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.

Measuring the entanglement-breaking index of a qubit collision channel on a quantum processor: protocol and worked example

Eran Kopel

erankopel@tauex.tau.ac.il

Tel Aviv University

Research Article

Keywords: entanglement-breaking channels, entanglement-breaking index, collision models, quantum process tomography, mid-circuit reset, superconducting quantum processors

Posted Date: September 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11186783/v1

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

Additional Declarations: No competing interests reported.

Measuring the entanglement-breaking index of a qubit collision channel on a quantum processor: protocol and worked example

Eran Kopel[1*]

[1*]Tel Aviv University, Tel Aviv, Israel.

Corresponding author(s). E-mail(s): erankopel@tauex.tau.ac.il;

Abstract

The entanglement-breaking index of a quantum channel is the smallest number of self-compositions after which the channel destroys all entanglement with any reference system. A photonic experiment has fixed an index of 2; to our knowledge no larger index has been located by measuring the n-round channel across rounds. We give a protocol for doing so on a programmable processor, for any qubit channel realised as a collision model whose ancillas are prepared in states diagonal in a known basis, so that the bath polarisation is set exactly by weighting basis-state configurations in the analysis. Single-qubit process tomography of the n-round channel gives the smallest eigenvalue of its partially transposed Choi matrix, read against zero at a pre-registered five standard deviations; entanglement breaking survives further rounds, so the readings are monotone in n and the index follows. Comparing each direct measurement with the composition of the measured single-round map tests the repeated-interaction assumption. We show that a readable step needs margins on both sides, show how depolarising and reset errors move the readings, and give a power-based pilot rule. As a worked example we take one round of a three-qubit feedback loop, compiled to eight CZ gates on a heavy-hex path with mid-circuit reset, with certified ideal indices 3, 2 and 2 for three baths. The run is estimated to cost a few hundred seconds of processor time, and simulations under a device noise model, with the pilot's shot raise at low noise, recover all three indices.

Keywords: entanglement-breaking channels, entanglement-breaking index, collision models, quantum process tomography, mid-circuit reset, superconducting quantum processors

1 Introduction

A quantum channel is entanglement breaking (EB) if its output is separable from any reference system, whatever the input [1, 2]. A channel Φ that is not EB may become so under composition, and the smallest n for which it does,

$$n_{\mathrm{EB}}(\Phi) \;=\; \min\{\, n \geq 1 : \Phi^n \text{ is EB} \,\}, \tag{1}$$

is its entanglement-breaking index, or EB order [3, 4]. Which channels are eventually EB, and how fast, is well studied [5–7]. For a qubit channel the question is a finite computation: Φ^n is EB exactly when the partial transpose of its Choi matrix [8] is positive semidefinite [1, 9, 10].

The experimental record is thinner. Cuevas et al. [11] built a photonic channel, certified by Choi-state tomography that its two-fold concatenation $\Phi \circ \Phi$ is EB, and showed that an intermediate unitary F restores entanglement transmission, $\Phi \circ F \circ \Phi$ not being EB. By the ideal property recalled in Sec. 2.1, the second result also shows that Φ itself is not EB, so their data fix $n_{\mathrm{EB}}(\Phi) = 2$, and they relate their construction to the EB order. Certification that a single channel is not EB, the basic test for a quantum memory, has been developed in theory [12] and demonstrated with photons [13, 14]. Entanglement sudden death [15] is a different phenomenon: a state loses its entanglement at a finite time under continuous noise, while the local channel is not itself EB. To our knowledge no experiment has located a larger index, that is, scanned the number of rounds beyond two and found the step at which the n-round channel becomes EB, nor compared such a direct measurement with the composition of the measured single-round map.

Collision models are the natural setting [16–18]. A system meets a fresh ancilla, or a freshly reset one, in each round, so one round is one compiled circuit block and n rounds are its n-fold repetition. Collision models have run on photonic, superconducting and trapped-ion hardware [19–22], and multi-step process characterisation that tests the Markov assumption has been demonstrated on a superconducting processor [23].

This paper gives a protocol that measures n_{EB} directly, for any qubit channel realised as a collision round whose ancillas are prepared in a state diagonal in a known basis (Secs. 2 and 3). Four features make it practical. The bath polarisation is set exactly, without preparing mixed states, by weighting basis-state configurations in the analysis (Lemma 1). Single-qubit process tomography of the n-round channel needs no reference qubit. Each reading is pre-registered at five standard deviations, and the readings must be monotone in n, so an inconsistent pattern is itself a signal. And the comparison of the direct n-round channel with the composition of the measured single-round map tests the repeated-interaction assumption on which collision models rest. Section 4 gives the design rules: a step of the staircase is readable only if both of its sides are, and noise thickens one side while eating the other. Section 5 works the protocol through for one round of a three-qubit feedback loop [24, 25] on a heavy-hex superconducting processor, from certified targets to compilation, budget and simulated runs. Section 6 summarises what the same machinery certifies beyond static powers and why those targets are out of reach today, and Sec. 7 states the limitations.

This is version 3 of arXiv:2609.09350. Version 2 [26] is the full certified design study of the same loop; the present version rewrites it around the measurement protocol, corrects several of its statements (Sec. 7), and keeps its certificates in the ancillary files. The use of generative AI in preparing this version is documented in Appendix C.

2 Setting

2.1 Qubit channels, the Choi test, and monotone readings

Write a qubit state as $\rho = \frac{1}{2}(\mathbb{I} + \mathbf{r}\cdot\boldsymbol{\sigma})$. A qubit channel acts as $\mathbf{r} \mapsto A\mathbf{r} + \mathbf{c}$, with a real 3×3 matrix A and a real vector $\mathbf{c}$; we call $(A, \mathbf{c})$ the affine pair. Composition is matrix algebra: $\Phi_2 \circ \Phi_1$ has the pair $(A_2A_1,\ A_2\mathbf{c}_1 + \mathbf{c}_2)$. With the maximally entangled state $|\Omega\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ on a reference R and the system, the Choi matrix is $J(\Phi) = (\mathrm{id}\otimes\Phi)(|\Omega\rangle\langle\Omega|)$, and its partial transpose on R has the closed form

$$H \;=\; J(\Phi)^{T_R} \;=\; \frac{1}{4}\Big[\,\mathbb{I}\otimes\mathbb{I} + \sum_j c_j\,\mathbb{I}\otimes\sigma_j + \sum_{j,k} A_{jk}\,\sigma_k\otimes\sigma_j\Big]. \tag{2}$$

Φ is EB if and only if $J(\Phi)$ is separable [1], and for two qubits this holds if and only if $H \geq 0$ [9, 10]. We therefore read the single number $\lambda_{\min}(H)$: negative means the channel preserves some entanglement (NPT), non-negative means it is EB (PPT). A two-qubit partial transpose has at most one negative eigenvalue [27], so $-\lambda_{\min}$ is the negativity of the Choi state whenever it is positive. A cheaper condition on the pair alone, $\|A\|_*^2 + |\mathbf{c}|^2 \leq 1$ for every qubit EB channel, with $\|\cdot\|_*$ the nuclear norm, is necessary but not sufficient. It is the covariance-matrix criterion [28] read on the Choi state, and it is sharp [29]. It can screen candidate channels but cannot replace the test on H.

