CONJURA · OPEN PROBLEM

No Tableau-Randomizing Self-Reduction for Learning Stabilizer with Noise

A barrier the source can exhibit but not close

Status. Statement: AI-written from Andrey Boris Khesin, Jonathan Z. Lu, Alexander Poremba, Akshar Ramkumar and Vinod Vaikuntanathan, *Average-Case Complexity of Quantum Stabilizer Decoding*, IACR ePrint 2025/1769. Not yet checked by a human. This is the source's Conjecture 5.20. It proves a conditional barrier via what it calls the scrambling gap, exhibits an analogous classical no-go, and conjectures that the barrier can be made unconditional for tableau-randomizing reductions. It is careful that this is not an impossibility for random self-reductions in general.

Category. *Quantum Cryptography.*

Conjecture (informal). *Classical coding theory has a prized tool: a worst-case to average-case reduction, which turns an arbitrary hard decoding instance into a random one. For quantum stabilizer codes nobody has one, and the source explains why it is harder than it looks. A stabilizer instance carries two Pauli objects — the code and the error — and a reduction has to scramble both at once, but making one maximally random destroys the other's decodability. The source turns that tension into a quantitative barrier and conjectures that, for reductions that work by randomizing the stabilizer tableau, the barrier is absolute.*

1 The setting

The source's headline result is that decoding a random stabilizer code with even a single logical qubit is at least as hard as decoding a random classical code at constant rate — the maximally hard regime — so any sub-exponential algorithm for a typical stabilizer code, at any rate, would be a breakthrough in cryptography. That is proved by a direct reduction from LPN to LSN, not by a self-reduction.

The self-reduction question is separate and stays open. Classically, Brakerski, Lyubashevsky, Vaikuntanathan and Wichs obtain

a random self-reduction to LPN, but only for a restricted subset of linear codes; the restriction is what circumvents the analogous barrier. In the stabilizer setting the source first proves that no random self-reduction from an arbitrary worst-case QNCP instance into LSN can exist, naming the obstruction precisely: "The culprit preventing the existence of such a reduction is exchange symmetry between the code and the error — both are Pauli strings which are randomly scrambled during the reduction, but one must be maximally random while the other cannot be maximally random (less we lose decodability), which is impossible." It then observes that an analogous no-go holds classically too, which is why the classical restriction trick is the thing to imitate.

What the source can prove beyond that is conditional. It defines a quantity it calls the *scrambling gap* $\eta(\varepsilon)$ and shows that if $\eta(\varepsilon) = O(\mathrm{poly}(n)\cdot\varepsilon^b)$ for some constant $b > 0$, then no number of repeated samples from a candidate reduction distribution yields a strong and non-trivial self-reduction. Conjecture 1 is the assertion that this route closes unconditionally in the tableau-randomizing case.

2 Notation and parameters

- n is the number of physical qubits and k the number of logical qubits; C_n is the Clifford group on n qubits and $\mathrm{Stab}(n,k)$ the set of $[[n,k]]$ stabilizer codes.
- $\mathcal{D}_p$ is the depolarizing channel with probability p, and $\mathcal{D}_{n,k,p} = \mathrm{Unif}(\mathrm{Stab}(n,k)) \otimes \mathcal{D}_p^{\otimes n}$.
- $\mathbf{C}$ is the Clifford circuit associated with a stabilizer tableau, $\mathbf{E}$ a Pauli error with $\mathrm{wt}(\mathbf{E}) \le w$, and $\mathbf{R}$ the Clifford drawn by the reduction.
- ε measures the quality of the reduction; it is called *strong* when $\varepsilon = \mathrm{negl}(n)$.
- p is the error rate in the random case. A reduction with $p = \frac{3}{4}$ is always possible and worthless: that error is equivalent to replacing the encoded state by the maximally mixed state. A reduction is *non-trivial* when $p = \frac{3}{4} - \frac{1}{\mathrm{poly}(n)}$.
- $\mathcal{V} \subseteq C_n$ is the restriction to a set of worst-case Clifford operators, the analogue of the classical restriction on the generator matrix.

3 The conjecture

Definition 1 (Random self-reduction distribution, [KLPRV25, Definition 5.7]). Fix n, k. A $(\mathcal{V}, \varepsilon, w, p)$ *random self-reduction distribution* is an efficiently sampleable distribution $\mathcal{R}_n$ over Cliffords C_n satisfying

$$\mathsf{TV}_{\mathbf{R}\sim\mathcal{R}_n}\left((\mathbf{RC}, \mathbf{RER}^\dagger), \mathcal{D}_{n,k,p}\right) \leq \varepsilon$$

for all $\mathbf{C}$ whose stabilizer code lies in $\mathcal{V}$, and all $\mathbf{E}$ with $\mathrm{wt}(\mathbf{E}) \leq w$.

