Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Adversarial channel equilibria for qLDPC benchmarks]{Adversarial channel equilibria for bivariate-bicycle qLDPC memories: minimax fidelity guarantees and reproducible decoder benchmarks

Saleh H. AlDaajeh

s.aldaajeh@gmail.com

Independent Researcher https://orcid.org/0000-0001-7810-9290

Saud Aldajah

saud.aldajah@sma.c.ae

College of International Transport and Logistics, Sharjah Maritime Academy https://orcid.org/0000-0001-6061-1004

Research Article

Keywords: quantum error correction, qLDPC codes, bivariate-bicycle codes, BP+OSD, minimum-weight perfect matching, reproducibility, adversarial channels

Posted Date: June 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10039013/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Article

Adversarial channel equilibria for bivariate-bicycle qLDPC memories: minimax fidelity guarantees and reproducible decoder benchmarks

Saleh H. AlDaajeh[1*] and Saud Aldajah[2]

[1*]Enterprise Command and Control Centre (EC3), Roads and Transportation Authority, Dubai, United Arab Emirates.
[2]College of International Transport and Logistics, Sharjah Maritime Academy, Khorfakkan, Sharjah, United Arab Emirates.

*Corresponding author(s). E-mail(s): S.aldaajeh@gmail.com;
Contributing authors: Saud.Aldajah@sma.ac.ae;

Abstract

Quantum low-density parity-check (qLDPC) memories promise lower-overhead fault tolerance, but reproducible finite-instance benchmarking has been hindered by the absence of explicit decoder representations and data provenance. We introduce adversarial channel equilibria (ACE), a bounded-Pauli channel game that organises fidelity guarantees, certificates and decoder benchmarks for bivariate-bicycle (BB) memories. Applying Sion's minimax theorem to compact convex sets of completely positive trace-preserving maps, we prove existence of a saddle-point equilibrium and derive a closed-form fidelity floor. We provide the first SHA-256-anchored three-code benchmark comparing BP+OSD against two graphified baselines—DEM-chain MWPM and UF-BFS—on BB72, BB144 and BB288, using 10,000 trials per point and deterministic seeds. BP+OSD achieves risk ratios of **595**× to **67.8**× over DEM-chain MWPM on BB72 across tested error rates; at larger code sizes BP+OSD failure counts are zero while the graphified baselines fail at rates above 50%. Repeated-cycle BP+OSD remains near zero through the tested range. A BB72 certificate verifies commutation, dimension, distance lower bound and radius-two correctability. Hardware-noise simulations are reported as model sensitivities only. All claims are mapped to verifiable data, code and SHA-256 hash artefacts archived at doi:10.5281/zenodo.20588303.

Keywords: quantum error correction, qLDPC codes, bivariate-bicycle codes, BP+OSD, minimum-weight perfect matching, reproducibility, adversarial channels

Introduction

Fault-tolerant quantum computation and post-quantum cryptography are now linked strategic problems. Shor's algorithm established the cryptanalytic relevance of scalable quantum computers [1], while the standardization of ML-KEM, ML-DSA and SLH-DSA has made quantum-resistant classical infrastructure an engineering priority [2–4]. These developments do not imply that quantum error correction is itself a cryptographic primitive. They do, however, motivate a careful study of quantum memories under explicitly adversarial noise models, because a future distributed quantum system must preserve encoded quantum information while operating in a hostile, observable and failure-prone environment.

Surface-code architectures have dominated practical fault-tolerance studies because their local checks are compatible with two-dimensional hardware layouts. Their locality, however, imposes demanding overheads for large computations. Quantum low-density parity-check (qLDPC) codes offer a complementary route: sparse checks, constant rate in good families and potentially much larger distances at fixed physical-qubit budgets [5, 6]. Bivariate-bicycle (BB) codes are particularly important because recent work demonstrated high-rate memory protocols with compelling overhead reductions [7]. Networked and distributed variants are also emerging, including teleported-check constructions and circuit-level BP+OSD studies [8]. A benchmark paper in this area must therefore be precise about the code instance, decoder representation, noise model, number of rounds, random seed policy and exact data provenance.