EB channels form a two-sided ideal: if Φ is EB then so are $\Psi\circ\Phi$ and $\Phi\circ\Psi$ for every channel Ψ [1]. Hence if Φ^n is EB so is Φ^m for every $m > n$, and if Φ^n is not EB neither is Φ^m for $m < n$. The readings of a power sequence are monotone: NPT up to round $n_{\mathrm{EB}} - 1$, PPT from round n_{EB} on. This is what lets an experiment read the integer from a few rounds, and it is also a test. Monotonicity survives rounds that differ but are memoryless, because the $(n+1)$-round channel is then $\Psi\circ\mathcal{E}_n$ for some channel Ψ. On hardware, however, each n is a separate circuit. A revival of entanglement at a later round therefore witnesses memory carried between rounds, a non-Markovian effect [30, 31], or a difference between circuits of different depth, such as drift or depth-dependent crosstalk.

2.2 Collision channels with diagonal ancillas

One round lets the system qubit M interact with m ancillas through a fixed unitary U, after which the ancillas are discarded (or reset and reused):

$$\Phi(\rho) \;=\; \mathrm{Tr}_{\mathrm{anc}}\big[\,U\big(\rho\otimes\tau^{\otimes m}\big)U^\dagger\big], \qquad \tau \;=\; \tau_S(p) \;=\; \tfrac{1}{2}\big(\mathbb{I} + p\,S\big), \tag{3}$$

where S is a Pauli operator, the bath axis, and $p \in [-1, 1]$ the polarisation. For a bath of splitting $\hbar\omega$ at temperature T along $S = Z$, $p = \tanh(\hbar\omega/2kT)$. The entries of $(A, \mathbf{c})$ are polynomials of degree m in p. The key property of diagonal ancillas is the following.

Lemma 1 (exact programming of the bath) *Let $|0_S\rangle, |1_S\rangle$ be the eigenbasis of S, with $\tau_S(p) = q_0|0_S\rangle\langle 0_S| + q_1|1_S\rangle\langle 1_S|$, $q_{0,1} = (1 \pm p)/2$. For a configuration $\mathbf{b} \in \{0, 1\}^{mn}$ of the ancilla basis states over n rounds, let $\mathcal{E}_{n,\mathbf{b}}$ be the n-round channel on M with those pure ancilla states. Then*

$$\Phi^n \;=\; \sum_{\mathbf{b}} w_p(\mathbf{b})\, \mathcal{E}_{n,\mathbf{b}}, \qquad w_p(\mathbf{b}) \;=\; \prod_{i=1}^{mn} q_{b_i}. \tag{4}$$

Proof With fresh (or perfectly reset) ancillas, the n-round channel is $\rho \mapsto \mathrm{Tr}_{\mathrm{anc}}[V(\rho \otimes \tau^{\otimes mn})V^\dagger]$ with V the product of the n round unitaries on disjoint ancilla sets. It is linear in the ancilla state, and $\tau^{\otimes mn} = \sum_{\mathbf{b}} w_p(\mathbf{b})\, |\mathbf{b}_S\rangle\langle \mathbf{b}_S|$. □

On hardware, each configuration is a basis-state preparation, one single-qubit rotation per ancilla. The expectation values of every $\mathcal{E}_{n,\mathbf{b}}$ are estimated, and the analysis combines them with the exact weights $w_p(\mathbf{b})$. The polarisation therefore carries no sampling error and needs no calibration beyond that of the basis-state preparation. At $p = \pm 1$ one configuration suffices. At $p = 0$ all 2^{mn} configurations enter with equal weight. In between, shots are allocated to configurations in proportion to their weights, with at least one replica each. Exact weighting needs all 2^{mn} configurations and is practical only for small mn (at most $2^6 = 64$ in the worked example). For many rounds one draws configurations at random with probabilities $w_p(\mathbf{b})$, which prepares $\tau^{\otimes mn}$ exactly on average but adds sampling error to the weights. The same reasoning covers any schedule of polarisations $(p_1, \ldots, p_n)$ across rounds, since $\Phi_{p_n} \circ \cdots \circ \Phi_{p_1}$ has the pair $A(p_n) \cdots A(p_1)$ and $\sum_j A(p_n) \cdots A(p_{j+1})\mathbf{c}(p_j)$.

2.3 Two noise mechanisms

Two imperfections change the channel in a way the analysis can follow exactly.

Depolarising noise on M. If each round is followed by $\rho \mapsto (1 - \lambda)\rho + \lambda\, \mathbb{I}/2$, the pair becomes $((1 - \lambda)A, (1 - \lambda)\mathbf{c})$, again a channel of the same kind. A single final depolarisation maps H to $(1 - \lambda)H + \lambda\, \mathbb{I}/4$ and so raises $\lambda_{\min}$, since $\lambda_{\min}(H) \leq 1/4$. Under per-round depolarisation the powers of the noisy round are still monotone in n, but it is not automatic that $\lambda_{\min}$ at fixed n increases with λ. We check this at every target.

Preparation and reset error of the ancillas. Suppose the reset leaves an ancilla in the wrong basis state with probability e, independently of the intended state. The prepared state is then $\tau_S((1-2e)p)$ instead of $\tau_S(p)$, so the channel is the ideal channel at the effective polarisation

$$p_{\mathrm{eff}} \;=\; (1 - 2e)\, p. \tag{5}$$

Two consequences are used below. The unpolarised bath $p = 0$ is immune to reset error, which makes it a control. At $p \neq 0$ the first-order change is $\Delta\lambda_{\min} \approx -2e\, p\, \partial\lambda_{\min}/\partial p$,

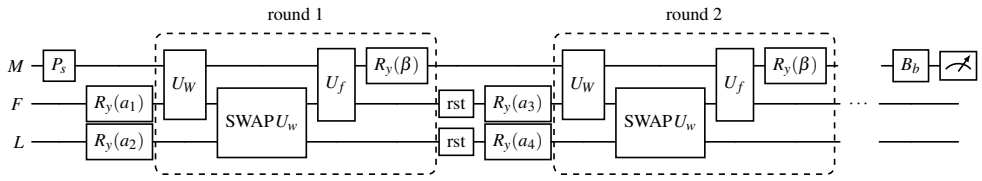


Fig. 1 Circuits of the protocol, shown for the worked example of Sec. 5. The message qubit M is prepared in a Pauli eigenstate P_s, passes n rounds, and is measured after a basis change B_b. The ancillas F and L receive a basis-state preparation $R_y(a_i)$ in every round and a mid-circuit reset (rst) before every round after the first. The angles are $a_i \in \{0, \pi\}$ for a Z bath and $a_i \in \{\pm\pi/2\}$ for an X bath, fixed by the configuration $\mathbf{b}$ and bound at run time. On a heavy-hex path the three blocks cost 2, 3 and 3 CZ gates (Sec. 5.3).

whose sign varies from channel to channel. The model assumes that a reset fails independently of the state before the reset. That state is correlated with M, so a state-dependent failure would couple the rounds. The composed-versus-direct test of Sec. 3 could detect it.

