabilities of the $V_{x,i}^{\text{otp}}$ from the history state.

¹⁰ Broadbent and Grilo use Kitaev’s original 5-local construction.

► **Lemma 4.3** (Variant of [19, Lemma 5.1]). *Let $V = U_T \cdots U_1$ be a quantum circuit on n qubits that only consists of 1-local gates and CZ gates followed and preceded by two Z gates, $n_1 \leq n$, and $\varepsilon \in n^{-O(1)}$. Then there exists a 2-local Hamiltonian H on $n + T$ qubits and $\|H\| \in \text{poly}(n)$, such that $\langle \phi | H | \phi \rangle = 0$ for all $|\phi\rangle \in \mathcal{S}_{\text{hist}}$, and if $\text{Tr}(H\rho) \leq 1$, then there exists a state σ in $\mathcal{S}_{\text{hist}}$ (pure if ρ is pure), such that $\|\rho - \sigma\|_{\text{tr}} \leq \varepsilon$, where for $|\hat{t}\rangle = |1^t 0^{T-t}\rangle$,*

$$\mathcal{S}_{\text{hist}} := \text{Span} \left\{ \sum_{t=0}^T U_t \cdots U_1 |x, 0^{n_2}\rangle |\hat{t}\rangle \mid x \in \{0, 1\}^{n_1} \right\}. \quad (5)$$

Note that circuits consisting of only 1-local gates and CZ gates followed and preceded by two Z gates are universal. We can thus assume that all circuits we consider have this form.

► **Lemma 4.4** (Extraction Lemma). *Let $|\psi_{\text{hist}}\rangle = (T + 1)^{-1/2} \sum_{t=0}^T |\psi_t\rangle |\hat{t}\rangle \in \mathcal{S}_{\text{hist}}$ as defined in Lemma 4.3 with $|\psi_t\rangle = U_t \cdots U_1 |\phi, 0^{n_2}\rangle$ for some $|\phi\rangle \in \mathbb{C}^{2^{n_1}}$. Let $i, j \in [T]$, $U_j = I$, $S = \{i, n + j\}$, and $\rho = \text{Tr}_{\bar{S}}(|\psi_{\text{hist}}\rangle\langle\psi_{\text{hist}}|)$. There exists a linear function f_{ex} (independent of the circuit) such that $f_{\text{ex}}(\rho) = \text{Tr}_{\bar{i}}(|\psi_j\rangle\langle\psi_j|)$ and $\|f_{\text{ex}}(A)\|_{\text{tr}} \leq (T + 1)\|A\|_{\text{tr}}$ for all A .*

Now suppose we have a state $|\psi\rangle$ approximately consistent with a set of 2-local density matrices $\{\rho_{ij}\}$ that has low energy with respect to the 2-local Hamiltonian $H = \sum_{ij} H_{ij}$ of Lemma 4.3. Thus $\text{Tr}(H|\psi\rangle\langle\psi|) \approx \sum_{ij} \text{Tr}(H_{ij}\rho_{ij}) \leq 1$ and $|\psi\rangle$ must be close to a history state. In general, we do not know how to compute local density matrices of an accepting history state. However, by the properties of the s -simulatable code, these *can* be computed for the encoded portion of the circuit. Using the Extraction Lemma, we also know the 1-local density matrices of an output qubit after decoding.

Let $|\psi_{\text{hist}}^{(s)}\rangle$ be the history state for $V_x^{(s)}$, such that the input $|\psi^{\text{otp}}\rangle$ satisfies every constraint in \mathcal{V}^{otp} perfectly:

$$|\psi_{\text{hist}}^{(s)}\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t \cdots U_1 |\psi^{\text{otp}}\rangle |0^{n_2'}\rangle |\hat{t}\rangle. \quad (6)$$

Using the properties of the simulatable code, it can be argued that reduced density matrices of $|\psi_{\text{hist}}^{(s)}\rangle$ can be classically efficiently computed.

▷ **Claim 4.5.** There exists a deterministic classical algorithm $\text{Sim}_{V^{(s)}}(x, S)$ that outputs a classical description of $\text{Tr}_{\bar{S}}(|\psi_{\text{hist}}^{(s)}\rangle\langle\psi_{\text{hist}}^{(s)}|)$ for $S \subset [N(n_1' + n_2') + T]$ with $|S| \leq s$ in time $\text{poly}(2^N, T)$.

To apply the Extraction Lemma we would like to force any state consistent with the output of the reduction to be close to a history state. To achieve this we use the following trick: the reduction computes the energy of the reduced density matrices it would output (this is possible since H is 2-local). If this energy is too high, the reduction simply outputs a trivial “No” instance¹¹. If the energy is sufficiently small, we can show that any consistent state must be close to a history state as desired.

We have now assembled all the pieces and are ready to sketch the actual reduction and hardness proof. Let $A \in \text{PSQMA} = \text{PureSuperQMA}$ be as at the beginning of Section 4, and let $V_x := V_x^{(2)} = U_T \cdots U_1$ be as defined below Lemma 4.2. Given an instance $x \in \{0, 1\}^n$, we compute a 2-PureCLDM instance as follows:

¹¹We would like to reject here, but as we are describing the reduction we cannot directly do this. Instead, we output a trivial “No” instance which has the same effect.

