How Hard Is It to Verify a Classical Shadow?

Georgios Karaiskos
Department of Computer Science, Paderborn University, Germany

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

Johannes Jakob Meyer
Freie Universität Berlin, Germany

Jens Eisert
Freie Universität Berlin, Germany

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

Abstract

Classical shadows are succinct classical representations of quantum states which allow one to encode a set of properties P of a quantum state ρ, while only requiring measurements on logarithmically many copies of ρ in the size of P. In this work, we initiate the study of verification of classical shadows, denoted classical shadow validity (CSV), from the perspective of computational complexity, which asks: Given a classical shadow S, how hard is it to verify that S predicts the measurement statistics of a quantum state? We first show that even for the elegantly simple classical shadow protocol of [Huang, Kueng, Preskill, Nature Physics 2020] utilizing local Clifford measurements, CSV is QMA-complete. This hardness continues to hold for the high-dimensional extension of said protocol due to [Mao, Yi, and Zhu, PRL 2025]. In contrast, we show that for the HKP and MYZ protocols utilizing global Clifford measurements, CSV can be "dequantized" for low-Frobenius norm observables, i.e., solved in randomized poly-time with standard sampling assumptions. Finally, we show that CSV for exponentially many observables is complete for a quantum generalization of the second level of the polynomial hierarchy, yielding the first natural complete problem for such a class.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Quantum complexity theory

Keywords and phrases classical shadows, quantum complexity theory, QMA, quantum polynomial hierarchy

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.123

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2510.08515` [23]

Funding *Jens Eisert*: Supported by the BMFTR (QSolid, Hybrid++, QuSol, MUNIQC-Atoms, QuSol, PasQuops, Hybrid++), the Munich Quantum Valley (K-4 and K-8), the Quantum Flagship (PasQuans2, Millenion), QuantERA (HQCC), the Clusters of Excellence MATH+ and ML4Q, the DFG (CRC 183, SPP 2514), Berlin Quantum, and the ERC (DebuQC).
Sevag Gharibian: Supported by the DFG under grant numbers 563388236 and 450041824, the BMFTR via project PhoQuant (grant number 13N16103), and the project PhoQC from the programme Profilbildung 2020 of the Ministry of Culture and Science of the State of North Rhine-Westphalia.

Acknowledgements We thank Asad Raza for helpful discussions.

53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026).
Editors: Sayan Bhattacharya, Danupon Nanongkai, Michael Benedikt, and Gabriele Puppis;
Article No. 123; pp. 123:1–123:23

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Fully classically describing a quantum state ρ has long been known to require exponential overhead, making characterizing the outputs of quantum devices a challenging task. Indeed, for a state ρ on n qubits, i.e., of dimension $D = 2^n$, a sample complexity of $\Theta(D^2)$ copies of ρ are known to be necessary and sufficient for full quantum state tomography [30, 18]. In general, however, one is not necessarily interested in learning *everything* about ρ, but only a specific set P of properties. Formally, we may model these as a set of M measurement operators $P = \{ P_i \}_{i=1}^{M}$, where one is interested in computing $\mathrm{Tr}(\rho P_i)$. The natural question is now: *Can one avoid full state tomography in this case?*

In 2018, Aaronson showed [1] the answer is *yes*: for set P of 2-outcome measurements, given k copies of ρ, one can produce estimates $b_1, \ldots, b_M \in [0, 1]$ such that with probability at least $1 - \delta$, one has $|\mathrm{Tr}(P_i\rho) - b_i| \leq \epsilon$ *for all* i. The magic here is that the *sample complexity*, k, can be chosen polylogarithmic in the dimension D and number of measurements M, i.e.,

$$k \in O\left(\log \frac{1}{\delta} \cdot \epsilon^{-5} \cdot \log^4 M \cdot \log D\right). \tag{1}$$

While this original protocol was not yet *time* efficient, it did not take long for the latter to be rectified, e.g., Brandão, Kalev, Li, Lin, Svore, Wu [7]. Indeed, soon after *Huang, Kueng and Preskill* (HKP) discovered [21] a remarkably simple and efficient classical shadow tomography procedure, which randomly samples unitary U from an "appropriate" ensemble $\mathcal{U}$ of unitaries, and measures $U\rho U^\dagger$ in the standard basis. Roughly, the resulting string can be thought of as a "snapshot" of ρ, and the set of all snapshots constitutes the classical shadow, S. A recovery procedure via median-of-means is then specified, so that given S, one can recover estimates for $\mathrm{Tr}(\rho P_i)$. For general P and $\mathcal{U}$, the procedure has sample complexity

$$O\left(\frac{\log(M)}{\epsilon^2} \max_{1 \leq i \leq M} \left\| P_i - \frac{\mathrm{Tr}(P_i)}{2^n} I \right\|_{\mathrm{shadow}}^2\right), \tag{2}$$

where the shadow norm depends on P and $\mathcal{U}$ (see Section 2 for details on the HKP protocol.) For $\mathcal{U}$ the set of global Cliffords and the set of k-local Cliffords, the shadow norm in Equation (2) is at most $3\,\mathrm{Tr}(P_i^2)$ and $4^k \|P_i\|_\infty^2$, respectively. Thus, for example, to predict measurement results for P the set of k-local Pauli strings[1], one obtains a sample *and* time efficient[2] protocol, which requires only $\log(M)$ copies, and as a bonus needs only to measure a single copy at a time.

Verifying classical shadows. This work initiates the study of the natural question:

> *Given as input a "classical shadow" S, what is the complexity of verifying that S actually "predicts" the measurement statistics of some ρ against P?*

As stated, this question is ill-posed, in the sense that we are not aware of a formal definition of a "classical shadow" in the literature. Thus, to remedy this, we first provide a general formal definition:

▶ **Definition 1.1** (Classical shadow). *A shadow on n qubits is a 4-tuple (S, O, A, χ), where*

- *(Shadow) $S = \{s_i\}_{i=1}^{N}$ is a multi-set of $N = \mathrm{poly}(n)$ strings, each of length $\mathrm{poly}(n)$.*

[1] A Pauli string Q is an element of $\{ I, X, Y, Z \}^n$. We say Q is k-local if it contains at most k non-identity terms.

[2] Since k-local Clifford measurements $\mathcal{U}$ are easy to implement.

- *(Observables)* $O = \{O_i\}_{i=1}^m$ *is a set of* n*-qubit observables satisfying* $\|O_i\|_\infty \leq 1$*, where* $1 \leq m \leq 2^{p(n)}$ *for polynomial* p*. Given index* i*, a* $\text{poly}(n)$*-bit description of* O_i *can be produced in* $\text{poly}(n)$*-time*[3]*. Moreover, there exists a* $\text{poly}(n)$*-time quantum algorithm which, for any* O_i *and any* n*-qubit state* ρ*, applies*[4] *measurement* O_i *to* ρ*.*
- *(Recovery algorithm)* A *is a* $\text{poly}(n)$*-time classical algorithm which, given* S *and* $i \in [m]$*, produces real number* $A(S, i) \in [-1, 1]$ *within* $\chi = \text{poly}(n)$ *bits of precision.*

This definition says nothing about prediction accuracy; it simply formalizes the idea that a "classical shadow" is a multi-set of strings S, *in principle* obtained via some set of efficient measurements on copies of a physical state ρ, coupled with a set of target observables O and an efficient recovery procedure A for "extracting predictions". An alternate possible definition might be not to give shadow S as a fixed sequence of strings, but rather to generate S on-the-fly by sampling from some unknown distribution (thus capturing the idea of measurement bases $\mathcal{U}$ as in HKP). In the full version [23], we formalize this sampled-shadow variant and show the complexity of verifying "classical shadows" versus "sampled classical shadows" is equivalent under randomized reductions. For simplicity, we thus work with Definition 1.1, as its input model is the standard one used in (e.g.) BQP and QMA.

Moving on, the task of checking the *validity* of a shadow, i.e., that the outputs of A correctly predict measurement statistics, is formalized as:

▶ **Definition 1.2** (Classical Shadow Validity (CSV))**.** *Given classical shadow* (S, O, A, χ)*, parameters* α *and* β *satisfying* $\beta - \alpha \geq 1/\text{poly}(n)$*, decide between the following two cases:*

- ***Yes****:* $\exists$ n*-qubit state* ρ *s.t.* $\forall\ i \in [m]$*,* $|\text{Tr}(O_i\rho) - A(S, i)| \leq \alpha$*.*
- ***No****:* $\forall$ n*-qubit states* ρ $\exists$ *some* $i \in [m]$ *s.t.* $|\text{Tr}(O_i\rho) - A(S, i)| \geq \beta$*.*

Since $A(S, i) \in [-1, 1]$ *and* $\|O_i\|_\infty \leq 1$*, it is natural to assume* $0 \leq \alpha < \beta \leq 2$*.*

The theme of this work is to characterize the complexity of this problem and its variants, including for the HKP protocol with local Clifford measurements, "dequantization" results for global Clifford measurements, and the case of exponentially many observables.

Comparison to and distinction from CONSISTENCY problem. Before proceeding, the reader familiar with quantum complexity theory may notice that, at least in the setting of *polynomially* many observables O_i, CSV is eerily similar to the QMA-complete CONSISTENCY problem of Liu [27]. In the latter, the input is a set of k-local reduced states ρ_i acting on a subset S_i of $k \in O(1)$ out of n qubits each, and the question is whether there exists an n-qubit state ρ such that for all i, $\text{Tr}_{[n]\setminus S_i}(\rho) \approx \rho_i$? Indeed, as our definition of classical shadows is intentionally very general, it includes as a special case the CONSISTENCY problem. From this, one immediately obtains that CSV is at least QMA-hard (Corollary 3.7). This is *not* the point of this paper!

The point is that classical shadow protocols used in practice typically do *not* produce local density operators as in CONSISTENCY, but rather highly non-local snapshots (e.g. n-local operators which are the tensor product of n non-trivial single qubit states)! Our goal is thus to characterize the complexity of CSV for precisely these experimentally relevant snapshots, to which the QMA-hardness of CONSISTENCY does not obviously apply. Indeed, as will be discussed shortly, a reduction from CONSISTENCY to CSV will have to overcome the well-known challenging problem of how to construct a *global* quantum snapshot from *local* overlapping reduced density operators (as in CONSISTENCY).

[3] This is the succinct access assumption.
[4] Formally, we can efficiently measure in the eigenbasis of O_i, and return the eigenvalue corresponding to the measurement result.

Motivation and application to near-term devices. (1) As near-term experimental devices remain noisy, verifying that a device *actually* outputs the state intended remains a major challenge. Well-known examples include verification of quantum advantage experiments such as Random Circuit Sampling [6] or Boson Sampling [2, 20]. In the same vein, it is arguably important to verify the validity of snapshots output by classical shadow experiments; this is true even if the experimenter fully trusts that the device has not been tampered with. More generally, in the distributed cloud setting where the user does *not* trust the device, CSV becomes yet more crucial.

(2) The study of CSV is important to the study of classes QMA versus QCMA (i.e. QMA with a classical proof [5]), as it underpins the central open question of whether quantum witnesses are fundamentally more powerful than classical ones. Specifically, *if* a QMA verifier's measurement falls into a class of observables O whose output statistics could be efficiently predicted by poly-size classical shadows, *and* if CSV for said shadows could be solved by a (uniformly generated) poly-time quantum circuit, then QMA = QCMA, which would be a breakthrough.

(3) We further motivate the study of CSV by framing it as a natural quantum analogue of the classical Sparse Representation problem (SR) under the lens of classical shadow protocols. In SR, one has to decide if a sparse vector, consistent with a given measurement sketch, exists. While SR is NP-hard in the general case [12], it becomes efficiently solvable (via convex relaxation) when the measurement matrix satisfies the Restricted Isometry Property (RIP) [9]. In the quantum setting an analogous compressed sensing phenomenon is known: under a low-rank promise on the state and suitable measurements, convex programs can efficiently reconstruct the state [17]. Our work studies the quantum SR question arising in the classical shadow framework, [1, 21].