3 The protocol

The protocol fixes, before any data are taken, a target channel (the circuit U, the bath axis S and polarisation p), the set of rounds $\mathcal{N}$ to be measured, the shots per setting, and the decision threshold. It then runs six steps.

1. Circuits.

For every $n \in \mathcal{N}$ and every configuration $\mathbf{b}$ with its shot allocation, prepare M in one of the six Pauli eigenstates $P_s \in \{z_\pm, x_\pm, y_\pm\}$, apply n compiled rounds, with the ancillas reset and re-prepared between rounds, and measure M in the x, y or z basis (Fig. 1). This gives 18 settings per configuration. Each job also contains a readout calibration of M (prepare $|0\rangle$ and $|1\rangle$, measure). The pilot job also contains a reset check of each ancilla qubit: prepare $|1\rangle$, reset, measure. The fraction of outcomes 1 includes the ancilla readout error. The pilot therefore also calibrates the readout of each ancilla and corrects the fraction for it.

2. Estimation.

Invert the readout of M with the calibration. For each setting combine the configurations with the exact weights of Lemma 1. This gives the output Bloch vectors $\mathbf{r}_s$ for the six inputs, and linear inversion [32] gives

$$\mathbf{c}_n = \tfrac{1}{2}(\mathbf{r}_{z+} + \mathbf{r}_{z-}), \qquad A_n \mathbf{e}_a = \tfrac{1}{2}(\mathbf{r}_{a+} - \mathbf{r}_{a-}), \quad a \in \{x, y, z\}, \tag{6}$$

and then $\lambda_{\min}$ of Eq. (2). A parametric bootstrap [33] gives its standard deviation σ. It resamples every count, including the readout calibration, as binomial with the observed frequency, and uses 1000 replicas. Linear inversion can return a pair slightly outside the set of channels. That does no harm here, because $\lambda_{\min}$ is a statistic, not a state.

3. Readings.

Each block (S, p, n) is read as NPT if $\lambda_{\min} \leq -5\sigma$, PPT if $\lambda_{\min} \geq 5\sigma$, and unresolved otherwise. The threshold is fixed in advance.

4. Index.

By the ideal property (Sec. 2.1) the index is k when round $k-1$ reads NPT and round k reads PPT. Rounds below an NPT reading and above a PPT reading need not be measured. If a round between the last NPT and the first PPT is unresolved, the result is the bracket "k to k'". A PPT reading below an NPT reading is inconsistent with a power sequence. It stops the analysis and is reported as evidence of memory between rounds or of circuit-to-circuit differences (Sec. 2.1).

5. Composed versus direct.

Round 1 starts from freshly initialised ancillas and later rounds start from reset ones. The prediction for n rounds is therefore the measured single-round map $\hat{\Phi}_1$, with pair $(A_1, \mathbf{c}_1)$, followed by $n-1$ copies of the shifted map $\tilde{\Phi}_1$, with pair $(A_1 + \Delta A, \mathbf{c}_1 + \Delta\mathbf{c})$. Here $(\Delta A, \Delta\mathbf{c})$ is the model change of the single-round pair from p to p_{eff} at the measured reset error. For each $n \geq 2$ compute

$$D_n \;=\; \lambda_{\min}\big[\text{direct } n\text{-round estimate}\big] \;-\; \lambda_{\min}\big[\tilde{\Phi}_1^{\,n-1} \circ \hat{\Phi}_1\big], \tag{7}$$

with the standard deviation from the two bootstraps. The pre-registered pass criterion is $|D_n| < 3\sigma_D$ for every block. A failure means that the n-round channel is not the composition of single rounds as modelled. The possible causes are memory carried across the reset, a reset failure that depends on the state, heating by repeated resets, crosstalk that grows with depth, or drift. A repeat of selected blocks later in the session separates drift from the rest. In either case the index readings stand, because they come from the direct measurement.

6. Pilot and shots.

With 18 settings, a readout error near one per cent, and S shots per setting, the tomographic standard deviation of $\lambda_{\min}$ is close to

$$\sigma \;\approx\; \kappa/\sqrt{S}, \qquad \kappa \approx 0.24 \text{ to } 0.33, \tag{8}$$

across all blocks of the worked example, noisy or ideal. Reading a margin m at 5σ with 95 per cent power needs an expected $|z| \geq 5 + 1.645$, hence

$$S_{\text{req}} \;\approx\; \big(6.645\,\kappa/m\big)^2 \;\approx\; 3.5/m^2 \quad (\kappa = 0.28). \tag{9}$$

A margin of 0.03 costs about 4×10^3 shots per setting, and a margin of 0.005 about 1.4×10^5. Because the device's margins are not known in advance, a short pilot measures the single-round maps first, together with the readout and reset errors. It predicts each registered reading at $n \geq 2$ from the measured single-round map. Round 1 uses

the map as measured. Later rounds use it shifted by the model change from p to p_{eff} of Eq. (5), since they start from reset ancillas. The pilot then applies a fixed rule:

- a block whose reading is not secured at the registered shots is raised to the smallest power of two that secures it, up to a cap. A non-core reading that the prediction cannot secure even at the cap is raised to the cap anyway, because the prediction inherits the tomographic error of the single-round map, which is of order σ_D;
- the run goes ahead (GO) if every core reading is secured and the raised run fits the budget;
- otherwise (NO-GO) the pilot is repeated on another qubit layout. Targets and thresholds are never changed after data have been seen.

4 Design rules

4.1 A step has two sides

To read $n_{\text{EB}} = k$ an experiment must resolve round $k-1$ as NPT and round k as PPT. The decision margin of a step is therefore the pair

$$m_- = -\lambda_{\min}(H_{k-1}), \qquad m_+ = \lambda_{\min}(H_k), \tag{10}$$

and the shot cost is set by the smaller of the two through Eq. (9). Version 2 of this record [26] listed m_-, with only a sign certificate for the PPT side. Table 1 adds the certified m_+ for its twelve decision points and prices both sides. Only two points have both sides readable below 10^4 shots per setting. Seven need 10^5 to 7×10^6 on the PPT side, even on ideal hardware. At every listed point the PPT side is the thinner one. The primary target of version 2, $g = 0.85$ at $p = 0.43$, is the most lopsided: $m_- = 0.109$ but $m_+ = 1.06 \times 10^{-3}$.