83:14 On the Pure-State Consistency of Local Density Matrices Problem

1. Compute the Hamiltonian $H = \sum_{i,j \in [n']} H_{i,j}$ with $n' := N(n'_1 + n'_2) + T$ by applying Lemma 4.3 to V_x for $\varepsilon_1 \in n^{-O(1)}$ (to be determined in the soundness proof).
2. For all $i, j \in [n']$, compute $\rho_{ij} = \text{Sim}_{V^{(s)}}(x, \{i, j\})$.
3. Compute $E = \sum_{i,j} \text{Tr}(H_{ij} \rho_{ij})$.
 - a. If $E = 0$, output $\{\rho_{ij}\}_{i,j \in [n']}$ and β (to be determined in the soundness proof).
 - b. Otherwise, output a NO-instance for 2-PureCLDM.

This reduction clearly runs in $\text{poly}(n)$ time.

If $x \in A_{\text{yes}}$, then the history state $|\psi_{\text{hist}}^{(2)}\rangle$ defined in Equation (6) is consistent with density matrices ρ_{ij} by Claim 4.5.

Using the Extraction Lemma it can be argued that if $x \in A_{\text{no}}$, that is, if there is a check that fails, *and* the 2-local density matrices are low-energy, then there is no consistent state, exactly as desired. Note that the fact that the local density matrices are low-energy is enforced by the reduction (if they are not, a trivial NO instance is output). ◀

5 k -PureCLDM is in PSPACE

In this section we set out to prove the containment of PureSuperQMA in PSPACE. However, the techniques we develop for this result turn out to be much more general and lead to a poly-logarithmic time parallel algorithm (i.e., NC) for solving quadratic systems with few constraints.

► **Theorem 1.8.** $\text{PureSuperQMA} \subseteq \text{PureSuperQMA}(\text{poly}, 0, \frac{1}{\text{poly}}) \subseteq \text{PSPACE}$.

At a high level, the proof consists of three steps. First, we reduce PureSuperQMA to deciding if $p(Q(X))$ has any zeros. Here $p: \mathbb{R}^{\text{poly}(n)} \rightarrow \mathbb{R}$ is a degree $\text{poly}(n)$ polynomial in $\text{poly}(n)$ variables and $Q: \mathbb{R}^{\text{exp}(n)} \rightarrow \mathbb{R}^{\text{poly}(n)}$ is a quadratic polynomial. Next, we use results by Grigoriev and Pasechnik [14] that reduce finding zeros of $p(Q(X))$ to finding limits of zeros of smaller systems. There will be $\text{exp}(n)$ smaller systems, each consisting of $\text{poly}(n)$ equations of degree at most $\text{exp}(n)$. Crucially, these equation will have at most $\text{poly}(n)$ variables. We modify this reduction to work in PSPACE. We then compute these limits efficiently in parallel using an algorithm for the first-order theory of the reals by Renegar [27], making the total computation a PSPACE computation and proving the theorem.

5.1 PureSuperQMA as a polynomial system

We first describe how to reduce PureSuperQMA to a *GP system*.

► **Definition 5.1** (Grigoriev-Pasechnik system (GP system) [14]). *A GP system is a polynomial of the form $p(Q(X))$ where $p: \mathbb{R}^k \rightarrow \mathbb{R}$ is a degree d polynomial and $Q = (Q_1, \dots, Q_k): \mathbb{R}^N \rightarrow \mathbb{R}^k$ is a quadratic map. Both are assumed to have coefficients in \mathbb{Z} .¹² We say a GP system is satisfiable if there exists $x \in \mathbb{R}^N$ such that $p(Q(x)) = 0$.*

Consider a super-verifier $\mathcal{V} = \{(V_i, r_i, s_i)\}_{[k]}$ where $k = \text{poly}(n)$. In this section we will let V_i denote the POVM measurement operator implemented by the circuit so that the acceptance probability $p(V, \psi) = \langle \psi | V_i | \psi \rangle$. In the YES-case, there exists some quantum state $|\psi\rangle \in \mathbb{C}^{2^n}$ such that for all i ,

$$|\langle \psi | V_i | \psi \rangle - r_i| \leq s_i \tag{7}$$

¹²The techniques by Grigoriev and Pasechnik are valid in the more general case where \mathbb{R} is replaced by an arbitrary real closed field \mathbb{K} and \mathbb{Z} with a computable subring of that field.

whereas in the NO-case there will be at least one i for which $|\langle \psi | V_i | \psi \rangle - r_i| > s_i$. Note that depending on the δ parameter in Definition 3.1, there could be more such i but this is not required for our methods. What *is* required for our method to work is that k , the total number of different (V_i, r_i, s_i) the super-verifier could output, is polynomial.

To distinguish the two cases, we write the n -qubit state $|\psi\rangle$ as an exponentially long vector in \mathbb{C}^N . Here and in the rest of this section $N = 2^n$. For each entry of the vector we introduce two variables: one for the real part and one for the complex part. That is, we write

$$|\psi\rangle = \begin{pmatrix} a_1 + b_1 i \\ \vdots \\ a_N + b_N i. \end{pmatrix} \tag{8}$$

Note that for any assignment of the variables $(a_1, b_1, \dots, a_N, b_N) \in \mathbb{R}^{2N}$ we get an (unnormalized) vector in $(\mathbb{C}^2)^{\otimes n}$. We now construct a GP system that is satisfiable if and only if an accepting proof state exists.