Our results. We organize our discussion[5] in terms of (1) polynomially many observables, (2) exponentially many observables, and (3) further variants of CSV with connections to QMA(2). For clarity, our main results involve (1) and (2). All hardness results are under poly-time many-one reductions.

1. Polynomially many observables: Hardness and dequantization. As previously stated, it is not difficult to see that CSV in its most general form is QMA-complete (Corollary 3.7). Here, we focus on the more challenging case of the HKP protocol [21] (instantiated with either local or global Clifford measurements), as well as a high-dimensional generalization thereof due to Mao, Yi and Zhu (MYZ) for odd-prime local dimension d [29].

To begin, we define $\mathsf{CSV}_{\mathsf{HKP}}$ as CSV for the HKP protocol instantiated with *local* Clifford measurements (Definition 4.2); roughly, the elements of S are n-bit strings, conjugated by Pauli strings in $\{X, Y, Z\}^n$, and the observables are k-local Pauli strings for $k \in O(1)$. We show the following statement.

▶ **Theorem 1.3** (Informal; see Proposition 3.4, Remark 4.1,Corollary 4.8)**.** $\mathsf{CSV}_{\mathsf{HKP}}$ *is* QMA*-complete, even for* 6*-local observables on a spatially sparse hypergraph.*

In words, deciding if a given HKP classical shadow based on local Clifford measurements is valid is intractable, even when the observables are 6-local and essentially arranged on a line (formally on a spatially sparse hypergraph (in the sense of Ref. [31]; Definition 2.1)).

[5] We remark that although we gave a fully general formal definition of classical shadows (Definition 1.1), most of our results are actually independent of the specific recovery algorithm A employed therein; thus, behind the scenes we often work with a simpler restatement of CSV, denoted Observable Consistency (ObsCon, Definition 3.1). Hence, while we informally state our results in terms of CSV here, our formal statements are often in terms of ObsCon.

Specifically, the hardness construction may be viewed as 1D nearest-neighbor on qudits of dimension 8. Each qudit is then decomposed into 3 qubits, and neighboring pairs of qudits (q_i, q_{i+1}) have a 6-local observable acting jointly on their constituent qubits.

Defining $\mathsf{CSV}_{\mathsf{MYZ}}$ (Definition 4.9) analogously for MYZ on odd prime local dimensions d, we next show:

▶ **Theorem 1.4** (Informal; see Proposition 3.4, Remark 4.1, Corollary 4.11)**.** $\mathsf{CSV}_{\mathsf{MYZ}}$ *is* QMA*-complete for every fixed odd prime local dimension* $d \geq 11$*, even for* 2*-local nearest-neighbor observables on a line.*

Here, since we are allowed to work with larger d, we cleanly obtain hardness with all observables acting on pairs of nearest neighbor qudits (q_i, q_{i+1}).

Finally, we study CSV for the HKP protocol instantiated with *global* Clifford measurements, denoted $\mathsf{CSV}_{\mathsf{GC}}$ (Definition 4.12). We show a "dequantization" result as follows, for Frobenius norm $\|A\|_F = \sqrt{\mathrm{Tr}(A^\dagger A)}$:

▶ **Theorem 1.5** (Informal (see Theorem 4.19))**.** $\mathsf{CSV}_{\mathsf{GC}}$ *is solvable in polynomial classical randomized time if (a)* $\|O_i\|_F \leq \mathrm{poly}(n)$ *for all* i*, and (b) we are given sampling and query access to each* O_i*.*

First, while the Frobenius norm bound above may *a priori* seem strong, this setting captures natural tasks such as (e.g.) fidelity estimation against pure and low rank states [21]. Moreover, the *global* Clifford HKP protocol itself is efficient precisely in this regime, i.e. when $\|O_i\|_F \leq \mathrm{poly}(n)$, coinciding with the regime in which Theorem 1.5 dequantizes $\mathsf{CSV}_{\mathsf{GC}}$. By a similar argument, this "dequantization" result extends to the MYZ protocol with *global* Clifford measurements, under the same assumptions (Theorem 4.21). Second, we dub this a "dequantization" result, in that conditions (a) and (b) are those in previous dequantization works, e.g., Tang [32] and Chia, Gilyén, Li, Lin, Tang and Wang [10], allowing randomized linear algebra techniques to be employed.

2. Exponentially many observables. We next consider CSV with *exponentially* many observables. Although *a priori* this setting may seem unrealistic, King, Gosset, Kothari and Babbush gave [24] an explicit *polynomial*-sample complexity shadow protocol for the set of all 4^n Pauli string observables, $\{I, X, Y, Z\}^n$. (Note the time complexity is still exponential, but in our setting, we do not produce the shadow, but receive it as input; thus, this overhead is not relevant.) What is also relevant is that Ref. [24] gives a *poly*-time recovery algorithm for the observable expectations, *assuming* one only demands *constant* additive error. We show the following statement.

▶ **Theorem 1.6** (Informal; follows from Lemma 3.17 and Corollary 3.15)**.** *CSV for exponentially many observables and constant additive error recovery precision is* $\mathsf{qc}\text{-}\Sigma_2$*-complete.*

Let us discuss strengths and weaknesses: The strengths are that (1) the result holds even if one need only recover constant precision approximations of observable predictions, and (2) Theorem 1.6 yields the first natural complete problem for a quantum generalization of (a level of) the *polynomial hierarchy* (PH). Specifically, $\mathsf{qc}\text{-}\Sigma_2$ (Definition 2.8) is a quantum generalization of Σ_2^p, the second level of PH, in which the first proof is a mixed quantum state, the second a classical string, and the verifier is quantum. We remark this is the first work studying $\mathsf{qc}\text{-}\Sigma_2$, though other variants of quantum PH have been studied in previous works [13, 14, 16, 3]. The weakness is that, unlike Theorem 1.3, we do not prove the result for the specific observable set of Ref. [24], i.e., for $\{I, X, Y, Z\}^n$.

3. Further variants of CSV *and connections to* QMA(2). For completeness, we also show the following for variants of CSV:

1. CSV where the consistent state must be *product*, i.e., $\rho = \rho_A \otimes \rho_B$, is QMA(2)-complete[6] for polynomially many observables, and qc-Σ_2(2)-complete (Definition 2.9) for exponentially many observables. As an intermediate step, the proof shows that $\mathsf{SuperQMA}(2)_{\mathrm{poly}} = \mathsf{QMA}(2)$, where $\mathsf{SuperQMA}(2)_{\mathrm{poly}}$ is a product state generalization of QMA^+ from Ref. [4].
2. Verifying if a *set* of classical shadows, each possibly with different observables, all correspond to the *same* state ρ is QMA-complete and qc-Σ_2-complete for polynomially many and exponentially many observables, respectively.

For these results see the full version of the paper [23].

Techniques. We focus on QMA-completeness of $\mathsf{CSV}_{\mathsf{HKP}}$ (Theorem 1.3) and completeness for qc-Σ_2 of CSV with exponentially many observables (Theorem 1.6).

QMA-completeness of $\mathsf{CSV}_{\mathsf{HKP}}$*.* Ideally, we wish to reduce the QMA-complete CONSISTENCY problem on k-local reduced density operators ρ_i to $\mathsf{CSV}_{\mathsf{HKP}}$. The challenge? Each ρ_i acts on only k-qubits. HKP classical shadows with local Clifford measurements, on the other hand, have shadow elements s_i which are highly *non-local* – each s_i is an n-qubit tensor product of eigenvectors of single-qubit Pauli matrices. And "stitching" together local information, i.e., the ρ_i, to obtain *globally* consistent information, i.e., the s_i, is a difficult task, reminiscent of the quantum marginal problem.

To overcome this requires a series of steps. First, we start with the 1D CONSISTENCY problem on qudits, so that it suffices to stitch together *nearest neighbor* reduced states $(\rho_{i,i+1}, \rho_{i+1,i+2})$ on the line. To this end, we first show[7] QMA-completeness of 1D CONSISTENCY with local dimension $d = 8$ via many-one reduction by combining the locally simulatable technique of Broadbent and Grilo [8] with the QMA-complete result for the 1D Local Hamiltonian problem with $d = 8$ of Hallgren, Nagaj, and Narayanaswami [19]. Then, we take the local nearest-neighbor reduced states ρ_i on qu-8-its from 1D CONSISTENCY, decompose each qudit q_i into a triple of qubits T_i, and consider all possible 6-*local* HKP shadows on pairs (T_i, T_{i+1}). Crucially, we know under the HKP protocol that any valid local shadow's expectation should exactly recover the corresponding state ρ_i. Using this fact and our 1D setup, we derive a *linear program* (LP) which captures "how much weight/probability" to put onto each local shadow, so that the "local probabilities" are consistent with some global HKP shadow if and only if the 1D CONSISTENCY instance we started with is a YES instance.

Unfortunately, solving this LP is itself not enough, because we next need to simulate the probability of a local shadow s_j occurring when measuring local Cliffords in HKP by *repeating* s_j an appropriate integer number of times in our shadow set S. We must, in fact, do this *exactly* to ensure consistency, and so we next "round" our LP into an *integer program* (IP) to give us *integer* weights on local shadows. This raises the potential roadblock that solving integer programs is NP-hard, but here we again crucially use the fact that we are working in 1D. Specifically, we exploit the 1D structure to design an efficient dynamic program to solve the IP, obtaining the desired integer weights on local shadows. Finally, we construct a list of *global* shadows by repeatedly carefully stitching together strings of local shadows under appropriate permutations given by a perfect matching.

---

[6] QMA(2) is QMA, but where the proof is promised to be in tensor product [26].

[7] One could alternatively use the 1D CONSISTENCY QMA-hardness result of Liu [28], but this would only yield hardness under Turing reductions, not many-one reductions.

qc-Σ_2*-completeness of CSV with exponentially many observables.* To connect CSV with qc-Σ_2, we first go through the QMA^+ formalism of Aharonov and Regev [4]. Roughly, in the latter one is given a set of polynomially many measurements Π_i and targets $r_i \in \mathbb{R}$, and asked if there is a state ρ such that $\mathrm{Tr}(\Pi_i \rho) \approx r_i$. While Aharonov and Regev showed $\mathsf{QMA}^+ = \mathsf{QMA}$, here we define the analogous class with *exponentially* many Π_i, denoted $\mathsf{SuperQMA}_{\mathrm{exp}}$. We then prove CSV with exponentially many observables is $\mathsf{SuperQMA}_{\mathrm{exp}}$-complete, and subsequently show that $\mathsf{SuperQMA}_{\mathrm{exp}} = \text{qc-}\Sigma_2$ to complete the proof. Intuitively, the latter holds because the existential quantum proof provides the globally consistent state, and the universal classical proof allows the verifier to iterate through all exponentially many measurement checks.

Open questions. We have initiated the study of the complexity of verifying classical shadows. For hardness, an important open question is whether other specific classical shadow protocols and observable sets have QMA-hard CSV problems? In the case of exponentially many observables, for example, can one give a qc-Σ_2-completeness proof of CSV for the protocol of King, Gosset, Kothari and Babbush [24]? The main bottleneck we faced here was that, unlike in our proof for HKP with polynomially many observables (Theorem 1.3), it is not clear how to start from an "exponential size" analogue of the QMA-complete 1D CONSISTENCY problem. A natural idea might be to start with *translationally invariant* 1D systems [15]. Such systems, however, act on *exponentially* many qudits, whereas our setting requires *polynomially* many qubits – the exponentiality occurs only in the number of observables for CSV. Finally, are there instances of CSV aside from our HKP, MYZ global Clifford results which can also be dequantized, or even better, solved classically without sampling assumptions?