Decoder representation is the central empirical issue addressed here. BP+OSD operates directly on sparse parity-check matrices and is a natural baseline for qLDPC codes [9, 10]. By contrast, minimum-weight perfect matching (MWPM) algorithms such as PyMatching and Sparse Blossom are designed for graphlike detector error mechanisms [11, 12]. A BB data-qubit error can activate three checks, so a graph-decoder baseline requires an explicit graphification of syndrome hyperedges. This representation step can discard correlations. Claims about MWPM, union-find or any graph-based decoder are therefore meaningful only after the graphification rule is specified. The present paper evaluates two such graphified reference baselines, DEM-chain MWPM and UF-BFS, and does not claim that all possible graph, matching or hypergraph-aware decoders are inferior.

The theoretical contribution is a deliberately modest formal wrapper. We introduce adversarial channel equilibria (ACE) as a game in which a defender selects an encoded recovery map and an adversary selects a bounded-weight Pauli channel, with entanglement fidelity as payoff. The minimax existence theorem is an application of Sion's theorem to compact convex completely positive trace-preserving sets [13]; the fidelity floor is a specialization of Knill-Laflamme correctability [14]; and the stochastic lower bound conditions on the binomial probability of staying within the correctable radius. This framework is useful because it separates theorem-backed claims from certificate-backed and Monte Carlo-backed claims. It is not a cryptographic proof of IND-CPA, IND-CCA or universal-composability security.

ORCID: Saleh H. AlDaajeh, 0000-0001-7810-9290; Saud Aldajah, 0000-0001-6061-1004.

The empirical contribution is a reproducible finite-size benchmark for BB72, BB144 and BB288. The study combines: a BB72 finite-instance certificate; single-round BP+OSD versus two graphified baselines at 10,000 trials per point; repeated-cycle BP+OSD with d syndrome rounds per trial; and hardware-noise-model sensitivity experiments. Each stochastic row records raw failures, total trials, deterministic SHA-256-derived seeds and matrix digests. The manuscript is organized in the format recommended for an npj Quantum Information Article: Results present the scientific findings, Discussion states scope and limitations without a separate conclusions section, and Methods contain the mathematical and computational details needed for replication.

1 Results

1.1 Adversarial channel equilibria give a bounded-noise fidelity floor

Let $\mathfrak{D}$ be a non-empty compact convex set of defensive completely positive trace-preserving maps and let $\mathfrak{A}$ be a non-empty compact convex set of adversarial channels on a finite-dimensional encoded Hilbert space. For a defense $\mathcal{D} \in \mathfrak{D}$ and adversary $\mathcal{A} \in \mathfrak{A}$, define the payoff $u(\mathcal{D}, \mathcal{A})$ as the entanglement fidelity after attack and recovery. The ACE value is

$$V^\star = \max_{\mathcal{D}\in\mathfrak{D}} \min_{\mathcal{A}\in\mathfrak{A}} u(\mathcal{D}, \mathcal{A}). \tag{1}$$

Theorem 1 (Minimax existence)**.** *For finite-dimensional compact convex CPTP strategy sets, u is affine and continuous in each argument, and*

$$\max_{\mathcal{D}\in\mathfrak{D}} \min_{\mathcal{A}\in\mathfrak{A}} u(\mathcal{D}, \mathcal{A}) = \min_{\mathcal{A}\in\mathfrak{A}} \max_{\mathcal{D}\in\mathfrak{D}} u(\mathcal{D}, \mathcal{A}). \tag{2}$$

Theorem 2 (Bounded-weight fidelity floor)**.** *Let $\mathcal{C} = [[n, k, d]]$ be a stabilizer code and $t = \lfloor (d-1)/2 \rfloor$. For channels whose Kraus operators lie in the span of Pauli errors of weight at most t, ideal recovery gives $F_e = 1$. If the implemented recovery differs from ideal recovery by at most $\varepsilon_{\mathrm{FT}}$ in diamond norm, then $F_e \geq 1 - \varepsilon_{\mathrm{FT}}$.*

Proposition 1 (I.i.d. stochastic Pauli lower bound)**.** *For i.i.d. Pauli attacks with physical error probability p_{atk} and implementation error δ_{FT},*

$$V^\star_{\mathrm{iid}} \geq (1 - \delta_{\mathrm{FT}}) \Pr[\mathrm{Bin}(n, p_{\mathrm{atk}}) \leq t]. \tag{3}$$

For the illustrative configuration $n = 128$, $d = 11$, $t = 5$ and $\delta_{\mathrm{FT}} = 0.1$, the lower bound is $0.9 \Pr[\mathrm{Bin}(128, p) \leq 5]$ (Table 1). This is not a threshold estimate; it is a closed-form, fixed-channel lower bound for the specified stochastic subclass.