Conjecture 1 ([KLPRV25, Conjecture 5.20]). *There is no $(\mathcal{V}, \varepsilon, w, p)$ tableau-randomizing random self-reduction operator, in the sense of Definition 1, for* LSN*, for any $\mathcal{V} \subseteq C_n$ and any $w \geq 1$, such that $\varepsilon = \mathrm{negl}(n)$ and $p = \frac{3}{4} - \frac{1}{\mathrm{poly}(n)}$.*

Remark 1 (What the source proves, and the shape of the gap). The barrier it establishes is conditional on the scrambling gap being polynomially related to ε: under $\eta(\varepsilon) = O(\mathrm{poly}(n)\varepsilon^b)$, for any k, w and $\mathcal{V} \subseteq \mathrm{Stab}(n, k)$, no positive integer d makes the distribution induced by d samples from $\mathcal{R}_n$ a strong and non-trivial self-reduction. So the missing step is a bound on the scrambling gap itself, and the source says what it expects to supply it: "analysis techniques generalizing the scrambling gap formalism will reveal unconditionally that a reduction cannot scramble the tableau."

Remark 2 (The two parameters are what make the statement non-vacuous). Both restrictions in Conjecture 1 are load-bearing and neither can be dropped. Drop $\varepsilon = \mathrm{negl}(n)$ and the claim is about a weaker object the source explicitly does not address. Drop $p = \frac{3}{4} - 1/\mathrm{poly}(n)$ and the statement is false, since $p = \frac{3}{4}$ is always achievable. A refutation therefore has to hit both targets at once, which is the same coupling that creates the barrier.

Remark 3 (Explicitly outside the statement). The source is careful about what its barriers do not do, and the same care belongs here. "[T]hese barriers do not definitely prove the non-existence of a random self-reduction." Three escapes it names, none of which Conjecture 1 rules out: reductions that are not tableau-randomizing;

reductions where the dimensions n, k increase by a polynomial factor, i.e. where $\mathbf{R}$ is an isometry rather than a unitary; and reductions with $\varepsilon = 1/\mathrm{poly}(n)$, for which non-triviality would require a smaller p, e.g. $p < \frac{3}{4} - \Omega(1)$. So proving the conjecture would close one route and leave those three open, and a resolution should say which it settles.

Remark 4 (Neighbouring open questions in the same section). Several, all distinct from this statement. Whether there is a reduction from $\mathsf{LSN}(k, n, p)$ to $\mathsf{LPN}(\Theta(np), \Theta(n), \Theta(p))$. Whether the source's search-to-decision reduction, which needs a $\mathsf{QNCP}(k, n, w')$ solver for *every* $w' \leq w$, can be weakened to one that works for weight exactly w. Whether a decision-to-search reduction exists for Search QNCP or recQNCP, where the difficulty is that proposed solutions cannot be efficiently verified. And whether the average-case search-decision equivalence for $\mathsf{LSN}_{\mathrm{poly}}$, which relies on obtaining multiple samples, holds with a single sample.

Remark 5 (Why a self-reduction is wanted even though hardness is already proved). The source establishes average-case hardness directly from LPN, so the self-reduction is not needed for that. What it would give is a statement of a different kind — hardness of typical instances from hardness of worst-case instances of the same problem, without importing a classical assumption — and, as the source puts it for the neighbouring reduction question, it "would help shed light on the the exact relationship between quantum and classical decoding more generally."

4 Bibliography

[KLPRV25] Andrey Boris Khesin, Jonathan Z. Lu, Alexander Poremba, Akshar Ramkumar and Vinod Vaikuntanathan. *Average-case complexity of quantum stabilizer decoding.* IACR ePrint 2025/1769. The source. Definition 5.7 is on page 52, the conditional barrier via the scrambling gap on page 57, and Conjecture 5.20 with the surrounding discussion of what it does and does not rule out on page 62.

[BLVW19] Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan and Daniel Wichs. *Worst-case hardness for LPN and cryptographic hashing via code smoothing.* EUROCRYPT 2019. The classical random self-reduction the source compares against, which circumvents the analogous barrier by restricting the set of possible worst-case instances.

[UNVERIFIED] — cited in the source as [BLVW19]; the bibliographic detail here is inferred.