Organization. Section 2 begins with preliminaries, including reviews of the HKP and MYZ classical shadow protocols. Section 3 studies the general CSV problem (i.e., not restricted to any particular shadow protocol), including the case of exponentially many observables. Finally, Section 4 studies the HKP and MYZ protocols, showing QMA-hardness and our dequantization result.

2 Preliminaries

Definitions. We use $A \leq B$ and $A \leq_r B$ to denote poly-time deterministic many-one and poly-time randomized reductions from A to B, respectively.

▶ **Definition 2.1** (Spatial sparsity [31]). *A spatially sparse hypergraph G on n vertices has:*

1. *every vertex participates in $O(1)$ hyper-edges, and*
2. *there is a straight-line drawing in the plane such that every hyper-edge overlaps with $O(1)$ other hyper-edges and the surface covered by every hyper-edge is $O(1)$.*

▶ **Definition 2.2** (QMA^+[4]). *A language $L \in \mathrm{QMA}^+$ if there exists a super-verifier*[8] *and polynomials p_1, p_2, p_3 such that:*

- $\forall x \in L \;\; \exists \rho \;\; \Pr_{V,r,s}\left(\left|\mathrm{Tr}\left(\Pi^{|1\rangle} V \rho V^\dagger\right) - r\right| \leq s\right) = 1$
- $\forall x \notin L \;\; \forall \rho \;\; \Pr_{V,r,s}\left(\left|\mathrm{Tr}\left(\Pi^{|1\rangle} V \rho V^\dagger\right) - r\right| \leq s + p_3(1/|x|)\right) \leq 1 - p_2(1/|x|)$

where probabilities are taken over the outputs V, r, s of the super-verifier and ρ is a density matrix over $p_1(|x|)$ qubits.

[8] A "super-verifier" is given by a classical polynomial-time randomized algorithm that given an input x outputs a description of a quantum circuit V and two numbers $r, s \in [0, 1]$.

▶ **Definition 2.3.** ($\mathsf{SuperQMA}(m,\epsilon)$). *A promise problem A is in* $\mathsf{SuperQMA}(m,\epsilon)$ *if there exists a super-verifier $V = \{(V_{x,i}, r_{x,i}, s_{x,i})\}_{i\in[m]}$ such that:*

- $\forall x \in A_{yes}, \exists \rho$: $\Pr_i\left(\left|\mathrm{Tr}(\Pi^{(1)} V_{x,i}\rho V_{x,i}^\dagger) - r_{x,i}\right| \le s_{x,i}\right) = 1.$
- $\forall x \in A_{no}, \forall \rho$: $\Pr_i\left(\left|\mathrm{Tr}(\Pi^{(1)} V_{x,i}\rho V_{x,i}^\dagger) - r_{x,i}\right| \le s_{x,i} + \epsilon\right) \le 1 - 1/m,$

where probabilities are taken over $i \in [m]$, ρ is a density matrix on $p(|x|)$ qubits, $r_{x,i}, s_{x,i} \in [0,1]$ and $1/\mathrm{poly}(n) \le \epsilon \le 1$. We additionally assume there exists a classical algorithm which, given any $i \in [m]$, efficiently computes $(V_{x,i}, r_{x,i}, s_{x,i})$ in time polynomial in the number of qubits.

Note our definition allows *exponentially* many checks, so long as each check can be efficiently generated on demand. We further remark that our definition $\mathsf{SuperQMA}_{\mathrm{poly}}$, i.e., with $m = \mathrm{poly}(|x|)$, coincide with QMA^+ from Ref. [4]. However, to the best of our knowledge, our $\mathsf{SuperQMA}_{\mathrm{exp}}$ with $m = \exp(|x|)$ has not been considered elsewhere before.

▶ **Definition 2.4.** ($\mathsf{SuperQMA}(2)(m,\epsilon)$). *We define it exactly as Definition 2.3 but with the promise that $\rho = \rho_A \otimes \rho_B$ where ρ_A and ρ_B are density matrices on $p_1(|x|)$ and $p_2(|x|)$ qubits, respectively.*

▶ **Definition 2.5** ($\mathsf{Q\Sigma_i}$ [14]). *A promise problem $A = (A_{\mathrm{yes}}, A_{\mathrm{no}})$ is in $\mathsf{Q\Sigma_i}(c,s)$ if there is a polynomial-time uniform quantum verifier $\{V_n\}$ and a polynomial p such that, on input $x \in \{0,1\}^n$, V_n receives i unentangled $p(n)$-qubit quantum proofs $\rho_1 \otimes \cdots \otimes \rho_i$ and satisfies:*

- ***Completeness:*** *If $x \in A_{\mathrm{yes}}$, then $\exists\rho_1\, \forall\rho_2\, \cdots\, Q_i\rho_i$ s.t. V_n accepts with probability $\ge c$.*
- ***Soundness:*** *If $x \in A_{\mathrm{no}}$, then $\forall\rho_1\, \exists\rho_2\, \cdots\, \overline{Q}_i\rho_i$ s.t. V_n accepts with probability $\le s$.*

Here $Q_i = \exists$ if i is odd and $Q_i = \forall$ if i is even. Finally, $\mathsf{Q\Sigma_i} = \bigcup_{c-s\in\Omega(1/\mathrm{poly}(n))} \mathsf{Q\Sigma_i}(c,s)$.

▶ **Definition 2.6** (Quantum polynomial hierarchy (QPH) [14]). $\mathsf{QPH} = \bigcup_{i\in\mathbb{N}} \mathsf{Q\Sigma_i}$.

▶ **Definition 2.7** ($\mathsf{qc\text{-}\Sigma_2}(c,s)$). *A promise problem $A = (A_{\mathrm{yes}}, A_{\mathrm{no}})$ is in $\mathsf{qc\text{-}\Sigma_2}(c,s)$ if there is a polynomial-time generated quantum verifier V_x which, on input $x \in \{0,1\}^n$, receives a polynomial-size quantum proof ρ and a polynomial-size classical proof y, and satisfies:*

- ***Completeness****: If $x \in A_{yes}$, then $\exists\ \rho$ such that $\forall\, |y\rangle$, $\Pr[V_x(\rho, |y\rangle) = 1] \ge c$.*
- ***Soundness****: If $x \in A_{no}$, then $\forall\rho$, $\exists\, |y\rangle$ such that $\Pr[V_x(\rho, |y\rangle) = 1] \le s$.*

▶ **Definition 2.8** ($\mathsf{qc\text{-}\Sigma_2}$). *We define $\mathsf{qc\text{-}\Sigma_2} := \mathsf{qc\text{-}\Sigma_2}(2/3, 1/3)$. In Lemma 3.12, we show that this constant-gap definition is equivalent to the inverse-polynomial-gap definition $\bigcup_{c-s\ge 1/\mathrm{poly}(n)} \mathsf{qc\text{-}\Sigma_2}(c,s)$.*

▶ **Definition 2.9** ($\mathsf{qc\text{-}\Sigma_2}(2)$). $\mathsf{qc\text{-}\Sigma_2}(2) := \bigcup_{c-s\ge 1/\mathrm{poly}(n)} \mathsf{qc\text{-}\Sigma_2}(2)(c,s)$, *where $\mathsf{qc\text{-}\Sigma_2}(2)(c,s)$ is defined as Definition 2.7 with the promise that ρ is a product state $\rho_A \otimes \rho_B$.*

Observables. In quantum mechanics, an observable is represented by a Hermitian operator O. Its real eigenvalues correspond to the possible outcomes of a measurement. Throughout this work, we assume without loss of generality that all observables O_i are normalized such that their operator norm $\|O_i\| \le 1$.

Succinct access assumption. When we say that we assume succinct access to a set $\{A_i^{(n)}\}_{i=1}^m$, with n the natural size parameter of the instance (e.g., number of qubits for observables or precision parameter for a real value), we mean that given an index i, a $\mathrm{poly}(n)$-bit description of A_i can be produced in $\mathrm{poly}(n)$-time.

Huang-Kueng-Preskill classical shadow framework. Here we briefly describe a classical shadow protocol proposed by Huang, Kueng and Preskill in Ref. [21]. For an unknown n-qubit state ρ fix an ensemble $\mathcal{U}$ of unitaries on n qubits. In each round do the following: sample $U \sim \mathcal{U}$, measure $U\rho U^\dagger$ in the computational basis to get a bitstring $b_\ell \in \{0,1\}^n$, and store a succinct classical description of $U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell$. The average channel

$$\mathcal{M}(\rho) = \mathbb{E}_{U\sim\mathcal{U}} \sum_b \langle b|U\rho U^\dagger|b\rangle \; U^\dagger |b\rangle \langle b| U$$

is invertible for tomographically complete $\mathcal{U}$, so a single-shot *snapshot* is

$$\hat{\rho}_\ell = \mathcal{M}^{-1}\big(U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell\big), \qquad \mathbb{E}[\hat{\rho}_\ell] = \rho.$$

For any observable O we use $\hat{o}_\ell(O) := \mathrm{Tr}[O\,\hat{\rho}_\ell]$. Partition the L rounds into K blocks $B_1, \ldots, B_K$ of (nearly) equal size, set

$$\overline{o}_k = \frac{1}{|B_k|} \sum_{\ell \in B_k} \hat{o}_\ell(O), \quad A(S,O) = \mathrm{median}\{\overline{o}_1, \ldots, \overline{o}_K\} \text{ (rounded to } \chi \text{ bits)}.$$

Since $\mathbb{E}[U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell] = \mathcal{M}(\rho)$, linearity gives $\mathbb{E}[\hat{\rho}_\ell] = \rho$ and thus $\mathbb{E}[\hat{o}_\ell(O)] = \mathrm{Tr}(O\rho)$ (unbiased). The median of means provides robustness. We need

$$L = O\left(\frac{\log(M)}{\epsilon^2} \max_{1\le i\le M} \left\| O_i - \frac{\mathrm{Tr}(O_i)}{2^n} I \right\|_{\text{shadow}}^2\right)$$

samples to estimate M observables up to error ϵ.

Local-Clifford (random Pauli). Here $\mathcal{U} = \mathrm{Cl}(2)^{\otimes n}$. Per round, sample independent single-qubit Cliffords $U_{j,\ell}$ (equivalently, pick Pauli bases $P_{j,\ell} \in \{X, Y, Z\}$) and measure to get bits $x_{j,\ell}$. Store the *measurement record*

$$s_\ell = \big((P_{1,\ell}, b_{1,\ell}), \ldots, (P_{n,\ell}, b_{n,\ell})\big), \quad b_{j,\ell} = (-1)^{x_{j,\ell}} \in \{\pm 1\}.$$

The average channel factorizes as $\mathcal{M} = D_{1/3}^{\otimes n}$ with inverse $D_{1/3}^{-1}(A) = 3A - \mathrm{Tr}(A)\, I_2$, hence the snapshot factorizes sitewise as

$$\hat{\rho}_\ell = \bigotimes_{j=1}^n \hat{\eta}_{j,\ell}, \quad \hat{\eta}_{j,\ell} = 3\, |\psi_{j,\ell}\rangle \langle \psi_{j,\ell}| - I_2,$$

where $|\psi_{j,\ell}\rangle := U_{j,\ell}^\dagger |x_{j,\ell}\rangle$ an eigenvector of X or Y or Z. To estimate M k-local observables up to error ϵ, it suffices to take $L = \mathcal{O}\big(\frac{\log M}{\epsilon^2} \max_i 4^k \|O_i\|_\infty^2\big)$ rounds.