1.2 BB72 is certified through the radius relevant to adversarial correction

The BB72 instance was verified with an archive-local certificate script. The certificate checks $H_X H_Z^T = 0$ over $\mathbb{F}_2$, confirms $k = 72 - 30 - 30 = 12$, performs a

Table 1 I.i.d. stochastic fidelity lower bound for $n = 128$, $d = 11$, $t = 5$ and $\delta_{\mathrm{FT}} = 0.1$.

p	LB	p	LB	p	LB
0.001	0.900	0.020	0.860	0.050	0.341
0.005	0.900	0.030	0.731	0.080	0.046
0.010	0.898	0.040	0.534	0.100	0.008

Table 2 Single-round decoder comparison. Values are failures per 10,000 trials with logical failure rate in parentheses. Full Wilson intervals, seeds and SHA-256 matrix digests are provided in the archive [15].

Code	Decoder	$p = 0.001$	$p = 0.004$	$p = 0.008$	$p = 0.012$
BB72	BP+OSD	0 (< 0.000384)	1 (0.00010)	14 (0.00140)	50 (0.00500)
BB72	MWPM-DEM	46 (0.00460)	595 (0.05950)	1894 (0.18940)	3392 (0.33920)
BB72	UF-BFS	–	614 (0.06140)	1903 (0.19030)	3623 (0.36230)
BB144	BP+OSD	0 (< 0.000384)	0 (< 0.000384)	0 (< 0.000384)	6 (0.00060)
BB144	MWPM-DEM	170 (0.01700)	1925 (0.19250)	4988 (0.49880)	7297 (0.72970)
BB288	BP+OSD	0 (< 0.000384)	0 (< 0.000384)	0 (< 0.000384)	0 (< 0.000384)
BB288	MWPM-DEM	625 (0.06250)	5059 (0.50590)	8667 (0.86670)	9669 (0.96690)

meet-in-the-middle search through weight five with zero undetected logical witnesses, and exhaustively verifies radius-two correctability over all 2,629 weight-at-most-two patterns. The certificate therefore supports $d \geq 6$ and $t = 2$ for the submitted matrix representation. It does not independently prove equality $d = 6$ by exhibiting a weight-six logical operator.

The certificate is important for the ACE interpretation: bounded-weight adversarial guarantees are theorem-backed only within the certified correction radius. Random Bernoulli simulations do not replace the exhaustive certificate; they evaluate decoder performance under stochastic sampling.

1.3 BP+OSD outperforms the two tested graphified baselines

The single-round benchmark uses one noiseless syndrome cycle, independent Bernoulli Pauli noise, 10,000 trials per point and the same code constructions for every decoder. Table 2 summarizes representative rows from the three-code dataset. BP+OSD uses the Tanner graph directly, whereas DEM-chain MWPM and UF-BFS use graphified representations of the same syndrome model.

For BB72, finite risk ratios of MWPM-DEM relative to BP+OSD are 595× at $p = 0.004$, 135.3× at $p = 0.008$ and 67.8× at $p = 0.012$. At higher code sizes several BP+OSD rows have zero failures, so the ratio is reported through one-sided upper bounds rather than interpreted as a finite asymptotic quantity. The valid conclusion is narrow: in this implementation and under this noise model, the two tested graphified baselines have substantially higher logical failure rates than BP+OSD. The result

Table 3 Repeated-cycle BP+OSD benchmark. Entries are failures/trials with logical failure rate in parentheses.

Code	rounds	$p = 0.006$	$p = 0.008$	$p = 0.012$
BB72	6	6/2000 (0.0030)	10/2000 (0.0050)	47/2000 (0.0235)
BB144	12	0/2000 (< 0.00192)	2/2000 (0.0010)	0/2000 (< 0.00192)
BB288	18	0/500 (< 0.00765)	0/500 (< 0.00765)	0/500 (< 0.00765)

Table 4 BB72 hardware-noise-model sensitivity at $p_{\text{base}} = 5 \times 10^{-4}$, 3,000 trials per condition. These are simulation-model ablations, not real-device causal estimates.