4.2 Noise thickens one side and thins the other

At every target of the worked example, per-round depolarising noise and reset error raise $\lambda_{\min}$ (checked in float64 on a grid of $\lambda \le 0.12$ and $e \le 0.1$). The NPT side of a step is eaten, and the PPT side thickens. The two sides of Table 1 therefore play different roles on hardware. A target should have a large ideal m_-, because that side is fragile. It can live with a thin ideal m_+, as long as the expected noise or a raised shot count covers it. The contrast between two index-3 targets makes the point. At $g = 1$ on the Z bath both ideal sides are comfortable (0.045 and 0.037), but in the model of Sec. 4.3 the NPT side survives only a per-round contraction up to 0.025 (at 4096 shots and $e = 0$). At $g = 1.2$ on the X bath the ideal PPT side is thin (0.0056), but the NPT side survives up to 0.07, and noise thickens the thin side.

4.3 Tolerance of the worked example

Table 2 gives, for the readings of the worked example, the range of per-round depolarising contraction λ over which each reading holds with 95 per cent power at 5σ. The model includes the reset error e of Eq. (5), readout error 0.01 and a one-time

Table 1 Two-sided decision margins of the staircase targets of version 2 [26] (coupling family of Sec. 5.1; grid polarisations as stored in binary64) and of the worked example (last three rows). $m_{\mp}$ are certified (two-sided enclosures of relative width about 10^{-9}; values rounded as shown). $S_{\mp}$ are the shots per setting needed to read each side at 5σ with 95 per cent power on the ideal channel (float64 tomography simulation, readout error 0.01).

g	bath	p	k	m_-	m_+	S_-	S_+
0.85	Z	0.430	3	0.1086	0.00106	280	3 490 760
0.85	Z	1	4	0.0414	0.01480	1 830	13 002
0.70	Z	0.747	5	0.0519	0.00452	1 320	185 635
0.70	Z	1	6	0.0328	0.00177	2 917	922 306
0.55	Z	0.683	8	0.0342	0.00073	3 265	7 022 614
0.55	Z	0.873	9	0.0246	0.00188	5 832	968 221
1	Z	1	3	0.0452	0.03709	1 653	2 265
1	X	0.873	3	0.0953	0.00241	366	542 934
1	X	1	4	0.0563	0.00443	952	142 784
1	Y	0.620	3	0.0883	0.00672	448	70 746
1	Y	0.747	4	0.0315	0.02201	3 097	5 777
1	Y	0.873	5	0.0249	0.00565	4 213	67 971
1.2	X	1	3	0.0721	0.00561	569	89 746
1.2	Z	1	2	0.1146	0.03623	264	2 345
1.2	any	0	2	0.0925	0.09303	436	495

Table 2 Noise tolerance of the readings of the worked example (planning model, float64). Entries give the range of per-round depolarising contraction λ (scanned from 0 to 0.12 in steps of 0.005) over which the reading holds with 95 per cent power at 5σ, for ancilla reset error e. "All" means the whole scanned range. The model applies e to every round, including the first. That is conservative for the NPT rows and optimistic for the PPT rows.

reading (shots per setting)	$e = 0$	$e = 0.01$	$e = 0.02$	$e = 0.03$
X, $p = 1$, $n = 2$ NPT (4096)	≤ 0.070	≤ 0.050	≤ 0.035	≤ 0.020
X, $p = 1$, $n = 3$ PPT (4096)	≥ 0.035	≥ 0.020	all	all
X, $p = 1$, $n = 3$ PPT (16384)	≥ 0.010	all	all	all
Z, $p = 1$, $n = 1$ NPT and $n = 2$ PPT (4096)	all	all	all	all
$p = 0$, $n = 1$ NPT and $n = 2$ PPT (4096)	all	all	all	all
$g = 0.85$, Z, $p = 1$, $n = 3$ NPT (16384)	≤ 0.030	≤ 0.025	≤ 0.020	≤ 0.020
$g = 0.85$, Z, $p = 1$, $n = 4$ PPT (4096)	≥ 0.015	≥ 0.010	≥ 0.005	all
$g = 0.85$, Z, $p = 0.43$, $n = 3$ PPT (8192)	≥ 0.030	≥ 0.030	≥ 0.030	≥ 0.030

preparation-and-measurement contraction of 0.01. Two things stand out. The only fragile core reading is the X-bath NPT reading at $n = 2$, and reset error is what makes it fragile: the ideal X-bath index drops from 3 to 2 once p_{eff} falls below about 0.85. And the Z-bath and $p = 0$ readings hold over the whole scanned range. The model is isotropic and therefore pessimistic. In the stressed simulation of Sec. 5.5, with a fitted contraction of 0.052 and reset errors near 0.01, the X-bath $n = 2$ reading still came out NPT at -10.7σ. The pilot rule for this reason works from measured single-round maps, not from this table.

The last three rows show why the $g = 0.85$ staircase (3 at $p = 0.43$, 4 at $p = 1$) is not a good first target. With perfect resets its $p = 1$ readings hold for λ between about 0.015 and 0.030, and its $p = 0.43$ PPT reading needs $\lambda \geq 0.030$: the whole staircase is readable only in a narrow window near $\lambda \approx 0.03$. With $e \geq 0.01$ it is readable at no noise level.

5 Worked example

5.1 The round

The channel is one round of a three-qubit loop in which the message qubit M meets two ancillas, F and L ($m = 2$), both prepared in $\tau_S(p)$ [24]. With $R_y(a) = e^{-iaY/2}$ and registers ordered (M, F, L),

$$U \;=\; \big(R_y(\beta) \otimes \mathbb{I} \otimes \mathbb{I}\big)\, U_f(\varphi)\, U_w(\kappa)\, U_W(\theta), \tag{11}$$

where

$$\begin{aligned} U_W(\theta) &= |0\rangle\langle 0|_M \otimes R_y(\pi - 2\theta)_F + |1\rangle\langle 1|_M \otimes R_y(2\theta)_F, \\ U_w(\kappa) &= \exp(-i\tfrac{\kappa}{2} Z_F \otimes Y_L), \\ U_f(\varphi) &= \exp(-i\varphi\, \mathrm{SWAP}_{ML}). \end{aligned}$$

U_W writes M into F by a controlled rotation, U_w couples F to L, and the partial swap U_f returns part of L to M. We use the one-parameter family

$$\theta = \tfrac{\pi}{4} + g\,(\theta_0 - \tfrac{\pi}{4}), \quad \varphi = g\varphi_0, \quad \kappa = g\kappa_0, \quad \beta = \beta_0, \tag{12}$$

with $(\theta_0, \varphi_0, \kappa_0, \beta_0) = (0.16345853,\, 0.20061939,\, 0.4323098,\, 0.23903823) \times \pi$, the certified reference point of Refs. [24, 25], and take $g = 1.2$. At the reference point ($g = 1$) with a Z bath the index is 3 at every polarisation $p \in [0, 1]$, certified over the continuum in Ref. [25]. That reference also studies the infinite-temperature floor of this loop and its thermal valleys.

5.2 Targets

Table 3 lists the certified ideal values of $\lambda_{\min}$ for three baths. The certification uses 256-bit ball arithmetic [34]. The upper bound is the Rayleigh quotient of a proposed eigenvector. The lower bound is Sylvester's criterion on $H - s\,\mathbb{I}$, which proves $\lambda_{\min} > s$ (Appendix A).