Global Clifford. Here $\mathcal{U} = \mathrm{Cl}(2^n)$. Per round, sample $U_\ell \in \mathrm{Cl}(2^n)$ uniformly at random and measure to get b_ℓ. Store the measurement record as $s_\ell = (\mathsf{Stab}_\ell, b_\ell)$ where Stab is the efficient classical representation of the global Clifford via the stabilizer formalism and $b \in \{0,1\}^n$ the measurement outcome of that round. The average channel is the global depolarizing map

$$\mathcal{M}(\rho) = D_{1/(2^n+1)}(\rho), \qquad \mathcal{M}^{-1}(A) = (2^n+1)A - \mathrm{Tr}(A) I_{2^n},$$

so the snapshot is $\hat{\rho}_\ell = (2^n+1)\,|\psi_\ell\rangle \langle \psi_\ell| - I_{2^n}$. To estimate M linear observables, one needs $L = \mathcal{O}\big(\frac{\log M}{\epsilon^2} \max_i \mathrm{Tr}(O_i^2)\big)$ rounds.

Mao–Yi–Zhu classical shadow framework. Mao, Yi and Zhu [29] extend the HKP classical-shadow protocol to n qudits of odd prime local dimension d. Let $\mathbb{F}_d$ be the finite field with d elements and $\omega = e^{2\pi i/d}$ a primitive d-th root of unity. Fix an ensemble $\mathcal{E}$ of unitaries on n qudits. In each round: sample $U \sim \mathcal{E}$, measure $U\rho U^\dagger$ in the computational basis to get an outcome $b_\ell \in \mathbb{F}_d^n$, and store a succinct classical description of $U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell$. The average channel

$$\mathcal{M}(\rho) = \mathbb{E}_{U\sim\mathcal{E}} \sum_b \langle b|U\rho U^\dagger|b\rangle \; U^\dagger |b\rangle \langle b| U$$

is invertible for tomographically complete $\mathcal{E}$, so a single-shot *snapshot* is

$$\hat{\rho}_\ell = \mathcal{M}^{-1}\big(U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell\big), \qquad \mathbb{E}[\hat{\rho}_\ell] = \rho.$$

For any observable O we use $\hat{o}_\ell(O) := \mathrm{Tr}[O\,\hat{\rho}_\ell]$. Partition the L rounds into K blocks and take the median of block-means (rounded to χ bits) as before. Since $\mathbb{E}[U_\ell^\dagger |b_\ell\rangle \langle b_\ell| U_\ell] = \mathcal{M}(\rho)$, linearity gives $\mathbb{E}[\hat{o}_\ell(O)] = \mathrm{Tr}(O\rho)$ (unbiased). Again,

$$L = O\left(\frac{\log(M)}{\epsilon^2} \max_{1\le i\le M} \left\| O_i - \frac{\mathrm{Tr}(O_i)}{d^n} I \right\|_{\text{shadow}}^2\right)$$

samples suffice to estimate M observables up to error ϵ.

Local-Clifford. Here $\mathcal{E} = \mathrm{Cl}(d)^{\otimes n}$. Per round, sample independent single-qudit Cliffords $U_{j,\ell}$ (equivalently, pick on each site one of the $d+1$ stabilizer bases and measure there). Store the *measurement record* $s_\ell = \big((\mu_{1,\ell}, b_{1,\ell}), \ldots, (\mu_{n,\ell}, b_{n,\ell})\big)$, where $\mu_{j,\ell} \in \mathbb{F}_d \cup \{\infty\}$ labels the basis ($\mu = \infty$: Z-eigenbasis; $\mu = t \in \mathbb{F}_d$: eigenbasis of $Z^t X$) and $b_{j,\ell} \in \mathbb{F}_d$ is the outcome label. The average channel factorizes as $\mathcal{M} = D_{1/(d+1)}^{\otimes n}$ with inverse $D_{1/(d+1)}^{-1}(A) = (d+1)A - \mathrm{Tr}(A)\, I_d$, hence the snapshot factorizes sitewise:

$$\hat{\rho}_\ell = \bigotimes_{j=1}^{n} \hat{\eta}_{j,\ell}, \qquad \hat{\eta}_{j,\ell} = (d+1)\; |\phi_{\mu_{j,\ell},b_{j,\ell}}\rangle \langle \phi_{\mu_{j,\ell},b_{j,\ell}}| - I_d,$$

where $|\phi_{\mu,a}\rangle$ is the eigenvector in the chosen stabilizer basis. To estimate M k-local observables up to error ϵ, it suffices to take $L = \mathcal{O}\Big(\frac{\log M}{\epsilon^2} \max_i d^{2k} \, \|O_i\|_\infty^2\Big)$.

Global Clifford. Here $\mathcal{E} = \mathrm{Cl}(d^n)$. Per round, sample $U_\ell \in \mathrm{Cl}(d^n)$ uniformly at random and measure to get $b_\ell \in \mathbb{F}_d^n$. Store the measurement record: $s_\ell = (\mathsf{Stab}_\ell, b_\ell)$, where Stab is the efficient classical representation (via stabilizer formalism) of the global Clifford and b the d-ary outcome string. The average channel is the global depolarizing map

$$\mathcal{M}(\rho) = D_{1/(d^n+1)}(\rho), \qquad \mathcal{M}^{-1}(A) = (d^n+1)A - \mathrm{Tr}(A) I_{d^n},$$

so the snapshot is $\hat{\rho}_\ell = (d^n+1)\,|\psi_\ell\rangle \langle \psi_\ell| - I_{d^n}$. To estimate M linear observables, one needs $L = \mathcal{O}\Big(\frac{\log M}{\epsilon^2} \max_i [(2d-3)\,\mathrm{Tr}(O_i^2) + 2\|O_i\|_\infty^2]\Big)$ rounds.

3 Complexity of ObsCon

Definition 1.2 is overloaded for the general hardness results we are about to present. For that reason we will now recast it in a more abstract form. Notice that we will come back to the full-fledged definition when we consider specific classical shadow protocols.

▶ **Definition 3.1** (Observable consistency (ObsCon)). *The input is a set of observables, as in Definition 1.1, along with their target expectation values* $(O_i, y_i)_{i=1}^m$*, for which we also assume succinct access, and parameters* α *and* β *satisfying* $\beta - \alpha \geq 1/\mathrm{poly}(n)$*. We further assume* $y_i \in [-1, 1]$ *and that* $0 \leq \alpha < \beta \leq 2$*. The output is to decide between the following cases:*

- ***Yes****:* $\exists$ *n-qubit state* ρ *such that* $\forall\ i \in [m]$, $|\mathrm{Tr}\,(O_i\rho) - y_i| \leq \alpha$.
- ***No****:* $\forall$ *n-qubit states* ρ, $\exists\ i \in [m]$ *such that* $|\mathrm{Tr}\,(O_i\rho) - y_i| \geq \beta$.

▶ **Lemma 3.2.** CSV *and* ObsCon *are equivalent under polynomial-time many-one reductions.*

Proof. Both directions are straightforward

- CSV $\leq$ ObsCon. Keep the same observables and define $y_i := A(S, i)$.
- ObsCon $\leq$ CSV. Keep the same observables, use a dummy shadow S and a recovery algorithm $A(S, i)$ that ignores S and outputs y_i. ◀

We will analyze the complexity of this problem in two regimes, distinguished by the number of observables m.

3.1 Polynomially many observables ($\mathsf{ObsCon}_{\mathsf{poly}}$)

▶ **Definition 3.3** ($\mathsf{ObsCon}_{\mathsf{poly}}$). *Same as Definition 3.1 with* $m = \mathrm{poly}(n)$.

▶ **Proposition 3.4.** $\mathsf{ObsCon}_{\mathrm{poly}} \in \mathsf{SuperQMA}_{\mathrm{poly}}$.

Proof.

Verification procedure. Given the state ρ the verifier picks $i \in [m]$ uniformly at random and measures the observable O_i on the state ρ. This will give one of its eigenvalues λ_j. Then define a biased coin that gives heads with probability $p_h = \frac{1+\lambda_j}{2}$ and tails with probability $p_t = 1 - p_h$. Flip the coin and accept on heads, reject on tails. The overall acceptance probability becomes $\Pr(\text{accept}|i) = \frac{1}{2}(1 + \mathrm{Tr}(\rho O_i))$. Set the target probability to be $r_{x,i} = \frac{1}{2}(1 + y_i)$ and the tolerance parameter $s_{x,i} = \frac{\alpha}{2}$, uniform $\forall i$. Then our protocol works with $\epsilon = \frac{\beta-\alpha}{4}$.

Completeness. From the promise of the YES case we have that $\forall i\ |\mathrm{Tr}(O_i\rho) - y_i| \leq \alpha$. So we find

$$\forall i\ \ |\Pr(\text{accept}|i) - r_{x,i}| = \frac{1}{2}|\mathrm{Tr}(O_i\rho) - y_i| \leq \frac{\alpha}{2} = s_{x,i}.$$

In other words $\Pr_i(|\Pr(\text{accept}|i) - r_{x,i}| \leq s_{x,i}) = 1$.

Soundness. From the promise of the NO case we have that there exists at least one i, say i^*, s.t. $|\mathrm{Tr}(O_{i^*}\rho) - y_{i^*}| \geq \beta$. For i^* we then have

$$|\Pr(\text{accept}|i^*) - r_{x,i^*}| = \frac{1}{2}|\mathrm{Tr}(O_{i^*}\rho) - y_{i^*}| \geq \frac{\beta}{2} > s_{x,i^*} + \epsilon$$

Where the last inequality holds since $s_{x,i^*} + \epsilon = \frac{\alpha+\beta}{4}$ and $\beta - \alpha \geq 1/\mathrm{poly}(n)$.

In other words $\Pr_i[|\Pr(\text{accept}|i) - r_{x,i}| \leq s_{x,i} + \epsilon] \leq 1 - \frac{1}{m}$. ◀

▶ **Proposition 3.5.** $\mathsf{ObsCon}_{\mathrm{poly}}$ *is* $\mathsf{SuperQMA}_{\mathrm{poly}}$*- hard.*

Proof. For input x the $\mathsf{SuperQMA}_{\mathrm{poly}}$ super-verifier provides $m = \mathrm{poly}(|x|)$ checks $\{(V_i, r_i, s_i)\}_{i=1}^m$, with $r_i, s_i \in [0,1]$ and a global gap parameter $1/\mathrm{poly}(n) \le \epsilon \le 1$. Define

$$\epsilon' := \frac{\epsilon}{2}, \quad \tau := \frac{\epsilon}{4}, \quad s_i' := \max\{s_i, \tau\}, \quad t_i = \frac{\tau}{s_i'}.$$

The reduction outputs the $\mathsf{ObsCon}_{\mathrm{poly}}$ instance $\{(O_i, y_i)\}_{i=1}^m$ with uniform thresholds α, β defined by

$$O_i := t_i\big(V_i^\dagger \Pi^{(1)} V_i\big), \quad y_i := t_i r_i, \quad \alpha := \tau, \quad \beta := \tau + \tau\epsilon'.$$

Notice that this choice of parameters gives us a $\beta - \alpha = \tau\epsilon' \ge \frac{\epsilon^2}{8} \ge \frac{1}{\mathrm{poly}(n)}$.

Completeness (YES case). If the original $\mathsf{SuperQMA}_{\mathrm{poly}}$ instance is YES, there exists a witness ρ such that

$$\forall i \quad : \big|\operatorname{Tr}(V_i^\dagger \Pi^{(1)} V_i\, \rho) - r_i\big| \le s_i \le s_i'.$$

Multiplying by t_i gives

$$\big|\operatorname{Tr}(O_i\rho) - y_i\big| = t_i\big|\operatorname{Tr}(V_i^\dagger \Pi^{(1)} V_i\, \rho) - r_i\big| \le t_i s_i' = \tau = \alpha.$$

Thus the same ρ satisfies $\big|\operatorname{Tr}(O_i\rho) - y_i\big| \le \alpha$ for all i, so the mapped instance is a YES-instance of $\mathsf{ObsCon}_{\mathrm{poly}}$.