First, to ensure that the a_j, b_j represent a normalized quantum state we define

$$Q_0(a_1, b_1, \dots, a_N, b_N) = \|\psi\|^2 - 1 = \left(\sum_{j=1}^N a_j^2 + b_j^2 \right) - 1 \tag{9}$$

Next, note that $\langle \psi | V_i | \psi \rangle$ is already a quadratic equations. Although $|\psi\rangle$ can have a complex part $\langle \psi | V_i | \psi \rangle$ will be real since V_i is positive semidefinite. We define

$$Q_i(a_1, b_1, \dots, a_N, b_N) = \langle \psi | V_i | \psi \rangle - r_i, \quad \text{for } i \in [k]. \tag{10}$$

In order to handle the inequalities in Equation (7) we will add slack variables c_1, \dots, c_k and define $Q_{k+i} = c_i$, for $1 \leq i \leq k$. We now put all constraints together and define

$$p(y_0, \dots, y_{2k}) = y_0^2 + \sum_{j=1}^k (y_j^2 + y_{k+j}^2 - s_j^2)^2. \tag{11}$$

We then have

$$p(Q(a_1, b_1, \dots, a_N, b_N, c_1, \dots, c_k)) = (\|\psi\|^2 - 1)^2 + \sum_{j=1}^k ((\langle \psi | V_j | \psi \rangle - r_j)^2 + c_j^2 - s_j^2)^2. \tag{12}$$

Note that $p(Q(X))$ will be zero only when all component parts are zero. The first term enforces that the norm of $|\psi\rangle$ is 1. Meanwhile, the j -th term in the sum makes sure that $(\langle \psi | V_j | \psi \rangle - r_j)^2$ can be made equal to s_j^2 by adding some nonnegative value c_j^2 . In other words, it ensures that $|\langle \psi | V_j | \psi \rangle - r_j| \leq s_j$. We conclude that $p(Q(a_1, b_1, \dots, a_N, b_N, c_1, \dots, c_k))$ has a zero if and only if there exists some quantum state $|\psi\rangle$ such that for all i , $|\langle \psi | V_i | \psi \rangle - r_i| \leq s_i$, which is exactly as we set out to construct.

5.2 Reducing the number of variables

In this subsection, we consider a general GP system $p(Q(X)) = \zeta$ and describe Grigoriev and Pasechnik's methods for reducing the number of variables. The full version [18] of the work contains a quite thorough treatment of their methods, reproducing large parts of their construction. We hope this makes their methods easily accessible to other quantum theorists.

83:16 On the Pure-State Consistency of Local Density Matrices Problem

Without loss of generality, we can write Q as

$$Q_j : X \mapsto \frac{1}{2} X^\top H_j X + b_j^\top X + c_j, \quad j \in [k], c_j \in \mathbb{Z}, b_j \in \mathbb{Z}^N, H_j = H_j^\top \in \mathbb{Z}^{N \times N}. \quad (13)$$

Define $p_i = \frac{\partial p(Y)}{\partial Y_i}$ for all $i \in [k]$.

For their main construction, Grigoriev and Pasechnik rely on some assumptions. These assumptions hold generically, but could fail in certain degenerate cases. We will later discuss how these assumptions are removed. We assume the following:

- (1) The set $Z = Z(p(Q(X)))$ of zeros of $p(Q(X))$ is bounded, i.e., there is some r such that $\|X\| > r \implies p(Q(X)) \neq 0$. In our case, this assumption is trivially satisfied because of the normalization constraint Equation (9).
- (2) ζ is a regular value of $p(Q(x))$ and $p(Y)$. That is, there exists no $X \in \mathbb{R}^n$ (respectively $Y \in \mathbb{R}^k$) with $p(Q(X)) = \zeta$ (respectively $p(Y) = \zeta$) for which $\nabla p(Q(X)) = 0$ (respectively $\nabla p(Y) = 0$).
- (3) The matrices \hat{H}_i obtained from H_i by deleting the first row are in r -general position. I.e.,

$$\text{rk} \left(\sum_{i=1}^k t_i \hat{H}_i \right) \geq r, \quad \forall t \in \{(p_1(Q(X)), \dots, p_k(Q(X))) \mid X \in Z\}. \quad (14)$$

Note that Assumption 2 implies that t is never 0 in the above.

We are now ready to summarize Grigoriev and Pasechnik's construction as the following theorem.

► **Theorem 5.2** ([14, Theorem 4.5]). *Let p, Q be a GP system satisfying the three assumptions above. Then, one can construct sets $V_c(U, W)$, such that $\bigcup V_c(U, W)$ intersects every connected component of Z .*

The sets $V_c(U, W)$, also called "pieces", are indexed by row sets $U \subseteq \{2, \dots, N\}$ and column sets $W \subset [N]$ such that $r \leq |U| = |W| \leq N - 1$. For such a W , let ϕ_W denote the polynomial

$$\phi_W : \mathbb{R}^N \rightarrow \mathbb{R}^{k+N-|W|}, \quad X \mapsto \begin{pmatrix} Q(X) \\ X_{\overline{W}} \end{pmatrix}. \quad (15)$$