- The X bath at $p = 1$ has index 3 and is the headline reading.
- The Z bath at the same population has index 2. The contrast between the two is a measurable instance of the dependence of the index on bath direction discussed in version 2.
- The unpolarised bath has index 2 and is the reset-immune control.

Table 3 Certified ideal $\lambda_{\min}$ of the n-round channel at $g = 1.2$ (enclosures of relative width about 10^{-9}, rigorous but not tight; six digits shown).

bath	$n = 1$	$n = 2$	$n = 3$	$n = 4$	n_{EB}
X, $p = 1$	-0.221643	-0.072137	$+0.005611$	$+0.043112$	3
Z, $p = 1$	-0.114649	$+0.036231$	$+0.091629$		2
$p = 0$	-0.092454	$+0.093030$	$+0.178514$		2

The core readings are X-bath $n = 2$ NPT and, for the other two baths, $n = 1$ NPT and $n = 2$ PPT. The X-bath $n = 3$ PPT reading completes the headline, and X-bath $n = 4$ is a backstop that brackets the index if $n = 3$ is unresolved.

5.3 Compilation

The three two-qubit blocks have analytic local-equivalence classes [35, 36]. U_W and U_w each have a single nonzero canonical coordinate, $|\pi/4 - \theta|$ and $\kappa/2$ (reduced to $\pi/2 - \kappa/2$ when $\kappa/2 > \pi/4$, as at $g = 1.2$), and need exactly two CNOTs. The partial swap has three equal nonzero coordinates $\varphi/2$ and needs three CNOTs [37]. A round therefore costs seven CNOTs with all-to-all connectivity, or five native arbitrary-angle two-qubit gates where these exist.

On a heavy-hex lattice [38] the three qubits sit on a path M–F–L with F in the middle, and U_f acts on the two ends. Routing it would cost a SWAP. The SWAP can instead be merged into the preceding block: $\mathrm{SWAP}_{FL}\, U_w$ is a three-CNOT gate. After it the state of L sits on the middle qubit, and U_f acts on two neighbours. The compiled round is $\mathrm{SWAP}_{FL} \cdot U$: M never moves, and the swap of the two ancillas is irrelevant because both are reset before the next round. The cost is $2 + 3 + 3 = 8$ CZ gates per round. Version 2 stated ten, which counts the SWAP separately.

We verified the compiled circuits exactly in noiseless simulation (Qiskit [39]), resets included. For n up to 3 the channel on M agrees with the model to about 10^{-16}. Transpiled for the heavy-hex device ibm_kingston (156-qubit Heron r2), each round costs eight CZ gates. On the path $(21, 22, 23)$, before echo padding, the three-round circuit has depth 113 with four mid-circuit resets. An X–X echo [40] is padded on M in idle windows, mainly during the resets. The padded, parametric circuits still reproduce the exact channel to about 10^{-15}. Mid-circuit reset [41] keeps the qubit count at three for any n. Fresh ancillas would cost $2n + 1$ qubits.

5.4 Budget

The run has three jobs:

- a pilot J1: $n = 1$ for the three baths, plus readout and reset checks;
- the main job J2: $n = 2, 3, 4$ for the X bath and $n = 2, 3$ for the other two;
- a repeat J3 of the X-bath $n = 2, 3$ blocks later in the same session.

All blocks of J1 to J3 start at 4096 shots per setting (Table 4).

Table 4 Jobs of the worked example at the registered shots (4096 per setting, except in J4, which uses 16384 for $p = 1$, $n = 3$ and 8192 for $p = 0.43$, $n = 3$). A binding is one circuit with one set of ancilla angles. The QPU seconds are estimates from the cost model of Sec. 5.4.

job	content	bindings	shots	QPU s
J1 pilot	$n = 1$, three baths; readout of M, F, L; reset check	116	253 952	72
J2 main	X: $n = 2, 3, 4$; Z and $p = 0$: $n = 2, 3$; readout	1 532	524 288	157
J3 repeat	X: $n = 2, 3$; readout	38	155 648	44
J4 optional	$g = 0.85$, Z: $p = 1$, $n = 3, 4$; $p = 0.43$, $n = 2, 3$	3 890	623 360	199

Table 5 Simulated runs (FakeKingston noise model, path $(140, 141, 142)$). The pilot columns give the fitted single-round contraction η, the measured reset errors of F and L, and the pilot decision. The index columns give the readings at 5σ. The reset errors are corrected for the ancilla readout. The last column is the largest $|D_n|/\sigma_D$ over all $n \geq 2$ blocks. The stressed scenario ran J1 and J2 only.

scenario	η	e_F, e_L	pilot	X, $p = 1$	Z, $p = 1$	$p = 0$	max $\lvert D_n \rvert$
nominal, 4096 shots	-0.001	0.005, 0.000	GO, raise	3 to 4	2	2	1.4σ
nominal, raised	-0.001	0.005, 0.000	GO	3	2	2	1.9σ
stressed	0.052	0.008, 0.012	GO	3	2	2	2.1σ

The processor time is estimated with a linear cost model, fitted to the author's earlier jobs on the same device: a per-binding overhead of 5.7 ms plus 0.27 ms per shot, with a 5 per cent margin. The three jobs cost 273 s, or 397 s with the raise the pilot asked for in the nominal simulation. That is a few hundred seconds of processor time. An optional job J4 for the $g = 0.85$ staircase costs a further 199 s.

5.5 Simulated runs

We ran the full pipeline on the Qiskit noise model of ibm_kingston (FakeKingston, qiskit-ibm-runtime 0.49), on the path $(140, 141, 142)$. That path ranked first under a per-round cost proxy computed from a calibration snapshot of the device. Two scenarios, meant to fall on either side of the device's actual noise, test the rules:

- *nominal*: the noise model as shipped;
- *stressed*: in addition, two-qubit depolarising noise of 0.004 on every CZ of the path and a bit flip of probability 0.01 after every reset.

Each scenario passed through the same pilot rule and analysis as a hardware run would (Table 5, Fig. 2).