Soundness (NO case). If the original $\mathsf{SuperQMA}_{poly}$ instance is NO, then for every state ρ there exists some index i^* with

$$\big|\operatorname{Tr}(V_{i^*}^\dagger \Pi^{(1)} V_{i^*}\, \rho) - r_{i^*}\big| > s_{i^*} + \epsilon \ge s_{i^*}' - \tau + \epsilon = s_{i^*}' + \frac{3\epsilon}{4} > s_{i^*}' + \epsilon'.$$

Multiplying by t_{i^*} yields

$$\big|\operatorname{Tr}(O_{i^*}\rho) - y_{i^*}\big| > t_{i^*} s_{i^*}' + t_{i^*}\epsilon' = \tau + \frac{\tau}{s_{i^*}'}\epsilon' \ge \tau + \tau\epsilon' = \beta.$$

Thus the mapped instance violates the uniform β-threshold for the index i^*, matching the $\mathsf{ObsCon}_{\mathrm{poly}}$ NO condition. ◀

▶ **Theorem 3.6** ([4])**.** $\mathsf{QMA} = \mathsf{SuperQMA}_{\mathrm{poly}}$.

▶ **Corollary 3.7.** $\mathsf{ObsCon}_{\mathrm{poly}}$ *is* QMA*-complete.*

Proof. This follows from Propositions 3.4 and 3.5 and Theorem 3.6. ◀

Note here that the QMA-hardness result need not go through the super-verifier machinery. We can directly reduce from CLDM problem which is known to be QMA complete under Karp reductions [8]. You can find this reduction in the full version of the paper. The reason we use this machinery is because it will become helpful in the exp regime that we analyze next.

3.2 Exponentially many observables ($\mathsf{ObsCon}_{\mathsf{exp}}$)

We now move to analyze the case where the observables can be exponentially many, albeit we have succinct access to them. Here the super-verifier machinery we developed for the poly regime will help us extract completeness results for $\mathsf{ObsCon}_{\mathrm{exp}}$ immediately.

▶ **Definition 3.8** ($\mathsf{ObsCon}_{\mathsf{exp}}$). *Same as Definition 3.1 with* $m = \exp(n)$.

▶ **Proposition 3.9.** $\mathsf{ObsCon}_{\mathrm{exp}} \in \mathsf{SuperQMA}_{\mathrm{exp}}$.

Proof. The proof follows in the same manner as in the poly-case, Proposition 3.4. The verifier only needs to generate and execute a single, randomly chosen check (O_i, y_i). Since the $\mathsf{ObsCon}_{\mathrm{exp}}$ instance guarantees that any such pair can be generated in polynomial time given the index i, the verifier remains efficient. The soundness guarantee of $1/m$ holds, where m is now exponential in the number of qubits. ◀

▶ **Proposition 3.10.** $\mathsf{ObsCon}_{\mathrm{exp}}$ *is* $\mathsf{SuperQMA}_{\mathrm{exp}}$*-hard.*

Proof. The proof again carries over from the poly case. Here for each one of the exponentially many checks of the $\mathsf{SuperQMA}_{\mathrm{exp}}$, the mapping in Proposition 3.5 gives, in polynomial time, one of the exp many pairs (O_i, y_i) of $\mathsf{ObsCon}_{\mathrm{exp}}$ along with the global parameters α, β. This is all we need since we assume succinct access to both the checks and the pairs. ◀

▶ **Corollary 3.11.** $\mathsf{ObsCon}_{\mathrm{exp}}$ *is* $\mathsf{SuperQMA}_{\mathrm{exp}}$*-complete.*

Proof. Follows from Propositions 3.9 and 3.10. ◀

We next show that $\mathsf{SuperQMA}_{\mathrm{exp}}$ coincides with the second level of a quantum-classical variant of the quantum polynomial hierarchy (QPH). By QPH we refer to the hierarchy $\mathsf{Q\Sigma_i}$ of Ref. [14] (see Definition 2.5). We call our variant $\mathsf{qc\text{-}\Sigma_2}$ (see Definition 2.8).

We first prove an amplification lemma showing that the constant-gap and inverse-polynomial-gap definitions of $\mathsf{qc\text{-}\Sigma_2}$ coincide.

▶ **Lemma 3.12** (Amplification for $\mathsf{qc\text{-}\Sigma_2}(c,s)$)**.** *Let* $L \in \mathsf{qc\text{-}\Sigma_2}(c,s)$*, where* $\Delta := c - s \geq 1/\mathrm{poly}(n)$*. Then for any polynomial* r*,* $L \in \mathsf{qc\text{-}\Sigma_2}(1 - 2^{-r(n)}, 2^{-r(n)})$.

Proof sketch. The idea here is to use Sion's minimax theorem in order to show that in the NO case there exists a "bad" distribution over the classical challenges, meaning $\lambda_{\max}(M_\mu) \leq s$, where $M_\mu = \mathbb{E}_{z\sim\mu} M_z$ with M_z the acceptance POVM of our verifier with z hardwired. Next we show that this distribution μ sparsifies, meaning there exists a polynomial list of challenges, $\Lambda = (z_1, \ldots, z_T)$, with $T = \mathrm{poly}(n)$, drawn from the distribution μ such that $\lambda_{\max}(M_\Lambda) \leq s + \eta =: s'$ with $M_\Lambda = \frac{1}{T}\sum_{j=1}^{T} M_{z_j}$ and $\eta := \frac{\Delta}{4}$. Now the amplified $\mathsf{qc\text{-}\Sigma_2}$ verifier expects an existential quantum proof of $R = \mathrm{poly}(n)$ blocks, each of the original proof size. On each block, the verifier chooses $j \in [T]$ uniformly at random, runs V_x with challenge z_j, and records the output bit. Accept iff the average acceptance rate is at least $\theta := \frac{c+s'}{2}$. Completeness follows because the YES witness is accepted with probability at least c for every challenge z, hence also for every list Λ. For soundness, the list above defines a fixed QMA verifier with soundness at most s'. By the standard parallel-repetition/amplification argument (see [25]), repetition remains sound against proofs entangled across the R registers, and choosing appropriate $R = \mathrm{poly}(n)$ gives the desired amplified soundness. ◀

▶ **Lemma 3.13.** $\mathsf{qc\text{-}\Sigma_2} \subseteq \mathsf{SuperQMA}_{\mathrm{exp}}$.

Proof. Let $L \in \mathsf{qc\text{-}\Sigma_2}$. By Lemma 3.12, we may assume the $\mathsf{qc\text{-}\Sigma_2}$ verifier has completeness $2/3$ and soundness $1/3$. Start by hardwiring the classical proof z into the verifier V of $\mathsf{qc\text{-}\Sigma_2}$, let us call it V_z. The super-verifier's checks are now parametrized by the classical strings $z \in \{0,1\}^{\ell(n)}$, i.e., $m = 2^{\ell(n)}$. Construct a super-verifier V' that on input x picks uniformly at random a challenge z and outputs the check $(V_z, r_z = 1, s_z = 1/3)$. This satisfies the definition of $\mathsf{SuperQMA}_{\mathrm{exp}}$ with $\epsilon = 1/6$.

Completeness. Let $x \in L$ then $\exists \rho$ such that $\forall c$ $\mathrm{Tr}(\Pi^{(1)} V_c \rho V_c^\dagger) \geq 2/3$. It is easy to see that the condition $|\mathrm{Tr}(\Pi^{(1)} V_c \rho V_c^\dagger) - 1| \leq 1/3$ is satisfied for all c.

Soundness. Let $x \notin L$ then $\forall \rho$ $\exists c$ s.t. $\mathrm{Tr}(\Pi^{(1)} V_c \rho V_c^\dagger) \leq 1/3$. This means that the condition $|\mathrm{Tr}(\Pi^{(1)} V_c \rho V_c^\dagger) - 1| > s + \epsilon = 1/3 + 1/6 = 1/2$ is satisfied for at least one c, for each ρ, so with probability $\geq \frac{1}{m}$. ◀

▶ **Lemma 3.14.** $\mathsf{SuperQMA}_{\mathrm{exp}} \subseteq \mathsf{qc}\text{-}\Sigma_2$.

Proof sketch. The $\forall$-prover names a check i that violates the super-verifier condition, and the $\mathsf{qc}\text{-}\Sigma_2$ verifier estimates the acceptance probability of V_i by running it on $k = O(n/\epsilon^2)$ proof registers and checking whether the empirical average lies within $s_i + \epsilon/2$ of r_i. Completeness follows by Hoeffding for honest k-copy witnesses, while soundness follows from the same Markov argument as in [4] , which applies even when the k registers are entangled. By Lemma 3.12, we can amplify to the standard constant-gap definition of $\mathsf{qc}\text{-}\Sigma_2$. ◀

▶ **Corollary 3.15.** $\mathsf{qc}\text{-}\Sigma_2 = \mathsf{SuperQMA}_{\mathrm{exp}}$.

Proof. Follows from Lemmas 3.13 and 3.14. ◀

An important variant of $\mathsf{ObsCon}_{\mathrm{exp}}$, motivated by the triply efficient classical shadow protocol for all the n-qubit Pauli observables [24], is the constant gap version.

▶ **Definition 3.16** ($\mathsf{ObsCon}^{\Theta(1)}_{\mathrm{exp}}$)**.** *Same as Definition 3.8 with* $\beta - \alpha = \Theta(1)$.

It is easy to see that even for a constant gap we still have $\mathsf{SuperQMA}_{\mathrm{exp}}$-completeness:

▶ **Lemma 3.17.** $\mathsf{ObsCon}^{\Theta(1)}_{\mathrm{exp}}$ *is* $\mathsf{SuperQMA}_{\mathrm{exp}}$*-complete.*

Proof sketch. Containment follows exactly as in Proposition 3.9. For hardness, use Corollary 3.15, $\mathsf{SuperQMA}_{\mathrm{exp}} = \mathsf{qc}\text{-}\Sigma_2$, and the amplification theorem for $\mathsf{qc}\text{-}\Sigma_2$,Lemma 3.12, so that any $L \in \mathsf{SuperQMA}_{\mathrm{exp}}$ has a $\mathsf{qc}\text{-}\Sigma_2$ verifier with completeness/soundness $2/3, 1/3$. For each classical challenge c, let M_c be the induced acceptance POVM of the amplified verifier with c hardwired, and output the $\mathsf{ObsCon}^{\Theta(1)}_{\mathrm{exp}}$ instance $\{O_c := M_c,\ y_c := 1,\ \alpha := 1/3,\ \beta := 2/3\}$ In the YES case, some ρ satisfies $\mathrm{Tr}(M_c \rho) \geq 2/3$ for all c, so $|\mathrm{Tr}(O_c \rho) - 1| \leq 1/3$. In the NO case, for every ρ some c satisfies $\mathrm{Tr}(M_c \rho) \leq 1/3$, so $|\mathrm{Tr}(O_c \rho) - 1| \geq 2/3$. Thus the constructed instance has constant gap. ◀

We conclude this section by showing an easy lower and upper bound for $\mathsf{qc}\text{-}\Sigma_2$.

▶ **Proposition 3.18.** $\mathsf{QMA} \subseteq \mathsf{qc}\text{-}\Sigma_2 \subseteq \mathsf{Q}\Sigma_2 \subseteq \mathsf{PSPACE}$.