Here $\overline{W} = [N] \setminus W$. For each $V_c(U, W)$, the set $\phi_W(V_c(U, W)) \subseteq \mathbb{R}^k \times \mathbb{R}^{N-|W|}$ is defined as the set of points satisfying the following equations:

$$p(Y) = \zeta, \quad (16a)$$

$$\Omega := \det \Phi(Y)_{UW} \neq 0, \quad (16b)$$

$$\Omega^2 Y = \Omega^2 Q(\phi_{UW}^{-1}(Y, T)), \quad (16c)$$

$$\Omega b(Y)_{\overline{U}} = \Omega \Phi(Y)_{\overline{U}W} \Phi(Y)_{UW}^{-1} \cdot b(Y)_U, \quad (16d)$$

$$\det \Phi(Y)_{U'W'} = 0 \quad \forall U', W' : |U'| = |W'| = |U| + 1, U \subset U' \subset \{2, \dots, N\}, \quad (16e)$$

$$W \subset W' \subset [N].$$

Here $\Phi(Y) = \sum_{j=1}^k p_j(Y) \hat{H}_j$ and $\phi_{UW}^{-1}(Y, T)$ is the inverse of ϕ_W on $\phi_W(V_c(U, W))$, given by

$$\phi_{UW}^{-1} : \mathbb{R}^k \times \mathbb{R}^{N-|W|} \rightarrow \mathbb{R}^N, \quad \begin{pmatrix} Y \\ T \end{pmatrix} \mapsto \begin{pmatrix} \Phi(Y)_{\overline{U}W}^{-1} (b(Y)_U - \Phi(Y)_{U\overline{W}} T) \\ T \end{pmatrix} \quad (17)$$

If $r = N - k + 1$ (which will be ensured later with Lemma 5.4), then there are $N^{O(k)}$ pieces, each defined by $O(k^2)$ polynomials of degree $O(N)$ in $O(k)$ variables.

Lastly, if the coefficients of p and Q are integers of bit length at most L , then ϕ_{UW}^{-1} and the equations defining the $V_c(U, W)$ can be computed from the coefficients of Q and p in time $\text{poly}(\log N, \log d, k, \log L)$ using $(dN)^{O(k)} L^{O(1)}$ parallel processors.

► **Remark 5.3.** Theorem 5.2 still holds with the same runtime and processor bounds when the GP system has coefficients in $\mathbb{Z}[C]$ of degree $O(1)$ (instead of \mathbb{Z}), for real constants $C = (C_1, \dots, C_m) \in \mathbb{R}^m$ and $m = O(k)$. The algorithm then also outputs the equations for the pieces with coefficients in $\mathbb{Z}[C]$. Note that the input and output representations treat these constants symbolically (i.e. like free variables), so that the output is independent of the actual value of C . In the proof of Theorem 5.2 the additional constants can be handled just like the variables.

5.2.1 Removing the assumptions

We now turn to the removal of the assumptions made in Theorem 5.2. To do so, Grigoriev and Pasechnik use limit arguments. The key idea is that if we perturb the initial system of polynomials slightly, there will be at most finitely many values of the perturbation for which the assumptions fail. Hence, if the perturbation is sufficiently small, all assumptions will hold. They then argue that the solutions to our initial system are equal to the solutions of the perturbed system if we let the perturbation go to zero. The assumptions will hold in the limit. We will later discuss how we can compute these limits.

The following lemma deals with the general position assumption.

► **Lemma 5.4** ([14, Lemma 5.2]). *Let p, Q be a GP system and write Q as*

$$Q_j : X \rightarrow \frac{1}{2} X^T H_j X + b_j^T X + c_j, \quad j \in [k], c_j \in \mathbb{Z}, b_j \in \mathbb{Z}^N, H_j = H_j^T \in \mathbb{Z}^{N \times N}. \quad (18)$$

Define $J(j)$ to be the $N \times N$ matrix with diagonal $(1^{j-1}, 2^{j-1}, \dots, N^{j-1})$ and perturb Q by defining

$$\tilde{Q}_j(X, \varepsilon) = Q_j(X) + \frac{\varepsilon}{2} X^T J(j) X. \quad (19)$$

Then, there is some constant ε' such that for $0 < \varepsilon < \varepsilon'$, the matrices $\tilde{H}_j(\varepsilon)$ obtained by deleting the first row of $H_j + \varepsilon J(j)$ will be in $N - k + 1$ general position. In fact, for all $0 \neq y \in \mathbb{R}^k$ the matrix

$$A(y, \varepsilon) = \sum_{j=1}^k Y_j(H_j + \varepsilon J(j)) \quad (20)$$

will have $\text{rk}(A) \geq N - k + 1$.

Next, Grigoriev and Pasechnik deal with the assumption 2 by showing that any $\zeta \neq 0$ that is sufficiently close to zero is a regular value of $p(Q(X))$ and $p(Y)$.

► **Lemma 5.5** ([14, Lemma 5.3]). *Let $p: \mathbb{R}^k \rightarrow \mathbb{R}$ and $Q: \mathbb{R}^n \rightarrow \mathbb{R}^k$ be polynomials. Then $p(Y)$ and $p(Q(X))$ each have at most finitely many critical values.*

We have now established that $p(\tilde{Q}(X, \varepsilon)) - \zeta$ satisfies all assumptions required for Theorem 5.2 when ε, ζ are sufficiently small. In order to use this to find solutions to our original equation we need to relate $Z = Z(p(Q(X)))$ to $\tilde{Z}(\zeta, \varepsilon) = Z(p(\tilde{Q}(X, \varepsilon)) - \zeta)$. Clearly,

if ζ and ε are very small and X^* is a zero of $p(\tilde{Q}(X, \varepsilon)) - \zeta$ then $p(Q(X^*))$ will be close to zero. This could suffice for finding approximate solutions, but to solve the problem exactly we need something more. Grigoriev and Pasechnik provide this by proving that Z coincides exactly with the limit of $\tilde{Z}(\zeta, \varepsilon)$ as $\zeta, \varepsilon \downarrow 0$. They consider the solutions of $p(\tilde{Q}(X, \varepsilon)) - \zeta$ as Puiseux series in ζ and ε and prove that the limits of these are exactly the zeros of $p(Q(X))$. For us, however, it will be more convenient to consider $\tilde{Z}(\zeta, \varepsilon)$ as sets depending on ζ and ε and take the limit of these sets as $\zeta, \varepsilon \downarrow 0$. In the rest of the paper we use the following notion of limits of sets. Strictly speaking, our definition is that of the Kuratowski limit inferior, which of course coincides with the Kuratowski limit if this exists.