Nominal. The pilot found a contraction consistent with zero and predicted the X-bath $n = 3$ PPT reading at only 1.6σ at 4096 shots, too thin to secure even at the cap. As the rule prescribes for a non-core reading, it raised that block and its repeat to the cap of 16384. That brings J2 to 219 s and J3 to 106 s, 397 s in all. At the registered shots the $n = 3$ block read $+3.0\sigma$ (unresolved), and the index was the bracket "3 to 4". With the raise it read $+5.4\sigma$, the repeat read $+7.9\sigma$, and the index was 3. The

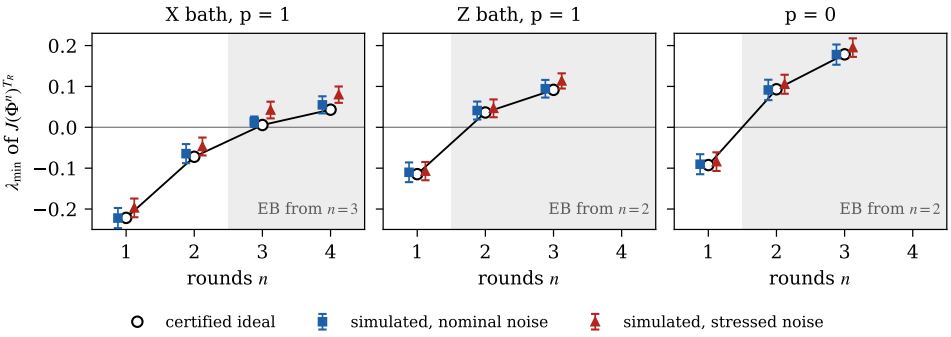


Fig. 2 Certified ideal $\lambda_{\min}$ (open circles) and simulated readings with $\pm 5\sigma$ bars, for nominal noise with the pilot's raise (squares) and stressed noise (triangles). The shaded region is where the ideal channel is EB. A reading counts as PPT or NPT when its whole bar lies above or below zero.

measured margin, 0.014, was twice the pilot's prediction of 0.006, a difference within the prediction's own uncertainty. The weakest reading of the headline sits at low noise, as Sec. 4.2 predicts.

Stressed. The fitted contraction was 0.052. The reset errors were 0.008 and 0.012 after the readout correction, against 0.024 and 0.016 as read. No raise was needed. The X-bath readings were -10.7σ at $n = 2$ and $+10.3\sigma$ at $n = 3$, giving index 3. Noise had thickened the thin side and eaten part of the robust one, and both still read.

Composed versus direct. In both scenarios no $|D_n|$ exceeded 2.1σ. The simulated rounds are memoryless and differ only by the modelled reset error, so this is the expected result. The reset shift matters: without it, the stressed X-bath blocks at $n = 3$ and 4 would have given D_n of $+1.5\sigma$ and $+2.1\sigma$, against $+0.4\sigma$ and $+0.9\sigma$ with it, while at $n = 2$ the shift moved D_n from -0.6σ to -1.4σ. The uncertainty of the measured e is not included in σ_D. It adds about 0.002 to D_n here. The test is a null test here. On hardware it is a measurement.

The $g = 0.85$ staircase. Run under nominal noise, J4 read both NPT sides clearly (-11.9σ at $p = 1$, $n = 3$, and -21.6σ at $p = 0.43$, $n = 2$). Both PPT sides were unresolved ($+3.4\sigma$ and $+3.2\sigma$), so it returned only lower bounds on the two indices, as Table 2 predicts.

These simulations test the pipeline and the rules, not the hardware. What they show is that the protocol, run exactly as registered, recovers the certified indices under two noise levels chosen to fall on either side of what current devices show, in a model without coherent errors, and that the pilot identifies the one reading that needs more shots.

6 Beyond static powers

Because composition is exact for any schedule of polarisations (Sec. 2.2), the same machinery certifies the index under time-dependent baths. Version 2 [26] explores this for the loop of Sec. 5.1 and deposits the certificates, which are kept in the ancillary files of this version. We summarise, because it marks where the direct measurement stops being practical.

Thermal valleys. Some circuits break entanglement sooner at an intermediate bath polarisation than at either extreme; the first certified examples are in Ref. [25]. The gated random search of version 2, over 33 million circuits, found 10,628 such valleys, and none reaches a contrast of 0.05. The contrast has an interior maximum near index 90, at a certified 2.85×10^{-3}.

Control results. Exact inhomogeneous composition turns a valley into a switch. At the best valley, held inside its window at $n = 88$ rounds:

- a centred pulse of exactly eight rounds restores entanglement;
- a 53-round pulse beats the static protocol by a certified margin;
- a 1.2 per cent coherent tilt of the ancilla opens the switch at fixed population.

Cost. These signals lie between about 10^{-4} (the minimal eight-round pulse) and 10^{-2} (the tilted ancilla), mostly near 3×10^{-3}, over 80 to 110 rounds. By Eq. (9) a margin of 3×10^{-3} costs about 4×10^5 shots per setting for each of the 18 settings. Version 2 estimated about 3×10^6, from an assumed error model for two-qubit state tomography. At intermediate polarisation, exact weighting over 2^{176} configurations is out of the question, and configurations must be sampled. Within a per-round depolarising model the signals die at a per-round contraction near 10^{-4}. They are certified statements about the model and a map of what better hardware could reach, not near-term measurement targets.

7 Discussion and limitations

What is certified, and what is not. The target values of Tables 1 and 3 are certified statements about the ideal channel. The shot requirements, the tolerance table and the simulated runs are float64 planning estimates. A hardware run measures the index of the channel it implements. Depolarising-type noise and reset error move that index down or leave it unchanged at the targets of the worked example, but coherent errors can move $\lambda_{\min}$ either way. The certified ideal index is therefore a comparison value, not a pass criterion.

Scope of the statistics. The bootstrap covers shot noise and readout calibration, not drift or systematic preparation error. Drift is addressed by the in-session repeat and the composed-versus-direct test. Preparation and measurement errors of M that act as channels compose the channel with a noisy channel before and after it. By the ideal property they can turn an NPT reading into PPT but never the reverse, so, like the other noise, they bias the reading towards a smaller index. The readout inversion removes most of the measurement part. An over- or under-correction in that inversion is not a channel, however, and could push the reading either way. The reset check bounds the ancilla reset error from above unless the ancilla readout is calibrated too, as it is in the pilot of the worked example. Maximum-likelihood tomography, or a tomography method with certified confidence regions, could replace linear inversion without changing the rest of the protocol.

Generality. Nothing in Secs. 2 and 3 uses the particular loop. The protocol applies to any qubit collision channel with ancillas diagonal in a known basis, on any platform with mid-circuit reset, or with $mn + 1$ qubits without it. For a system of dimension

$d > 2$ the Choi matrix is $d^2 \times d^2$ and a positive partial transpose no longer implies separability. The PPT reading then bounds the index of the PPT property, and an EB test needs a separability criterion or witness.

Corrections to version 2. The following statements of version 2 are corrected here.

- The heavy-hex count is eight CZ gates per round, not ten (Sec. 5.3). The derived counts change accordingly: 72 CZ gates, not 90 CNOTs, for the deepest staircase target, and 704, not 880, for the 88-round valley.
- The decision margin of a step has two sides. At every listed target the side that version 2 did not list is the thinner one (Table 1). The shot range of version 2, about 2×10^3 to 4×10^4 per setting for the selected targets, came from its own calibration of the NPT side. In the present calibration the NPT sides need 280 to 5,832 shots per setting, and the PPT sides up to 7×10^6.
- Exact programming of p holds for fresh ancillas. With mid-circuit reuse a reset error shifts p to $(1 - 2e)p$ (Eq. (5)), so the reset does enter the ancilla state.
- Version 2 said that no experiment had measured an entanglement-breaking index. That is not correct as stated. The data of Ref. [11] fix an index of 2 for a photonic channel. What remains unmeasured is a larger index located by a direct scan over rounds.