Proof. The first inclusion follows since the $\mathsf{qc}\text{-}\Sigma_2$ verifier can simply ignore the proof from the $\forall$ prover and run the QMA verifier. The second follows since the verifier of $\mathsf{Q}\Sigma_2$ can measure the $\forall$ proof in the computational basis, essentially rendering the quantum proof to a classical one, or rather a distribution of classical ones, and then simulate the verifier of $\mathsf{qc}\text{-}\Sigma_2$. As for the third inclusion it is proven in Ref. [14]. The proof was based on the observation that $\mathsf{Q}\Sigma_2 = \mathsf{QRG}(1)$, where $\mathsf{QRG}(1)$ and its containment in PSPACE are presented in Ref. [22]. ◀

4 Complexity of specific protocol classical shadows

The QMA-completeness of $\mathsf{ObsCon}_{\mathrm{poly}}$, and so of the fully general version of $\mathsf{CSV}_{\mathrm{poly}}$ from Definition 1.2, demonstrates the problem's fundamental difficulty. We now further explore the complexity of this problem by casting it on specific, structured measurement protocols. We show that the hardness persists for two such protocols, namely the HKP with a local Clifford ensemble protocol, given in Ref. [21] and the MYZ which is its qudit generalization, given in Ref. [29]. Additionally we give an efficient algorithm result for the HKP, MYZ protocols with global Clifford ensemble.

▶ Remark 4.1. For a fixed shadow protocol P, the reduction $\mathsf{CSV}_P \leq \mathsf{ObsCon}_P$ is immediate by keeping the set of observables the same and setting $y_i := A(S, i)$; the converse direction is protocol dependent and need not be trivial.

4.1 HKP classical shadows

First we focus on the Huang, Kueng and Preskill protocol using local Clifford measurements [21] (for details on the protocol check Section 2). We will call the CSV problem that is based on this protocol $\mathsf{CSV}_{\mathrm{HKP}}$.

▶ **Definition 4.2** ($\mathsf{CSV}_{\mathrm{HKP}}$). *The definition is the same as Definition 1.2 only now our classical shadow has the structure dictated by the HKP local Clifford measurement protocol. That means the following:*

- *The shadow S consists of $L = \mathrm{poly}(n)$ strings, $\{s_i\}_{i=1}^{L}$, each of which encodes the Pauli-basis measurement and the measurement outcome of each round of the protocol. More formally, each string will be of the form $(P, b)_1, \ldots, (P, b)_n$ with $P \in \{X, Y, Z\}$ and $b \in \{-1, 1\}$, where $(P, b)_i$ denotes the basis and the measurement outcome of the i-th qubit.*
- *O is a set of $m = \mathrm{poly}(n)$ k-local observables on n-qubits, with $k = O(1)$.*
- *The recovery algorithm applies the inverse channel to extract the snapshot operators $\hat{\eta}$ from S and aggregates estimates via the* median of means *(MoM) technique.*

We sometimes speak of the snapshot operator associated to a stored string (P, b); it is not stored explicitly but computed in recovery. The map $s \to \hat{\eta}$ is a bijection onto the set of achievable snapshots, so storing strings or storing snapshots are equivalent representations, hence we freely use "strings" and "snapshots" interchangeably when no confusion can arise.

▶ **Definition 4.3** ($1D$-LH). *Given a local Hamiltonian on a chain of n qu-d-its $H = \sum_{i=1}^{n-1} H_{i,i+1}$ and thresholds α, β with $\beta - \alpha \geq 1/\mathrm{poly}(n)$, decide*

- *YES: $\lambda_{\min}(H) \leq \alpha$.*
- *NO: $\lambda_{\min}(H) \geq \beta$.*

▶ **Definition 4.4** ($1D$-CLDM). *Given local density matrices $\sigma_{i,i+1}$ for $i \in [n-1]$ for a system of n qu-d-its and parameters α, β with $\alpha \leq \mathsf{negl}(n)$, $\beta - \alpha \geq 1/\mathrm{poly}(n)$, decide*

- *YES: $\exists \rho \in \mathscr{D}(d^n)\ \forall i \in [n-1]\colon \|\mathrm{Tr}_{\overline{i,i+1}}(\rho) - \sigma_{i,i+1}\|_{\mathrm{Tr}} \leq \alpha$.*
- *NO: $\forall \rho \in \mathscr{D}(d^n)\ \exists i \in [n-1]\colon \|\mathrm{Tr}_{\overline{i,i+1}}(\rho) - \sigma_{i,i+1}\|_{\mathrm{Tr}} \geq \beta$.*

▶ **Theorem 4.5.** $1D$-CLDM *on 8-level qudits is* QMA*-complete.*

Proof sketch. The high level idea is to combine the results of Ref. [8], where they prove that the CLDM problem is QMA-complete under Karp reductions via the machinery of simulatable codes and of Ref. [19] where they show that $1D$-LH on a chain of 8-level qudits is still QMA-complete. For details, see the full version of the paper [23]. ◀

▶ **Theorem 4.6.** $1D\text{-}\mathsf{CLDM}_{d=2^\ell} \leq \mathsf{CSV}_{HKP}$.

Proof. Here we assume qudits of dimension $d = 2^\ell \in O(1)$, so that we can treat each qudit as ℓ qubits. We use the HKP shadow protocol with an ensemble of local Clifford operators, so that we end up applying random Pauli measurements, i.e., the product of ℓn single-qubit Paulis. Let $\sigma_{1,2}, \ldots, \sigma_{n-1,n} \in \mathscr{D}(d^2)$ be the input density matrices. We now describe the reduction from the $\sigma_{i,i+1}$ to a shadow. Let $\{\,\hat{\eta}_j\,\}_{j\in[m]}$ be the set of all $m = 6^\ell$ possible snapshots on ℓ qubits. For clarity, we are measuring each qubit in one of Pauli X, Y, or Z uniformly at random, therefore the set of possible snapshots on a single qubit are of form $3\,|\psi\rangle\langle\psi| - I$ for $|\psi\rangle$ an eigenvector of X, Y, or Z. In turn, a snapshot on a given qu*d*it is a tensor product of ℓ such terms.
Suppose that there exists a state $\rho \in \mathscr{D}(d^n)$ whose local marginals $\tau_{i,i+1} := \mathrm{Tr}_{\overline{i,i+1}}(\rho)$ satisfy $\|\tau_{i,i+1} - \sigma_{i,i+1}\|_{\mathrm{Tr}} \leq \alpha_{\mathsf{CLDM}}$. Since $\mathbb{E}[\hat{\rho}] = \rho$, we have

$$\tau_{i,i+1} = \mathrm{Tr}_{\overline{i,i+1}}(\rho) = \mathrm{Tr}_{\overline{i,i+1}}(\mathbb{E}[\hat{\rho}]) = \sum_j p_j \,\mathrm{Tr}_{\overline{i,i+1}}(\hat{\rho}^{(j)}) = \mathbb{E}[\hat{\rho}_{i,i+1}], \tag{3}$$

where $\{\,\hat{\rho}^{(j)}\,\}_j$ is a collection of all possible snapshots and p_j is the probability of obtaining that snapshot from the shadow protocol. Note that with our choice of shadow protocol, we have $\hat{\rho} = \hat{\rho}_1 \otimes \cdots \otimes \hat{\rho}_n$, where each $\hat{\rho}_i$ is a local snapshot on ℓ-qubits. Thus, by Equation (3), the ideal local snapshot distribution reconstructs $\tau_{i,i+1}$ exactly, and hence reconstructs the target $\sigma_{i,i+1}$ up to the original CLDM completeness error. Equivalently, for each edge there are probabilities $\{p_{i,j,k}\}$ such that $\sum_{j,k} p_{i,j,k}\hat{\eta}_j \otimes \hat{\eta}_k$ is within α_{CLDM} of $\sigma_{i,i+1}$ in trace norm, for i the left qudit index in a neighboring pair of qudits, and j and k indexing the possible snapshots on the left and right qudit, respectively.

The preceding discussion shows that, at the level of the ideal HKP snapshot distribution, a globally consistent state induces compatible local snapshot distributions on every neighboring pair. Since the reduction must output a finite classical shadow, we work directly with integer counts rather than real probabilities. Let $L = \mathrm{poly}(n)$ be chosen sufficiently large, and let $n_{i,j,k}$ denote the number of times the two-qudit snapshot $\hat{\eta}_j \otimes \hat{\eta}_k$ appears on edge $(i, i+1)$. We impose exact overlap-count constraints, which will allow the local shadows to be stitched into global strings, and an approximate reconstruction constraint with tolerance ϵ. The tolerance ϵ is chosen to absorb this CLDM error together with the finite-count approximation needed to represent the ideal snapshot distribution by L integer counts.

$$\left\| \sigma_{i,i+1} - \frac{1}{L} \sum_{j,k\in[m]} n_{i,j,k}\hat{\eta}_j \otimes \hat{\eta}_k \right\|_{\mathrm{Tr}} \leq \epsilon \qquad \forall i \in [n-1], \tag{4a}$$

$$\sum_{j\in[m]} n_{ijt} = \sum_{k\in[m]} n_{i+1,tk} \quad \forall i \in [n-2]\ \forall t \in [m], \tag{4b}$$

$$n_{i,j,k} \in \mathbb{Z}_{\geq 0} \qquad \forall i \in [n-1]\ \forall j \in [m]\ \forall k \in [m], \tag{4c}$$

$$\sum_{j,k\in[m]} n_{i,j,k} = L \qquad \forall i \in [n-1]. \tag{4d}$$

If the $1D\text{-}\mathsf{CLDM}$ instance is YES, then for sufficiently large $L = \mathrm{poly}(n)$ the system in Equation (4) is feasible. Indeed, drawing L HKP snapshots from a consistent state ρ gives edge counts satisfying the overlap constraints exactly, and by the HKP concentration bound the corresponding empirical edge averages are within the allowed tolerance. If Equation (4) is unsatisfiable then the reduction outputs a trivial NO-instance.

▶ **Proposition 4.7.** *There is a* dynamic programming algorithm *that efficiently solves the integer program defined in Equation* (4)*.*

Proof sketch. This is a classical constraint satisfiability problem on a path $x_1 - x_2 - \cdots - x_{n-1}$. For the DP algorithm and its runtime analysis see the full version of the paper [23]. ◀

We define "local shadows" $S_i = \{s_{il}\}_{l\in[L]}$ by taking $n_{i,j,k}$ copies of $\hat{\eta}_j \otimes \hat{\eta}_k$. We can now compute permutations $f_i \in \mathrm{S}_L$, such that $\mathrm{Tr}_1(s_{il}) = \mathrm{Tr}_2(s_{i+1,f_i(l)})$ via a perfect matching, which we are guaranteed it exists by construction (see Equation (4b)). Finally, we assemble the local shadows to a global shadow

$$S = \{s_l\}_{l\in[L]}, \quad s_l = s_{1,l} \otimes \mathrm{Tr}_A(s_{2,f_1(l)}) \otimes \mathrm{Tr}_A(s_{3,f_2(f_1(l))}) \otimes \cdots \otimes \mathrm{Tr}_A(s_{n-1,(f_{n-2}\circ\cdots\circ f_1)(l)}). \quad (5)$$

By construction, we have

$$\left\| \mathrm{Tr}_{\overline{i,i+1}} \left(\frac{1}{L} \sum_{l\in[L]} s_l \right) - \sigma_{i,i+1} \right\|_{\mathrm{Tr}} \leq \epsilon. \quad (6)$$

For the $\mathsf{CSV}_{\mathsf{HKP}}$ instance, we use K identical copies of the global shadow S (or to be precise the string equivalent of the snapshots), which will serve as the buckets for the MoM aggregation, essentially rendering this to an empirical average. As for the set of observables O, for each neighboring qudit pair $(i, i+1)$ we include all Pauli operators supported only on those 2ℓ qubits that comprise the pair. The recovery algorithm reconstructs the snapshots and aggregates estimations via MoM technique.