Condition	Failures	LFR	Wilson 95% CI
Phenomenological baseline	3	0.0010	[0.00034, 0.00294]
CZ full model	30	0.0100	[0.00701, 0.01424]
CZ no SWAP overhead	15	0.0050	[0.00303, 0.00823]
CZ no idle noise	30	0.0100	[0.00701, 0.01424]
CZ no measurement noise	7	0.0023	[0.00113, 0.00481]
CZ gate decoherence only	0	< 0.00128	Wilson upper bound
MS full model	7	0.0023	[0.00113, 0.00481]
MS CZ-matched gate error	19	0.0063	[0.00406, 0.00987]

should not be read as a proof that every possible graph, matching or correlation-aware decoder is unsuitable for BB codes.

1.4 Repeated-cycle BP+OSD shows finite-size evidence of sub-threshold ordering

The repeated-cycle benchmark uses BP+OSD on a space-time syndrome representation with d rounds per trial: six rounds for BB72, twelve for BB144 and eighteen for BB288. Table 3 reports the finite-size observations. The evidence is consistent with sub-threshold ordering across the tested points but does not constitute a formal threshold study, because only three code instances and finite shot budgets are used.

The most defensible interpretation is that increasing BB code size improves BP+OSD performance under the submitted phenomenological repeated-cycle protocol. The result is finite-size evidence only. It should not be described as an independently established threshold or finite-size scaling collapse.

1.5 Hardware-noise simulations identify model sensitivities rather than real-device causality

The hardware section is restricted to BB72 and to simulation-model sensitivity. The archive does not supply compiled circuits, schedule JSONs, routing files or device-calibration records; therefore this result is not presented as a hardware experiment or hardware-compiled benchmark. Table 4 reports the controlled ablation at $p_{\text{base}} = 5 \times 10^{-4}$.

Within this simulation model, removing SWAP overhead lowers CZ LFR from 0.0100 to 0.0050, while removing measurement noise lowers it to 0.0023. Replacing MS gate parameters by CZ-matched gate error raises MS LFR from 0.0023 to 0.0063. These shifts show that routing, measurement readout and backend gate parameters are all material in the specified model. No single-cause explanation or real-hardware causal attribution is claimed.

2 Discussion

The central result is not a new asymptotic threshold theorem, nor a universal statement about MWPM. The paper provides a finite-instance, reproducible comparison in which BP+OSD uses the BB Tanner graph directly, whereas DEM-chain MWPM and UF-BFS use graphified reductions. The large observed gap is scientifically meaningful because it quantifies the cost of two explicit graphified baselines on the same BB instances. It is also deliberately limited: a correlation-aware, hypergraph-aware or differently weighted matching construction could perform differently and should be compared in future work.

The ACE formalism helps clarify evidence categories. Theorems cover bounded-weight Pauli channels and ideal or approximately implemented recovery; the BB72 certificate supports the correction radius used by those theorems; Monte Carlo data support inite-size decoder comparisons; and hardware-noise simulations support only model sensitivities. This separation is crucial for a quantum-information audience because logical failure rate is not a diamond-norm bound, and entanglement-fidelity guarantees are not cryptographic security definitions. The work therefore complements, rather than replaces, formal quantum-cryptographic security frameworks.

The repeated-cycle dataset is encouraging but not final. BB144 and BB288 show low observed failure rates at the tested points, but the shot budgets are finite and the study does not perform a scaling collapse. The proper claim is finite-size evidence consistent with sub-threshold operation under the submitted protocol. Establishing an independent threshold would require larger code families, larger shot budgets, explicit stopping rules and a finite-size scaling analysis.

The hardware-noise section should be interpreted conservatively. The ablations show that measurement noise, SWAP overhead and gate parameters can each move the observed LFR by comparable amounts in the specified model. Because compiled circuits, routing schedules and calibration data are not supplied, the manuscript avoids hardware-compiled and real-device causal language. A future hardware-focused paper should include schedule files, compiled circuits, calibrated noise JSONs, raw logs and a matched ablation design for each architecture.

The main limitation is scope. The adversary is bounded-weight or i.i.d. stochastic Pauli noise, not an adaptive cryptanalytic adversary. The BB72 certificate proves $d \geq 6$ for the submitted representation but not an exact equality witness. The decoder-dependent ACE gap uses LFR as an empirical proxy and is not a theorem about diamond norm. These limitations are not defects if kept explicit; they are the boundary conditions under which the reported results can be trusted.