The worked example is ready to run. It needs three qubits, eight CZ gates per round, four rounds at most, and an estimated few hundred seconds of processor time.

Acknowledgements. The computations used Arb through python-flint, NumPy, Qiskit and Qiskit Aer.

Declarations

Funding. No external funding was received for this work.

Competing interests. The author has no competing interests to declare.

Ethics approval and consent to participate. Not applicable.

Consent for publication. Not applicable.

Data availability. No experimental data were generated. The certificates, the simulated run records and analyses, and the complete search output of version 2 are in the ancillary files of arXiv:2609.09350 (Appendix B).

Materials availability. Not applicable.

Code availability. The circuit builders, run and analysis scripts, the certification scripts and the re-derivation script are in the ancillary files of arXiv:2609.09350. The certification library from which the primitives derive is archived at Zenodo [42].

Author contribution. E.K. is the sole author. He conceived and designed the study, carried out and checked the analysis with the assistance described in Appendix C, and wrote the manuscript.

Use of generative AI. See Appendix C.

Appendix A Certification

All certified numbers are two-sided enclosures in Arb ball arithmetic [34] at 256 bits, through python-flint. Circuit parameters and polarisations enter as exact binary64 numbers, so each certificate concerns the circuit as stored. The round unitary is built from closed-form sines and cosines, since $(Z \otimes Y)^2 = \mathrm{SWAP}^2 = \mathbb{I}$, and the affine pair by tracing out the ancillas. Composites follow exactly from the pair algebra of Sec. 2.1. For $\lambda_{\min}$ of the 4×4 matrix H:

- the upper bound is the Rayleigh quotient of a float64-proposed eigenvector, certified in ball arithmetic;
- the lower bound s comes from Sylvester's criterion on $H - s\,\mathbb{I}$ (all leading principal minors certified positive), with s backed off from the float64 value until the criterion passes;
- an enclosure that straddles zero counts as a failed sign certificate.

No floating-point eigensolver is trusted. The worked-example targets are produced by `certify_g120.py`, the two-sided margins of Table 1 by `two_sided_margins.py`, and every certified number of version 2 by `paperV_checks.py`.

Appendix B Deposited files

The ancillary directory of arXiv:2609.09350v3 keeps all files of version 2 unchanged: certificates, search chassis and output, schedule engine, and the re-derivation script `paperV_checks.py`. It adds a folder `staircase/` for this version:

- `staircase_circuits.py`: circuit builders and exactness checks;
- `run_staircase.py`: build, price, simulate or submit the jobs;
- `analyze_staircase.py`: estimation, bootstrap, readings, index, pilot rule and composed-versus-direct test;
- `choose_layout.py`: layout ranking;
- `tolerance.py`: Table 2;
- `certify_g120.py` and `two_sided_margins.py`: the certified targets and margins, with their outputs;
- the records, counts and analyses of every simulated run in Sec. 5.5, and `make_fig2.py` for Fig. 2;
- a `README.md` with the commands that reproduce every table and figure of this version.

Appendix C Methods note: use of generative AI

This version was prepared with Claude (Anthropic), used through the Claude desktop app between 15 and 26 September 2026 (model versions Claude Fable 5.1 and Claude Opus 5.5). It was used to:

- draft and revise the manuscript text;

- write and run the scripts of the worked example (circuit construction and verification, noisy simulations, noise-tolerance scans, analysis, and the new certification scripts);
- check the references against Crossref and arXiv.

The author designed the study, directed and checked all of this work, verified the results, and takes full responsibility for the content.

References

[1] Horodecki, M., Shor, P.W., Ruskai, M.B.: Entanglement breaking channels. Rev. Math. Phys. **15**, 629–641 (2003) https://doi.org/10.1142/S0129055X03001709

[2] Ruskai, M.B.: Qubit entanglement breaking channels. Rev. Math. Phys. **15**, 643–662 (2003) https://doi.org/10.1142/S0129055X03001710

[3] De Pasquale, A., Giovannetti, V.: Quantifying the noise of a quantum channel by noise addition. Phys. Rev. A **86**, 052302 (2012) https://doi.org/10.1103/PhysRevA.86.052302

[4] Lami, L., Giovannetti, V.: Entanglement-breaking indices. J. Math. Phys. **56**, 092201 (2015) https://doi.org/10.1063/1.4931482

[5] Rahaman, M., Jaques, S., Paulsen, V.I.: Eventually entanglement breaking maps. J. Math. Phys. **59**, 062201 (2018) https://doi.org/10.1063/1.5024385

[6] Christandl, M., Müller-Hermes, A., Wolf, M.M.: When do composed maps become entanglement breaking? Ann. Henri Poincaré **20**, 2295–2322 (2019) https://doi.org/10.1007/s00023-019-00774-7

[7] Hanson, E.P., Rouzé, C., Stilck França, D.: Eventually entanglement breaking Markovian dynamics: structure and characteristic times. Ann. Henri Poincaré **21**, 1517–1571 (2020) https://doi.org/10.1007/s00023-020-00906-4

[8] Choi, M.-D.: Completely positive linear maps on complex matrices. Linear Algebra Appl. **10**, 285–290 (1975) https://doi.org/10.1016/0024-3795(75)90075-0

[9] Peres, A.: Separability criterion for density matrices. Phys. Rev. Lett. **77**, 1413–1415 (1996) https://doi.org/10.1103/PhysRevLett.77.1413

[10] Horodecki, M., Horodecki, P., Horodecki, R.: Separability of mixed states: necessary and sufficient conditions. Phys. Lett. A **223**, 1–8 (1996) https://doi.org/10.1016/S0375-9601(96)00706-2

[11] Cuevas, Á., De Pasquale, A., Mari, A., Orieux, A., Duranti, S., Massaro, M., Di Carli, A., Roccia, E., Ferraz, J., Sciarrino, F., Mataloni, P., Giovannetti, V.: Amending entanglement-breaking channels via intermediate unitary operations. Phys. Rev. A **96**, 022322 (2017) https://doi.org/10.1103/PhysRevA.96.022322

[12] Rosset, D., Buscemi, F., Liang, Y.-C.: Resource theory of quantum memories and their faithful verification with minimal assumptions. Phys. Rev. X **8**, 021033 (2018) https://doi.org/10.1103/PhysRevX.8.021033

[13] Mao, Y., Zhen, Y.-Z., Liu, H., Zou, M., Tang, Q.-J., Zhang, S.-J., Wang, J., Liang, H., Zhang, W., Li, H., You, L., Wang, Z., Li, L., Liu, N.-L., Chen, K., Chen, T.-Y., Pan, J.-W.: Experimentally verified approach to nonentanglement-breaking channel certification. Phys. Rev. Lett. **124**, 010502 (2020) https://doi.org/10.1103/PhysRevLett.124.010502