For notational convenience, define

$$\rho_{i,i+1} := \mathrm{Tr}_{\overline{i,i+1}}(\rho), \quad \widetilde{\sigma}_{i,i+1} := \mathrm{Tr}_{\overline{i,i+1}} \left(\frac{1}{L} \sum_{l\in[L]} s_l \right).$$

Completeness. If the $1D$-CLDM instance is a YES instance, there exists a state ρ such that

$$\|\rho_{i,i+1} - \sigma_{i,i+1}\|_{\mathrm{Tr}} \leq \alpha \qquad \forall i \in [n-1].$$

By construction of the integer counts and the stitched shadow, we have

$$\|\widetilde{\sigma}_{i,i+1} - \sigma_{i,i+1}\|_{\mathrm{Tr}} \leq \epsilon \qquad \forall i \in [n-1].$$

Therefore, by the triangle inequality, $\|\rho_{i,i+1} - \widetilde{\sigma}_{i,i+1}\|_{\mathrm{Tr}} \leq \alpha + \epsilon$. For every Pauli observable P on the qubits of the neighboring pair $(i, i+1)$, we have $\|P\|_\infty = 1$, and hence by Hölder's inequality,

$$\begin{aligned} |\mathrm{Tr}(P\rho) - A(S,P)| &= |\mathrm{Tr}\left(P\left(\rho_{i,i+1} - \widetilde{\sigma}_{i,i+1}\right)\right)| \\ &\leq \|\rho_{i,i+1} - \widetilde{\sigma}_{i,i+1}\|_{\mathrm{Tr}} \leq \alpha + \epsilon. \end{aligned}$$

Thus the constructed $\mathsf{CSV}_{\mathsf{HKP}}$ instance is a YES instance with $\alpha_{\mathrm{CSV}} := \alpha + \epsilon$.

Soundness. For soundness, suppose for contradiction that there exists a state ρ such that, for every Pauli observable P supported on a neighboring pair, $|\mathrm{Tr}(P\rho) - A(S,P)| \le \beta_{\mathrm{CSV}}$. Equivalently,

$$|\mathrm{Tr}\left(P\left(\rho_{i,i+1} - \widetilde{\sigma}_{i,i+1}\right)\right)| \le \beta_{\mathrm{CSV}} \qquad \forall i \in [n-1].$$

Since the Pauli observables on the 2ℓ qubits of a neighboring pair form a tomographically complete operator basis, and since the local dimension $d^2 = 2^{2\ell}$ is constant, there exists a constant $c_d > 0$ such that

$$\|\rho_{i,i+1} - \widetilde{\sigma}_{i,i+1}\|_{\mathrm{Tr}} \le c_d \beta_{\mathrm{CSV}} \qquad \forall i \in [n-1].$$

By construction of the integer counts and the stitched shadow,

$$\|\widetilde{\sigma}_{i,i+1} - \sigma_{i,i+1}\|_{\mathrm{Tr}} \le \epsilon \qquad \forall i \in [n-1].$$

Therefore, by the triangle inequality,

$$\|\rho_{i,i+1} - \sigma_{i,i+1}\|_{\mathrm{Tr}} \le c_d \beta_{\mathrm{CSV}} + \epsilon \qquad \forall i \in [n-1].$$

Choose β_{CSV} so that $c_d\beta_{\mathrm{CSV}} + \epsilon < \beta$, where β is the NO threshold of the input $1D$-CLDM instance. Then ρ would be a state whose every neighboring marginal is within distance strictly less than β of the corresponding $\sigma_{i,i+1}$, contradicting the NO case of the $1D$-CLDM instance. Hence the constructed $\mathsf{CSV}_{\mathsf{HKP}}$ instance is a NO instance. ◀

▶ **Corollary 4.8.** $\mathsf{QMA} \le \mathsf{CSV}_{HKP}$*, even for* 6*-local observables on a spatially sparse hypergraph.*

Proof. Follows from Theorems 4.5 and 4.6. Since here $d = 8$ and so $\ell = 3$, the observables are 6-local on qubits. It is easy to verify that the resulting hypergraph (where each qubit is a vertex, and each Pauli operator acting non-trivially on a set $V' \subseteq V$ of vertices is represented by a hyperedge) is spatially sparse, as per Definition 2.1. ◀

4.2 MYZ classical shadow

Our hardness result is not limited to qubit translated systems. A recent protocol by Mao, Yi, and Zhu [29] generalizes the local Clifford measurement framework to qudits of odd prime dimension d. Their protocol uses the ensemble $\mathcal{E} = \mathrm{Cl}(d)^{\otimes n}$, where $\mathrm{Cl}(d)$ is the single-qudit Clifford group, leading to snapshots that are tensor products of single-qudit operators. Each such operator is derived from one of the $d(d+1)$ single-qudit stabilizer states (for details see Section 2).

▶ **Definition 4.9** ($\mathsf{CSV}_{\mathsf{MYZ}}(d)$)**.** *The definition is the same as Definition 1.2, only now the classical shadow has the structure dictated by the MYZ local-Clifford protocol on odd-prime d. Concretely:*

- *Shadow S consists of $L = \mathrm{poly}(n)$ strings. Each string is of the form $((\mu,b)_1, \ldots, (\mu,b)_n)$ with $\mu \in \mathbb{F}_d \cup \{\infty\}$ the measurement basis label and $b \in [d]$ the measurement outcome.*
- *O is a set of k-local observables on n qudits, for fixed $k = O(1)$.*
- *The recovery algorithm applies the inverse channel of the measurement protocol on S to get the snapshots and then aggregates via MoM.*

This protocol works for odd prime d. Notice that we can always pad the local dimensions of our chain and add projector terms in our Hamiltonian and so we can trivially get a QMA-completeness under Karp reductions result for a $d \ge 8$-level $1D$-CLDM problem.

▶ **Theorem 4.10.** $1D\text{-CLDM}_{d=odd\ prime} \leq \mathsf{CSV}_{MYZ}(d)$.

Proof sketch. The proof is analogous with Theorem 4.6, only here the local dimension of the qudits is an odd prime. Let us first quickly summarize the differences of the two:

- **Single-site snapshot types ($\hat{\eta}$):** In MYZ local-Clifford ensemble, each site is measured in one of the $d+1$ stabilizer basis- the eigenbases of Z and Z^tX for $t \in \mathbb{F}_d$. For basis label $\mu \in \mathbb{F}_d \cup \{\infty\}$ and outcome label $b \in \mathbb{F}_d$, the single-site snapshot operator is

$$\hat{\eta}_{\mu,b} = (d+1)\,|\phi_{\mu,b}\rangle\,\langle\phi_{\mu,b}| - I_d$$

  where $|\phi_{\mu,b}\rangle$ is the eigenvector in the basis labelled by μ with outcome label b. The alphabet size is now $m' = d(d+1)$, so still constant for fixed d.
- **Observables:** Our observables will now be all the Hermitian real and imaginary parts of generalized Pauli/Weyl operators supported on adjacent qudits.

With these changes in mind we can see that our proof follows directly. Since the alphabet m' is still constant we can solve Equation (4) system efficiently, via the same DP algorithm. After that, we use the same "stitching the local shadows" argument to create a global shadow which alongside our observables and the known recovery algorithm will form the $\mathsf{CSV}_{\mathsf{MYZ}}$ instance. ◀

▶ **Corollary 4.11.** $\mathsf{QMA} \leq \mathsf{CSV}_{\mathsf{MYZ}}(d)$ *for every fixed odd prime local dimension* $d \geq 11$, *even for* 2*-local nearest-neighbor observables on a line.*

Proof. Follows from Theorems 4.5 and 4.10. ◀

4.3 "Dequantizing" HKP, MYZ for global Clifford measurements

Recall now that classical shadows constructed using global Clifford operations allow for an efficient recovery of the expectation values of observables whose Frobenius norm $\|O\|_F = \sqrt{\mathrm{Tr}[O^\dagger O]}$ is bounded. Interestingly, in this setting, we can solve the validity problem in polynomial time if we have sampling and query access to the target observables. Briefly, this is done by invoking the bounded Frobenius norm semidefinite programming (SDP) dequantization result of [10] (see also [11], which previously handled the low-rank case).

We begin by defining the Global Clifford version of CSV.

▶ **Definition 4.12** ($\mathsf{CSV}_{\mathsf{GC}}$)**.** *The definition is the same as Definition 1.2, only now our classical shadow has the structure dictated by the global Clifford measurement protocol presented in Ref. [21]. That means the following:*

- *The shadow S consists of $L = \mathrm{poly}(n)$ strings, $\{s_i\}_{i=1}^{L}$, each of which encodes the random n-qubit Clifford used in that round and the measurement outcome. More formally, each string will be of the form $s_i = (\mathsf{Stab}_i, b_i)$ where Stab is the efficient classical representation of the global Clifford via the stabilizer formalism and $b \in \{0,1\}^n$ the measurement outcome of that round.*
- *O is any set of $\mathrm{poly}(n)$ observables with bounded Frobenius norm, i.e., $\|O_i\|_F \leq \mathrm{poly}(n)$. Those observables are possibly highly non-local.*
- *The recovery algorithm applies the global inverse depolarizing channel to extract the snapshot operators $\hat{\eta}$ from S and aggregates the estimates via the* median of means *(MoM) technique.*

It will be easier to work in the abstract definition which we denote $\mathsf{ObsCon}_{\mathsf{F}}$ and define as:

▶ **Definition 4.13** ($\mathsf{ObsCon}_{\mathsf{F}}$)**.** *The input is a set of observables along with their respective expectation values $(O_i, y_i)_{i=1}^{m=\mathrm{poly}(n)}$, with $\|O_i\|_F \leq \mathrm{poly}(n)$, and parameters α and β satisfying $\beta - \alpha \geq 1/\mathrm{poly}(n)$. The output is to decide between the following cases:*

- ***Yes:*** *$\exists$ n-qubit state ρ s.t. $\forall\ i \in [m]$, $|\mathrm{Tr}\,(O_i\rho) - y_i| \leq \alpha$.*
- ***No:*** *$\forall$ n-qubit states ρ $\exists$ some $i \in [m]$ s.t. $|\mathrm{Tr}\,(O_i\rho) - y_i| \geq \beta$.*

We assume $y_i \in [-1, 1]$.

Let us now properly define what a sampling and query access to the target observables mean.

▶ **Definition 4.14** (Sampling and query access [10])**.** *For a vector $v \in \boldsymbol{C}^N$, sampling and query access, denoted $\mathsf{SQ}v$, means that we can query entries $v(i)$, sample indices $i \in [N]$ with probability $\frac{|v(i)|^2}{\|v\|_2^2}$, and compute $\|v\|_2$. Query access alone, denoted $\mathsf{Q}(v)$, means that we can query entries $v(i)$. For a matrix $A \in \boldsymbol{C}^{M\times N}$, query access, denoted $\mathsf{Q}(A)$, means that given $(i, j) \in [M] \times [N]$ one can compute $A(i, j)$; sampling and query access, denoted $\mathsf{SQ}(A)$, means that we have SQ-access to each row of A and SQ-access to the vector of row norms of A.*

With our sampling and query access definitions in hand, we can define:

▶ **Definition 4.15** ($\mathsf{ObsCon}_{\mathsf{F},\mathsf{Samp}}$)**.** *Defined as $\mathsf{ObsCon}_{\mathsf{F}}$ (see Definition 4.13) but additionally with sampling and query access (see Definition 4.14) to each observable O_i.*

We now recall the definition of the (SDP ε-feasibility)-problem [11, 10].