3 Methods

3.1 Code constructions and matrix provenance

The three BB code instances are constructed over $\mathbb{F}_2[x, y]/(x^\ell - 1, y^m - 1)$. BB72 uses $(\ell, m) = (6, 6)$ with $A = x^3 + y + y^2$ and $B = y^3 + x + x^2$. BB144 uses $(\ell, m) = (12, 6)$ with the same polynomials. BB288 uses $(\ell, m) = (12, 12)$ with $A = x^3 + y^2 + y^7$ and $B = y^3 + x + x^2$, following the BB288 instance reported by Bravyi et al. [7]. Every data row stores full SHA-256 digests of H_X and H_Z matrix bytes. The certificate and stochastic benchmarks reject mismatched matrices.

3.2 Decoder implementations

BP+OSD is implemented through QEC Emulator v2.0.0 [16] using the `ldpc` package [10]. The configuration uses OSD-CS order 7, min-sum belief propagation with scaling 0.625 and 100 maximum iterations. Logical failure is counted when the decoded correction disagrees with the sampled logical effect.

DEM-chain MWPM is implemented as a graphified reference baseline compatible with graphlike matching infrastructure [11, 12]. Weight-three syndrome columns are reduced to chain edges and boundary handling is specified in the archive. UF-BFS uses a different greedy graph-decoding rule but the same weight-three-to-edge reduction. Both are therefore graphified baselines, not hyperedge-aware qLDPC decoders.

3.3 Monte Carlo protocols and seeds

Single-round trials use one noiseless syndrome cycle with independent Bernoulli Pauli noise at each tested p. Repeated-cycle trials use d syndrome rounds per trial for the corresponding code instance. Hardware-noise simulations use a radius-two lookup decoder for BB72 and the model conditions in Table 4. Each task uses a deterministic 32-bit seed obtained from the first four bytes of SHA-256 applied to the context string containing the code, decoder, protocol and physical error rate. This avoids Python's process-dependent built-in hash randomization.

3.4 Statistical analysis

Logical failure rate is the number of failed trials divided by total trials. Binomial confidence intervals are Wilson intervals unless a zero-failure row is reported only through its Wilson upper bound. Risk ratios are reported only where the numerator and denominator are interpretable from the raw count table; zero-denominator rows are treated as one-sided evidence rather than finite ratios. The study does not use non-overlapping confidence intervals as a substitute for a formal threshold claim.

3.5 Certificate verification

The BB72 certificate evaluates four obligations: $H_X H_Z^T = 0$ over $\mathbb{F}_2$; $k = 72 - \mathrm{rank}(H_X) - \mathrm{rank}(H_Z) = 12$; absence of undetected logical witnesses through weight

five; and exhaustive correction of all weight-at-most-two patterns. The distance-lower-bound search uses a meet-in-the-middle enumeration. The radius-two check enumerates all 2,629 patterns and confirms zero correction failures.

3.6 Reproducibility

All results in this manuscript are independently reproducible from the public archive at doi:10.5281/zenodo.20588303 [15]. The archive contains the manuscript source, compiled PDF, all primary CSV datasets, SHA-256 manifest (42 files), verification scripts, and the declared software environment (`ldpc==2.4.1`, `pymatching==2.4.0`, `numpy==2.4.4`). Running `python reproduce_all.py` in the declared environment regenerates all stochastic tables. Running `python verify_hashes.py` confirms the 42-file manifest. Running `python verify_certificate.py` confirms all four BB72 certificate obligations.

3.7 Ethics statement

This study used no human participants, human data, animals or biological material. No ethics approval was required.

Data availability. All primary datasets generated and analyzed in this study are publicly available in the Zenodo reproducibility archive at doi:10.5281/zenodo.20588303 [15]. The archive contains `single_round_all_codes.csv`, `single_round_uf_decoder.csv`, `repeated_cycle_all_codes.csv`, `hardware_ablations.csv`, `hardware_three_condition_comparison.csv`, `minimax_lower_bounds.csv`, `bb72_certificate_report.json` and the SHA-256 manifest. All datasets are publicly accessible without restriction.

Code availability. The QEC Emulator v2.0.0 source code used to generate all results in this manuscript is publicly available under an MIT license at `github.com/Saleh-aldajah/qec-emulator` and permanently archived at doi:10.5281/zenodo.20588172 [16]. The package includes code-construction, decoder, noise-model, benchmark, and provenance modules; a REST API server; a command-line interface; a Dockerfile for containerized execution; and a `reproduce_all.py` script that reproduces all stochastic results in the declared environment (`ldpc==2.4.1`, `pymatching==2.4.0`, `numpy==2.4.4`). The reproducibility archive containing all primary CSV datasets, SHA-256 manifest, and runner scripts is archived at doi:10.5281/zenodo.20588303 [15].