[14] Graffitti, F., Pickston, A., Barrow, P., Proietti, M., Kundys, D., Rosset, D., Ringbauer, M., Fedrizzi, A.: Measurement-device-independent verification of quantum channels. Phys. Rev. Lett. **124**, 010503 (2020) https://doi.org/10.1103/PhysRevLett.124.010503

[15] Almeida, M.P., de Melo, F., Hor-Meyll, M., Salles, A., Walborn, S.P., Souto Ribeiro, P.H., Davidovich, L.: Environment-induced sudden death of entanglement. Science **316**, 579–582 (2007) https://doi.org/10.1126/science.1139892

[16] Scarani, V., Ziman, M., Štelmachovič, P., Gisin, N., Bužek, V.: Thermalizing quantum machines: dissipation and entanglement. Phys. Rev. Lett. **88**, 097905 (2002) https://doi.org/10.1103/PhysRevLett.88.097905

[17] Ciccarello, F., Lorenzo, S., Giovannetti, V., Palma, G.M.: Quantum collision models: open system dynamics from repeated interactions. Phys. Rep. **954**, 1–70 (2022) https://doi.org/10.1016/j.physrep.2022.01.001

[18] Campbell, S., Vacchini, B.: Collision models in open system dynamics: a versatile tool for deeper insights? EPL **133**, 60001 (2021) https://doi.org/10.1209/0295-5075/133/60001

[19] Cuevas, Á., Geraldi, A., Liorni, C., Bonavena, L.D., De Pasquale, A., Sciarrino, F., Giovannetti, V., Mataloni, P.: All-optical implementation of collision-based evolutions of open quantum systems. Sci. Rep. **9**, 3205 (2019) https://doi.org/10.1038/s41598-019-39832-9

[20] García-Pérez, G., Rossi, M.A.C., Maniscalco, S.: IBM Q Experience as a versatile experimental testbed for simulating open quantum systems. npj Quantum Inf. **6**, 1 (2020) https://doi.org/10.1038/s41534-019-0235-y

[21] Cattaneo, M., Rossi, M.A.C., García-Pérez, G., Zambrini, R., Maniscalco, S.: Quantum simulation of dissipative collective effects on noisy quantum computers. PRX Quantum **4**, 010324 (2023) https://doi.org/10.1103/PRXQuantum.4.010324

[22] Wang, Z., Teske, J.D., Bhardwaj, A., Takahashi, M.O., Yunoki, S.: Simulating the dynamics of Markovian quantum processes by quantum collision models on

quantum computers. arXiv:2606.27856 (2026)

[23] White, G.A.L., Hill, C.D., Pollock, F.A., Hollenberg, L.C.L., Modi, K.: Demonstration of non-Markovian process characterisation and control on a quantum processor. Nat. Commun. **11**, 6301 (2020) https://doi.org/10.1038/s41467-020-20113-3

[24] Kopel, E.: Certified coherent, informative, and non-entanglement-breaking fixed points of future-referential quantum feedback. arXiv:2608.13764 (2026)

[25] Kopel, E.: The quantum lifetime of a future-referential feedback loop: certified index, architecture floor, and thermal bonus. arXiv:2609.27908 (2026)

[26] Kopel, E.: Certified targets for measuring and controlling an entanglement-breaking index on programmable hardware. arXiv:2609.09350v2. Version 2 of this record, with the full valley, pulse and tilt analysis (2026)

[27] Sanpera, A., Tarrach, R., Vidal, G.: Local description of quantum inseparability. Phys. Rev. A **58**, 826–830 (1998) https://doi.org/10.1103/PhysRevA.58.826

[28] Gittsovich, O., Gühne, O., Hyllus, P., Eisert, J.: Unifying several separability conditions using the covariance matrix criterion. Phys. Rev. A **78**, 052319 (2008) https://doi.org/10.1103/PhysRevA.78.052319

[29] Kopel, E.: A sharp norm inequality for entanglement-breaking channels. arXiv:2609.27906 (2026)

[30] Rivas, Á., Huelga, S.F., Plenio, M.B.: Entanglement and non-Markovianity of quantum evolutions. Phys. Rev. Lett. **105**, 050403 (2010) https://doi.org/10.1103/PhysRevLett.105.050403

[31] Rivas, Á., Huelga, S.F., Plenio, M.B.: Quantum non-Markovianity: characterization, quantification and detection. Rep. Prog. Phys. **77**, 094001 (2014) https://doi.org/10.1088/0034-4885/77/9/094001

[32] Chuang, I.L., Nielsen, M.A.: Prescription for experimental determination of the dynamics of a quantum black box. J. Mod. Opt. **44**, 2455–2467 (1997) https://doi.org/10.1080/09500349708231894

[33] Efron, B., Tibshirani, R.J.: An Introduction to the Bootstrap. Chapman and Hall/CRC, New York (1994). https://doi.org/10.1201/9780429246593

[34] Johansson, F.: Arb: efficient arbitrary-precision midpoint-radius interval arithmetic. IEEE Trans. Comput. **66**, 1281–1292 (2017) https://doi.org/10.1109/TC.2017.2690633

[35] Makhlin, Y.: Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations. Quantum Inf. Process. **1**, 243–252 (2002)

https://doi.org/10.1023/A:1022144002391

[36] Zhang, J., Vala, J., Sastry, S., Whaley, K.B.: Geometric theory of nonlocal two-qubit operations. Phys. Rev. A **67**, 042313 (2003) https://doi.org/10.1103/PhysRevA.67.042313

[37] Vatan, F., Williams, C.: Optimal quantum circuits for general two-qubit gates. Phys. Rev. A **69**, 032315 (2004) https://doi.org/10.1103/PhysRevA.69.032315

[38] Chamberland, C., Zhu, G., Yoder, T.J., Hertzberg, J.B., Cross, A.W.: Topological and subsystem codes on low-degree graphs with flag qubits. Phys. Rev. X **10**, 011022 (2020) https://doi.org/10.1103/PhysRevX.10.011022

[39] Javadi-Abhari, A., Treinish, M., Krsulich, K., Wood, C.J., Lishman, J., Gacon, J., et al.: Quantum computing with Qiskit. arXiv:2405.08810 (2024)

[40] Viola, L., Knill, E., Lloyd, S.: Dynamical decoupling of open quantum systems. Phys. Rev. Lett. **82**, 2417–2421 (1999) https://doi.org/10.1103/PhysRevLett.82.2417

[41] Córcoles, A.D., Takita, M., Inoue, K., Lekuch, S., Minev, Z.K., Chow, J.M., Gambetta, J.M.: Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits. Phys. Rev. Lett. **127**, 100501 (2021) https://doi.org/10.1103/PhysRevLett.127.100501

[42] Kopel, E.: quantum-future-feedback: code and machine-verified certificates for future-referential quantum feedback, version 1.2.0. Zenodo (2026). https://doi.org/10.5281/zenodo.22046245