► **Definition 5.6.** For $\varepsilon \in (0, 1)$ let $S_\varepsilon \subseteq \mathbb{R}^n$. Then we define

$$\lim_{\varepsilon \downarrow 0} S_\varepsilon = \{X \in \mathbb{R}^n : \forall r \in \mathbb{R}_{>0} \exists \varepsilon_0 \forall \varepsilon < \varepsilon_0 \exists Y \in S_\varepsilon \text{ s.t. } \|X - Y\| < r\}. \quad (21)$$

► **Theorem 5.7** ([14, Theorem 5.4]). For fixed ε, ζ , let $\tilde{Z}(\varepsilon, \zeta)$ be the zeros of $p(\tilde{Q}(X, \varepsilon)) - \zeta$ and let $Z = Z(p(Q(X)))$ be the zeros of $p(Q(X))$. Then we have

$$Z = \lim_{\zeta \downarrow 0} \lim_{\varepsilon \downarrow 0} \tilde{Z}(\varepsilon, \zeta). \quad (22)$$

To test whether Z is empty, it suffices to check whether there is a sequence of nonempty \tilde{Z} .

► **Lemma 5.8.** $Z \neq \emptyset$ iff there exists a sequence $\zeta_n, \varepsilon_n \rightarrow 0$ such that $\tilde{Z}(\varepsilon_n, \zeta_n) \neq \emptyset$.

Proof. If $Z \neq \emptyset$, then by Theorem 5.7, for every sufficiently small $\zeta > 0$, there exists an ε_ζ , such that $\tilde{Z}(\varepsilon, \zeta) \neq \emptyset$ for all $\varepsilon \in (0, \varepsilon_\zeta]$. Thus, we can easily construct the sequence $\zeta_n = O(1/n), \varepsilon_n = \varepsilon_{\zeta_n}$ with $\tilde{Z}(\varepsilon_n, \zeta_n) \neq \emptyset$.

Now given a sequence ζ_n, ε_n with $\tilde{Z}(\varepsilon_n, \zeta_n) \neq \emptyset$, we can select $X_n \in \tilde{Z}(\varepsilon_n, \zeta_n)$, and since $\tilde{Z}(\varepsilon, \zeta)$ is bounded (for sufficiently small ε, ζ), we can apply the Bolzano-Weierstrass theorem to pick a subsequence $X_n \rightarrow X^*$. We have $p(\tilde{Q}(X_n, \varepsilon_n)) = \zeta_n$. Due to continuity it follows $\lim_n p(\tilde{Q}(X_n, \varepsilon_n)) = p(Q(X^*)) = 0$ and $X^* \in Z$. ◀

5.3 Solving the smaller systems using algorithms for the first-order theory of the reals

We now have a way to determine if $p(Q(X))$ has a zero: first, we consider the perturbed system $p(\tilde{Q}(X, \varepsilon)) - \zeta$. By Lemma 5.8 it suffices to check if there is a sequence of nonempty \tilde{Z} . For sufficiently small ζ, ε we can use Theorem 5.2 with Remark 5.3 (setting $C := (\zeta, \varepsilon)$) to find sets $V_c(U, W, \zeta, \varepsilon)$ such that $\bigcup_{U, W} V_c(U, W, \zeta, \varepsilon)$ intersects all connected components of $\tilde{Z}(\zeta, \varepsilon)$. In particular, a sequence of nonempty \tilde{Z} exists iff there exists a sequence of nonempty $\phi_W(V_c(U, W, \zeta, \varepsilon), \varepsilon)$. Here

$$\phi_W : (X, \varepsilon) \mapsto \begin{pmatrix} \tilde{Q}(X, \varepsilon) \\ X_{\overline{W}} \end{pmatrix}. \quad (23)$$

The sets $\phi_W(V_c(U, W, \zeta, \varepsilon), \varepsilon)$ will be defined as the solutions to

$$p(Y) = \zeta, \quad (24a)$$

$$\Omega := \det \Phi(Y, \varepsilon)_{UW} \neq 0, \quad (24b)$$

$$\Omega^2 Y = \Omega^2 \tilde{Q}(\phi_{UW}^{-1}(Y, T, \varepsilon), \varepsilon), \quad (24c)$$

$$\Omega b(Y)_{\overline{U}} = \Omega \Phi(Y, \varepsilon)_{\overline{U}W} \Phi(Y, \varepsilon)_{\overline{U}W}^{-1} \cdot b(Y)_U, \quad (24d)$$

$$\det \Phi(Y, \varepsilon)_{U'W'} = 0 \quad \forall U', W' : |U'| = |W'| = |U| + 1, U \subset U' \subset \{2, \dots, n\}, \quad (24e)$$

$$W \subset W' \subset [n],$$

where we have written $\tilde{Q}_j(X, \varepsilon) = \frac{1}{2}X^\top H_j(\varepsilon)X + b_j^\top X + c_j$ (see Lemma 5.4) and $\Phi(Y, \varepsilon) = \sum_{j=1}^k p_j(Y) \hat{H}_j(\varepsilon)$ for $\hat{H}_j(\varepsilon)$ the matrix $H_j(\varepsilon)$ with the first row deleted.

We will write the existence of such a sequence of $\phi_W(V_c(U, W, \zeta, \varepsilon), \varepsilon)$ as a sentence in the first-order theory of the reals.