Acknowledgements. The authors thank Joschka Roffe (developer of the `ldpc` package), Oscar Higgott (developer of PyMatching), and Craig Gidney (developer of Stim) for open-source software on which this work builds. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author contributions. Saleh H. AlDaajeh: Conceptualization, Methodology, Software, Formal analysis, Validation, Data curation, Writing – original draft, Writing

– review & editing. Saud Aldajah: Conceptualization, Methodology, Formal analysis, Validation, Writing – review & editing.

Competing interests. The authors declare no competing interests.

Additional information. Large-language-model assistance (Claude, Anthropic) was used for editorial refinement. The authors remain solely responsible for all scientific claims, mathematical proofs, experimental data, code, and conclusions. No LLM is listed as an author. Supplementary information accompanies this paper at doi:10.5281/zenodo.20588303.

References

[1] Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing **26**(5), 1484–1509 (1997) https://doi.org/10.1137/S0097539795293172 arXiv:quant-ph/9508027

[2] National Institute of Standards and Technology: FIPS 203: Module-lattice-based key-encapsulation mechanism standard. Federal Information Processing Standards Publication 203, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.203 . https://csrc.nist.gov/pubs/fips/203/final

[3] National Institute of Standards and Technology: FIPS 204: Module-lattice-based digital signature standard. Federal Information Processing Standards Publication 204, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.204 . https://csrc.nist.gov/pubs/fips/204/final

[4] National Institute of Standards and Technology: FIPS 205: Stateless hash-based digital signature standard. Federal Information Processing Standards Publication 205, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.205 . https://csrc.nist.gov/pubs/fips/205/final

[5] Panteleev, P., Kalachev, G.: Asymptotically good quantum and locally testable classical LDPC codes. In: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. STOC 2022, pp. 375–388. Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3519935.3520017

[6] Leverrier, A., Zémor, G.: Quantum Tanner codes. In: Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science. FOCS 2022, pp. 872–883. IEEE, ??? (2022). https://doi.org/10.1109/FOCS54457.2022.00117

[7] Bravyi, S., Cross, A.W., Gambetta, J.M., Maslov, D., Rall, P., Yoder, T.J.: High-threshold and low-overhead fault-tolerant quantum memory. Nature **627**, 778–782 (2024) https://doi.org/10.1038/s41586-024-07107-7 arXiv:2308.07915 [quant-ph]

[8] Shaw, S., Rengaswamy, N.: Networked realization of quantum LDPC codes. arXiv preprint arXiv:2604.25026 (2026) arXiv:2604.25026 [quant-ph]

[9] Roffe, J., White, D.R., Burton, S., Campbell, E.T.: Decoding across the quantum low-density parity-check code landscape. Physical Review Research **2**, 043423 (2020) https://doi.org/10.1103/PhysRevResearch.2.043423 arXiv:2005.07016 [quant-ph]

[10] Roffe, J.: LDPC: a Python package for low-density parity-check codes. https://roffe.eu/software/ldpc. Open-source software, MIT licence (2022)

[11] Higgott, O., Gidney, C.: Sparse blossom: correcting a million errors per core second with minimum-weight matching. Quantum **9**, 1600 (2025) https://doi.org/10.22331/q-2025-01-20-1600

[12] Higgott, O., Gidney, C.: PyMatching v2: sparse blossom back-end. https://pymatching.readthedocs.io (2025)

[13] Sion, M.: On general minimax theorems. Pacific Journal of Mathematics **8**(1), 171–176 (1958) https://doi.org/10.2140/pjm.1958.8.171

[14] Knill, E., Laflamme, R.: Theory of quantum error-correcting codes. Physical Review A **55**, 900–911 (1997) https://doi.org/10.1103/PhysRevA.55.900 arXiv:quant-ph/9604034

[15] AlDaajeh, S.H.: Adversarial Channel Equilibria for qLDPC Memories: Reproducibility Archive v16. Zenodo. SHA-256-anchored reproducibility archive, 42 files (2026). https://doi.org/10.5281/zenodo.20588303 . https://doi.org/10.5281/zenodo.20588303

[16] AlDaajeh, S.H.: QEC Emulator V2.0.0: Open-source qLDPC Benchmark Suite. https://doi.org/10.5281/zenodo.20588172 . GitHub: https://github.com/Saleh-aldajah/qec-emulator. MIT licence. https://doi.org/10.5281/zenodo.20588172