▶ **Definition 4.16** (SDP ε-feasibility [10])**.** *Given an $\varepsilon > 0$, m real numbers $b_1, \ldots, b_m \in \mathbb{R}$, and Hermitian $N \times N$ matrices $SQ(A^{(1)}), \ldots, SQ(A^{(m)})$ such that $-I \preceq A^{(i)} \preceq I$ for all $i \in [m]$, we define $\mathcal{S}_\varepsilon$ as the set of all ρ satisfying*

$$\mathrm{Tr}[A^{(i)}\rho] \leq b_i + \varepsilon,\ \ \forall i \in [m] \tag{7a}$$

$$\rho \succeq 0 \tag{7b}$$

$$\mathrm{Tr}[\rho] = 1 \tag{7c}$$

If $\mathcal{S}_\varepsilon = \emptyset$, output "infeasible". If $\mathcal{S}_0 \neq \emptyset$, output a $\rho \in \mathcal{S}_\varepsilon$.

▶ **Lemma 4.17.** $\mathsf{ObsCon}_{\mathsf{F},\mathsf{Samp}} \leq$ *SDP ε-feasibility.*

Proof. We start with a $\mathsf{ObsCon}_{\mathsf{F},\mathsf{Samp}}$ instance: $\{(\mathsf{SQ}(O_i), y_i)\}_{i=1}^m$ with $\|O_i\|_\infty \leq 1$, $\|O_i\|_F \leq \mathrm{poly}(n)$ and parameters α, β with $\beta - \alpha \geq 1/\mathrm{poly}$. Now define:

$$A_{i,\pm} := \pm O_i,\ \ b_{i,\pm} := \alpha \pm y_i,\ \ \varepsilon := \frac{\beta - \alpha}{2}$$

Consider now the following SDP

$$\mathrm{Tr}[A_{i,\pm}\rho] \leq b_{i,\pm} + \varepsilon,\ \ \forall i \in [m] \tag{8a}$$

$$\rho \succeq 0 \tag{8b}$$

$$\mathrm{Tr}[\rho] = 1 \tag{8c}$$

Since $\|O_i\|_\infty \leq 1$, one has $-I \preceq A_{i,\pm} \preceq I$ for all i. Thus, the above SDP is a valid instance of the SDP ε-feasibility problem. We now show correctness.

Completeness. Assume the $\mathsf{ObsCon}_{\mathsf{F},\mathsf{Samp}}$ instance is a YES instance, so that there exists a state ρ s.t. $|\,\mathrm{Tr}(O_i\rho) - y_i| \leq \alpha\ \ \forall\ i \in [m]$. Equivalently:

$$\exists \rho\ \text{ s.t. } \mathrm{Tr}(O_i\rho) \leq y_i + \alpha,\ \ \mathrm{Tr}(-O_i\rho) \leq \alpha - y_i\ \ \forall\ i \in [m].$$

In terms of the SDP constraints: $\exists\ \rho$ s.t. $\mathrm{Tr}(A_{i,\pm}\rho) \leq b_{i,\pm}\ \ \forall\ i \in [m]$. Therefore the associated SDP instance is feasible with zero slack, i.e., $\mathcal{S}_0 \neq \emptyset$.

Soundness. Assume $\mathsf{ObsCon}_{\mathsf{F,Samp}}$ instance is a NO instance. Then for every state ρ there exists some index $i^* \in [m]$ such that $|\operatorname{Tr}(O_{i^*}\rho) - y_{i^*}| \geq \beta$. Let $g := \beta - \alpha$. Then for the index i^*, one of the following must hold

$$\operatorname{Tr}(O_{i^*}\rho) - y_{i^*} \geq \alpha + g \quad \text{or} \quad -(\operatorname{Tr}(O_{i^*}\rho) - y_{i^*}) \geq \alpha + g$$

Rewriting these in terms of the SDP constraints we get: $\operatorname{Tr}(A_{i^*,\pm}\rho) \geq b_{i^*,\pm} + g$. So every ρ violates at least one SDP constraint by at least $g = \beta - \alpha = 2\varepsilon$. Therefore $\mathcal{S}_\varepsilon = \emptyset$. ◀

▶ **Lemma 4.18** (Corollary 6.25 [10]). *Let $F \geq \max_{j\in[m]}(\|A_j\|_F)$, and suppose $F = \Omega(1)$. Then we can solve 4.16 with success probability $\geq 1 - \delta$ in cost*

$$\widetilde{\mathcal{O}}\left(\left(\frac{F^{18}}{\varepsilon^{40}}\log^{20}(N)\boldsymbol{sq}(A) + \frac{F^{22}}{\varepsilon^{46}}\log^{23}(N) + m\frac{F^{8}}{\varepsilon^{18}}\log^{8}(N)\boldsymbol{q}(A) + m\frac{F^{14}}{\varepsilon^{28}}\log^{13}(N)\right)\log^{3}\frac{1}{\delta}\right)$$

providing sampling and query access to a solution.

▶ **Theorem 4.19.** $\mathsf{ObsCon}_{\mathsf{F,Samp}}$*, and hence* $\mathsf{CSV}_{\mathsf{GC}}$ *under query and sampling access, is solvable in randomized classical polynomial time.*

Proof. From Lemmas 4.17 and 4.18 and Remark 4.1, since $\|O_i\|_F \leq \operatorname{poly}(n)\ \forall\ i \in [m]$. ◀

We now state the qudit analogue for the global n-qudit Clifford ensemble protocol (see Section 2).

▶ **Definition 4.20** ($\mathsf{ObsCon}^{(d)}_{\mathsf{F,Samp}}$). $\mathsf{ObsCon}^{(d)}_{\mathsf{F,Samp}}$ *is the qudit analogue of* $\mathsf{ObsCon}_{\mathsf{F,Samp}}$ *(see Definition 4.15): inputs $(O_i, y_i)_{i=1}^m$ with Hermitian n-qudit O_i satisfying $\|O_i\|_F \leq \operatorname{poly}(n)$ and $\beta - \alpha \geq 1/\operatorname{poly}(n)$, together with sampling and query access.*

▶ **Theorem 4.21.** $\mathsf{ObsCon}^{(d)}_{\mathsf{F,Samp}}$ *is solvable in randomized classical polynomial time.*

Proof. The proof of Theorem 4.19 is dimension-independent. If N denotes the Hilbert-space dimension of the SDP variable, then the SDP solver depends only polylogarithmically on N. Replacing the qubit dimension $N = 2^n$ by the qudit dimension $N = d^n$ therefore changes the logarithmic dimension factor from $\log(2^n) = n$ to $\log(d^n) = n\log d$. Hence, for fixed local dimension d and $\|O_i\|_F \leq \operatorname{poly}(n)$, the same reduction to SDP ε-feasibility gives a randomized classical polynomial-time algorithm for $\mathsf{ObsCon}^{(d)}_{\mathsf{F,Samp}}$. ◀

References

1 Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 325–338, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3188745.3188802.

2 Scott Aaronson and Alex Arkhipov. The Computational Complexity of Linear Optics. In *Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 333–342, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993682.

3 Avantika Agarwal, Sevag Gharibian, Venkata Koppula, and Dorian Rudolph. Quantum polynomial hierarchies: Karp-Lipton, error reduction, and lower bounds. In *49th International Symposium on Mathematical Foundations of Computer Science (MFCS 2024)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.MFCS.2024.7.

4 D. Aharonov and O. Regev. A lattice problem in quantum NP. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003.*, pages 210–219, 2003. doi:10.1109/SFCS.2003.1238195.

5 Dorit Aharonov and Tomer Naveh. Quantum np - a survey, 2002. arXiv:quant-ph/0210077.

6 Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018. doi:10.1038/s41567-018-0124-x.

7 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2019.27.

8 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, August 2022. doi:10.1137/21m140729x.

9 E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theor.*, 52(2):489–509, 2006. doi:10.1109/TIT.2005.862083.

10 Nai-Hui Chia, András Pal Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *J. ACM*, 69(5), 2022. doi:10.1145/3549524.

11 Nai-Hui Chia, Tongyang Li, Han-Hsuan Lin, and Chunhao Wang. Quantum-Inspired Sublinear Algorithm for Solving Low-Rank Semidefinite Programming. In Javier Esparza and Daniel Král', editors, *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, volume 170 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2020.23.

12 Simon Foucart and Holger Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser/Springer, New York, 2013. doi:10.1007/978-0-8176-4948-7.

13 Sevag Gharibian and Julia Kempe. Hardness of approximation for quantum problems. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming*, Lecture Notes in Computer Science, pages 387–398, Berlin, Heidelberg, 2012. Springer. doi:10.1007/978-3-642-31594-7_33.

14 Sevag Gharibian, Miklos Santha, Jamie Sikora, Aarthi Sundaram, and Justin Yirka. Quantum generalizations of the polynomial hierarchy with applications to QMA(2). In Igor Potapov, Paul Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018)*, volume 117 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 58:1–58:16, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2018.58.

15 Daniel Gottesman and Sandy Irani. The quantum and classical complexity of translationally invariant tiling and Hamiltonian problems. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 95–104, 2009. doi:10.1109/FOCS.2009.22.

16 Sabee Grewal and Justin Yirka. The entangled quantum polynomial hierarchy collapses. In *39th Computational Complexity Conference (CCC 2024)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.CCC.2024.6.

17 David Gross, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105:150401, October 2010. doi:10.1103/PhysRevLett.105.150401.

18 Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 913–925, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2897518.2897585.

19 Sean Hallgren, Daniel Nagaj, and Sandeep Narayanaswami. The local hamiltonian problem on a line with eight states is qma-complete. *Quantum Info. Comput.*, 13(9–10):721–750, 2013. doi:10.26421/QIC13.9-10-1.

20 Craig S. Hamilton, Regina Kruse, Linda Sansoni, Sonja Barkhofen, Christine Silberhorn, and Igor Jex. Gaussian Boson Sampling. *Phys. Rev. Lett.*, 119(17):170501, 2017. doi:10.1103/PhysRevLett.119.170501.

21 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, June 2020. doi:10.1038/s41567-020-0932-7.

22 Rahul Jain and John Watrous. Parallel approximation of non-interactive zero-sum quantum games, 2008. arXiv:0808.2775.

23 Georgios Karaiskos, Dorian Rudolph, Johannes Jakob Meyer, Jens Eisert, and Sevag Gharibian. How hard is it to verify a classical shadow?, 2025. doi:10.48550/arXiv.2510.08515.

24 Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. Triply efficient shadow tomography. *PRX Quantum*, 6(1):010336, 2025. doi:10.1103/PRXQuantum.6.010336.

25 A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.

26 Hirotada Kobayashi, Keiji Matsumoto, and Tomoyuki Yamakami. Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur? In Toshihide Ibaraki, Naoki Katoh, and Hirotaka Ono, editors, *Algorithms and Computation*, Lecture Notes in Computer Science, pages 189–198, Berlin, Heidelberg, 2003. Springer. doi:10.1007/978-3-540-24587-2_21.

27 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*, Lecture Notes in Computer Science, pages 438–449, Berlin, Heidelberg, 2006. Springer. doi:10.1007/11830924_40.

28 Yi-Kai Liu. The local consistency problem for stoquastic and 1-D quantum systems, 2007. arXiv:0712.1388.

29 Chengsi Mao, Changhao Yi, and Huangjun Zhu. Qudit shadow estimation based on the clifford group and the power of a single magic gate. *Physical Review Letters*, 134(16), April 2025. doi:10.1103/physrevlett.134.160801.

30 Ryan O'Donnell and John Wright. Efficient quantum tomography, 2015. arXiv:1508.01907.

31 R. Oliveira and B. M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Information & Computation*, 8(10):0900–0924, 2008. doi:10.26421/QIC8.10-2.

32 Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 217–228, 2019. doi:10.1145/3313276.3316310.