► **Definition 5.9** (First-order theory of the reals). *The first-order theory of the reals is concerned with sentences of the form*

$$Q_1 x_1 \in \mathbb{R}^{n_1} Q_2 x_2 \in \mathbb{R}^{n_2} \dots Q_q x_q \in \mathbb{R}^{n_q} P(x_1, \dots, x_q), \quad (25)$$

where

- the Q_i are alternating quantifiers
- P is a quantifier-free Boolean formula in atomic formulas of the form $f(x_1, \dots, x_q) \Delta 0$. Here $f: \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_q} \rightarrow \mathbb{R}$ is a polynomial with integer coefficients and Δ is one of the relation symbols $\leq, <, =, >, \geq, \neq$.

We write the set of all true sentences of this form as $\text{Th}(\mathbb{R})$.

► **Theorem 5.10** ([27, Theorem 1.1]). *There is an algorithm that, given a sentence φ of the form of Equation (25) involving only polynomials with integer coefficients¹³, decides whether $\varphi \in \text{Th}(\mathbb{R})$. The algorithm uses $L^2(md)^{2^{O(q)}} \prod_{i=1}^q n_i$ parallel processors and requires $\log(L) (2^q (\prod_{i=1}^q n_i) \log(md))^{O(1)} + \text{Time}(P)$ time. Here L denotes the maximal bit size of the coefficients of the polynomials in φ , m the number of polynomial (in)equalities in φ , d the maximal degree of these polynomials, and $\text{Time}(P)$ the worst case time required for computing P when the atomic formulas are substituted for Boolean values.*

Combining this result with Lemma 5.8 we get the following result.

► **Theorem 1.13.** *Let p, Q be a GP system. There is a parallel algorithm for deciding whether $p(Q(X))$ has a root. The algorithm uses $L^2(k^2Nd)^{O(k)}$ parallel processors and needs $\text{poly}(\log N, \log d, k, \log L)$ time. In particular, if $\log N, \log d, \log L, k \leq \text{poly}(n)$ the computation can be done in $\text{NC}(\text{poly}) = \text{PSPACE}$. If $N, d, L \leq \text{poly}(n)$ and $k = O(1)$ then the algorithm is in NC.*

Proof. By Lemma 5.8, it suffices to check whether there exists a sequence of $\varepsilon, \zeta \rightarrow 0$, such that $\tilde{Z}(\varepsilon, \zeta) \neq \emptyset$. By Lemma 5.4, assumption (3) is satisfied for sufficiently small ε (independent of ζ). By Lemma 5.5, for every ε there exist only a finite number of critical values (or “bad” ζ). Thus, if there exist a suitable sequence ε, ζ , we may also assume that the assumptions of Section 5.2 hold. By Theorem 5.2, there must be a subsequence $(\varepsilon_n, \zeta_n) \rightarrow 0$ and U, W such that $\tilde{Z}(\varepsilon_n, \zeta_n) \cap V_c(U, W, \varepsilon_n, \zeta_n) \neq \emptyset$ for all n . Therefore $Z \neq \emptyset$ is equivalent to the following FOTR sentence being true for some U, W , which we can test in parallel:

$$\exists_{U, W} := \forall \gamma > 0 \exists \varepsilon, \zeta \in (0, \gamma) \exists Y, T: (Y, T) \in \phi_W(V_c(U, W, \varepsilon, \zeta)), \quad (26)$$

where $(Y, T) \in \phi_W(V_c(U, W, \zeta, \varepsilon))$ is shorthand for (Y, T) satisfies Equation (24), which then implies $\phi_{U, W}^{-1}(Y, T, \varepsilon) \in \tilde{Z}(\varepsilon, \zeta)$. We can test this sentence using Theorem 5.10. The complexity bounds follow from those of Theorems 5.2 and 5.10. ◀

The proof of Theorem 1.8 now follows directly from applying Theorem 1.13 to the discussion in Section 5.1.

¹³To deal with algebraic coefficients we can add additional variables and polynomials enforcing that these variables take the value of the required algebraic numbers.

6 Approximation and Optimization

In the previous section, we showed how to solve the (decision) PureCLDM problem. Next, we give a parallel polynomial time algorithm (i.e. (function) NC(poly)) to compute an approximate description of the consistent state. It is worth emphasizing that our algorithm computes an approximation to an exactly consistent state, and not just a state that is approximately consistent.

If the density matrices are slightly perturbed, there may not be a *perfectly* consistent state. We can compute the state that minimizes the “error” (e.g. the minimum α such that the given instance is a YES-instance per Definition 2.1) by solving OptGP as defined in Definition 1.14.

► **Theorem 1.15.** *Let r, p, Q be an OptGP instance. There is a parallel algorithm that computes a δ -approximation to $\theta = \min_{X \in Z} r(Q(X))$ and to some X^* with $r(Q(X^*)) = \theta$ and $p(Q(X^*)) = 0$. The algorithm uses $\text{poly}(L, |\log \delta|, (kNd)^{O(k^2)})$ parallel processors and needs $\text{poly}(\log N, \log d, k, \log L, \log |\log \delta|)$ time.*

6.1 Computing the minimum

The first step in solving the approximation problem of Theorem 1.15 is to compute θ . For that, we will use Renegar’s algorithm for finding approximate solutions to first-order theory of the reals formulas.

► **Theorem 6.1** ([29, Theorem 1.2]). *Let $\varphi(y)$ be a formula as in Equation (25) involving free variables $y \in \mathbb{R}^l$, let $0 < \varepsilon < r$ be powers of 2 and define $\text{sol}(\varphi, r) = \{y \in \mathbb{R}^l : \varphi(y) \wedge \|y\| \leq r\}$. Then there exists an algorithm that constructs a set $\{y_i\}_i$, such that for every connected component of $\text{sol}(\varphi, r)$, there is at least one y_i within distance ε of the component.*

The algorithm can be implemented using $(L + |\log \varepsilon| + |\log r|)^{O(1)} (md)^{2^{O(q)} l} \prod_{j=1}^q n_j$ parallel processors. It then requires time $[2^{ql} (\prod_{j=1}^q n_j) \log (mdL + |\log \varepsilon| + |\log r|)]^{O(1)} + \text{Time}(P)$. Here m is the number of polynomials in φ , d their maximum degree, L the maximum bit length of the coefficients of these polynomials, q the number of quantifiers and $\text{Time}(P)$ the worst case parallel time to compute P when the atomic formulas are substituted for Boolean values.

It is relatively straightforward to compute θ with the above theorem. To that end, define $p_\theta(Y) := (p(Y))^2 + (r(Y) - \theta)^2$. We will compute θ by finding the minimum θ for which $Z_\theta := Z(p_\theta(Q(X))) \neq \emptyset$. Since Z is compact, the minimum $\theta = \min_{X \in Z} r(Q(X))$ must exist. Hence, there must exist a piece U, W , such that the FOTR sentence $\Xi_{U,W}(\theta)$ is true, where $\Xi_{U,W}(\theta)$ is defined as in Equation (26), but with the additional parameter θ (see Remark 5.3). Note that since θ is minimal, $\Xi_{U,W}(\theta') \wedge \theta' < \theta$ cannot hold. Formally, θ is minimal iff the sentence following sentence is true:

$$\Xi_{U,W} \wedge \forall \theta' : \Xi_{U,W}(\theta') \rightarrow (\theta' \geq \theta) \quad (27)$$

Since t is a free variable in the above sentence, we can approximate it using Theorem 6.1 for all pieces in parallel, and then take the minimum.

6.2 Computing the solution vector

Next, we are going to show how to compute X^* such that $p(Q(X^*)) = 0$ and $r(Q(x^*)) = \theta$ and thereby prove Theorem 1.15. Naively, we could use Theorem 6.1 to compute X^* . However, that will not lead to the desired runtime since the dimension of X is exponentially

large. We modify $\Xi_{U,W}$ from Equation (26) to find “good enough” $(\tilde{Y}, \tilde{T}) \in \tilde{Z}(\varepsilon, \zeta)$ with $\tilde{X} := \phi_{U,W}^{-1}(\tilde{Y}, \tilde{T}, \varepsilon)$, such that there exists some $X \in Z$ with $\|X - \tilde{X}\| \leq \delta$ for some desired accuracy parameter δ , which can be chosen exponentially small in N . The existence of \tilde{Y}, \tilde{T} follows by a similar to the proof of Lemma 5.8.

We will use a univariate representation to compute \tilde{Y}, \tilde{T} exactly and then Theorem 6.1 to obtain approximations to the entries of \tilde{X} in parallel.¹⁴ The next theorem shows how to compute univariate representations of the solutions to an FOTR sentence. It is a straightforward combination of [27, Proposition 2.3.1 and 3.8.1] and [28, Proposition 6.2.2].

► **Theorem 6.2.** *Let $\varphi(y)$ be a formula of the form Equation (25) with free variables $y \in \mathbb{R}^l$. Then, there exists a set $\mathcal{P}(\varphi)$ of $(md)^{2^{O(q)l}} \prod_{j=1}^q n_j$ pairs of polynomials (p, F) where $p: \mathbb{R} \rightarrow \mathbb{R}$ and $F: \mathbb{R} \rightarrow \mathbb{R}^{l+1}$ with the following property. For every connected component C of $\{y \in \mathbb{R}^l: \varphi(y)\}$, there is a $(p, F) \in \mathcal{P}(\varphi)$ such that for some root t^* of p , $\text{Aff}(F(t^*))$ is well defined and in C . Here $\text{Aff}(F(t^*))$ denotes the affine image $\frac{1}{F_{l+1}(t^*)}(F_1(t^*), \dots, F_l(t^*))$*

The sets $\mathcal{P}(\varphi)$ can be constructed in parallel using $L^2(md)^{2^{O(q)l}} \prod_{j=1}^q n_j$ processors and time $(\log L) \left(2^{ql} \left(\prod_{j=1}^q n_j\right) \log(md)\right)^{O(1)}$. Here m is the number of polynomial (in)equalities appearing in φ , and d is their maximum degree.

We have now established all the components to prove Theorem 1.15. The proof can be found in the full version [18].

References

- 1 Scott Aaronson, Salman Beigi, Andrew Drucker, Bill Fefferman, and Peter Shor. The power of unentanglement. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 223–236, 2008. doi:10.1109/CCC.2008.5.
- 2 D. Aharonov and O. Regev. A lattice problem in quantum NP. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 210–219, 2003. doi:10.1109/SFCS.2003.1238195.
- 3 Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, April 2009. doi:10.1017/cbo9780511804090.
- 4 Roozbeh Bassirian, Bill Fefferman, Itai Leigh, Kunal Marwaha, and Pei Wu. Quantum Merlin-Arthur with an internally separable proof, 2024. doi:10.48550/arXiv.2410.19152.
- 5 Roozbeh Bassirian, Bill Fefferman, and Kunal Marwaha. Quantum Merlin-Arthur and proofs without relative phase. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:19, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2024.9.
- 6 Hugue Blier and Alain Tapp. A quantum characterization of NP. *computational complexity*, 21(3):499–510, September 2012. doi:10.1007/s00037-011-0016-2.
- 7 Allan Borodin. On relating time and space to size and depth. *SIAM Journal on Computing*, 6(4):733–744, 1977. doi:10.1137/0206054.
- 8 Fernando G. S. L. Brandão and Aram W. Harrow. Quantum de finetti theorems under local measurements with applications. *Communications in Mathematical Physics*, 353(2):469–506, July 2017. doi:10.1007/s00220-017-2880-3.
- 9 Anne Broadbent and Alex Bredariol Grilo. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022. doi:10.1137/21M140729X.

¹⁴One could also compute \tilde{X} numerically, but then numerical errors need to be considered.

- 10 A. J. Coleman. Structure of fermion density matrices. *Rev. Mod. Phys.*, 35:668–686, July 1963. doi:10.1103/RevModPhys.35.668.
- 11 Bill Fefferman and Cedric Lin. Quantum Merlin Arthur with exponentially small gap, 2016. arXiv:1601.01975.
- 12 Sevag Gharibian, Miklos Santha, Jamie Sikora, Aarthi Sundaram, and Justin Yirka. Quantum generalizations of the polynomial hierarchy with applications to QMA(2). *computational complexity*, 31(2):13, September 2022. doi:10.1007/s00037-022-00231-8.
- 13 Dima Grigoriev. Polynomial complexity of solving systems of few algebraic equations with small degrees. In Vladimir P. Gerdt, Wolfram Koepf, Ernst W. Mayr, and Evgenii V. Vorozhtsov, editors, *Computer Algebra in Scientific Computing*, pages 136–139, Cham, 2013. Springer International Publishing. doi:10.1007/978-3-319-02297-0_11.
- 14 Dima Grigoriev and Dmitrii V. Pasechnik. Polynomial-time computing over quadratic maps i: sampling in real algebraic sets. *Comput. Complex.*, 14(1):20–52, April 2005. doi:10.1007/s00037-005-0189-7.
- 15 Alex B. Grilo, William Slofstra, and Henry Yuen. Perfect zero knowledge for quantum multiprover interactive proofs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 611–635, 2019. doi:10.1109/FOCS.2019.00044.
- 16 Aram W. Harrow and Ashley Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *J. ACM*, 60(1), February 2013. doi:10.1145/2432622.2432625.
- 17 Fernando Granha Jeronimo and Pei Wu. The power of unentangled quantum proofs with non-negative amplitudes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 1629–1642, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585248.
- 18 Jonas Kamminga and Dorian Rudolph. On the complexity of pure-state consistency of local density matrices, 2025. arXiv:2411.03096.
- 19 Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. doi:10.1137/S0097539704445226.
- 20 A. Kitaev, A. Shen, and M. Vyalii. *Classical and Quantum Computation*. American Mathematical Society, May 2002. doi:10.1090/gsm/047.
- 21 Alexander Klyachko. Quantum marginal problem and representations of the symmetric group, 2004. arXiv:quant-ph/0409113.
- 22 Yi-Kai Liu. Consistency of local density matrices is QMA-complete. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 438–449, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11830924_40.
- 23 Yi-Kai Liu, Matthias Christandl, and F. Verstraete. Quantum computational complexity of the N -representability problem: QMA complete. *Phys. Rev. Lett.*, 98:110503, March 2007. doi:10.1103/PhysRevLett.98.110503.
- 24 Chris Marriott and John Watrous. Quantum Arthur–Merlin games. *computational complexity*, 14(2):122–152, June 2005. doi:10.1007/s00037-005-0194-x.
- 25 David A Mazziotti. Pure- N -representability conditions of two-fermion reduced density matrices. *Physical Review A*, 94(3):032516, 2016.
- 26 Attila Pereszlényi. Multi-prover quantum Merlin-Arthur proof systems with small gap, 2012. arXiv:1205.2761.
- 27 James Renegar. On the computational complexity and geometry of the first-order theory of the reals. part I: Introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *Journal of Symbolic Computation*, 13(3):255–299, 1992. doi:10.1016/S0747-7171(10)80003-3.
- 28 James Renegar. On the computational complexity and geometry of the first-order theory of the reals. part II: The general decision problem. preliminaries for quantifier elimination. *Journal of Symbolic Computation*, 13(3):301–327, 1992. doi:10.1016/S0747-7171(10)80004-5.

- 29 James Renegar. On the computational complexity of approximating solutions for real algebraic formulae. *SIAM Journal on Computing*, 21(6):1008–1025, 1992. doi:10.1137/0221060.
- 30 Yaoyun Shi and Xiaodi Wu. Epsilon-net method for optimizations over separable states. *Theoretical Computer Science*, 598:51–63, 2015. doi:10.1016/j.tcs.2015.03.031.
- 31 Tzu-Chieh Wei, Michele Mosca, and Ashwin Nayak. Interacting boson problems can be QMA hard. *Phys. Rev. Lett.*, 104:040501, January 2010. doi:10.1103/PhysRevLett.104.040501.
- 32 Xiao-Dong Yu, Timo Simnacher, Nikolai Wyderka, H. Chau Nguyen, and Otfried Gühne. A complete hierarchy for the pure state marginal problem in quantum mechanics. *Nature Communications*, 12(1):1012, February 2021. doi:10.1038/s41467-020-20799-5.