

The Capacity–Backaction Frontier

An Information-Theoretic Coordinate for Cryogenic Quantum Error Correction

Aernoud Dekker

June 2026

Draft for peer review. DOI: 10.17605/OSF.IO/G293X.

Abstract

Superconducting quantum error correction already operates on a physical frontier: syndrome information must be extracted fast enough to protect a logical state, while the measurements, readout photons, amplification, decoding, and feedback activity that provide that information can themselves add heat, leakage, quasiparticles, photon shot-noise dephasing, crosstalk, latency, and correlated errors. This paper does not introduce a new backaction mechanism. It defines an operational aggregate coordinate for comparing useful syndrome capacity against the physical instability induced by obtaining and using it.

We define the useful online correction capacity C_{eff} as the mutual-information rate by which syndrome data are measured, decoded, and incorporated into a timely Pauli-frame or feedback update. We define the effective instability rate h_{eff} as the aggregate rate of physical and estimator uncertainty that must be controlled, including a backaction term $g_N(C_{\text{eff}})$ induced by information extraction:

$$h_{\text{eff}}(N, C_{\text{eff}}) = h_0(N) + h_{\text{corr}}(N) + g_N(C_{\text{eff}}).$$

The proposed dimensionless coordinate is

$$\rho_{\text{CB}} = \frac{\epsilon_{\text{QEC}} C_{\text{eff}} \ln 2}{h_{\text{eff}}(N, C_{\text{eff}})},$$

where ϵ_{QEC} denotes the code-specific margin converting tracking capacity into tolerated logical protection. The measured object is the Pareto frontier between useful information flow and induced physical instability. If ρ_{CB} improves monotonically with added capacity, increasing timely syndrome throughput moves the operating point favorably. If $g'_N(C_{\text{eff}})$ becomes large enough that extra measurement/control activity increases h_{eff} as fast as or faster than useful $C_{\text{eff}} \ln 2$, logical performance can saturate or degrade despite greater nominal throughput.

The contribution is organizational and predictive. QEC teams already optimize capacity–backaction tradeoffs component by component; this paper proposes an aggregate coordinate that compares readout cadence, decoder latency, measurement strength, cryogenic control, correlated error generation, and code performance on the same rate-accounting axis. We specify estimators, scaling questions, experimental designs, and falsification criteria for testing whether this coordinate improves prediction over raw engineering variables alone.

Scope. The analysis assumes the standard QEC setting in which readout backaction, leakage, quasiparticles, crosstalk, timing, and decoder latency are component-level engineering variables. The proposed aggregate description compresses those variables into a measurable capacity–backaction frontier and asks how that frontier scales with processor size, code distance, readout multiplexing, wiring density, cryogenic control, and correlated error channels.

Relation to finite-capacity control. The formalism is motivated by the finite-capacity control structure developed in Ref. [1]: a controller can only stabilize or track degrees of freedom at the rate supplied by its useful information channel. In the QEC setting studied here, that principle is applied conservatively to syndrome extraction, decoding, feedback, and the physical disturbance induced by obtaining timely information.

Contents

1	Introduction	5
1.1	Structure of the Paper	5
2	Known Engineering Frontier	6
2.1	Syndrome Extraction Is Physical	6
2.2	Backaction Channels Are Standard	6
2.3	Why Aggregate Coordinates Are Useful	6
3	Operational Variables	7
3.1	Useful Correction Capacity	7
3.2	Effective Instability Rate	7
3.3	Capacity–Backaction Ratio	8
4	Frontier Model	8
4.1	A Minimal Rate Model	8
4.2	Operating-Point Condition	9
4.3	Backaction Functional Forms	9
4.4	Relation to Data-Rate Language	10
5	Scaling Question	10
5.1	Local Frontier Versus Scaling Frontier	10
5.2	Exponent Accounting	10
5.3	Why Below-Threshold QEC Matters	10
5.4	Architecture as Frontier Engineering	11
6	Experimental Program	11
6.1	Retrospective Audit	11
6.2	Prospective Capacity–Backaction Sweep	12
6.3	Decision Outcomes	13
7	Falsification and Weakening Criteria	13
8	Discussion	14
8.1	Aggregate Rate Coordinates	14

8.2	Relation to Standard Error Budgets	14
8.3	Relation to Standard Quantum Dynamics	14
8.4	Implications for Architectures	14
8.5	Implications for Reporting	15
9	Conclusion	15
A	Appendix: Minimal Preregistered Analysis Template	16
B	Appendix: Glossary	16
	References	17

1 Introduction

Fault-tolerant quantum computing requires a continuous exchange of information between a fragile quantum substrate and a classical control system. In superconducting processors this exchange is highly physical: syndrome extraction uses microwave readout resonators and amplifiers; reset, feedback, and Pauli-frame tracking require time-critical digital processing; qubits sit at millikelvin temperatures with limited cooling power; and additional measurement or control activity can itself inject noise into the device.

The resulting tradeoff is well known. If syndrome extraction is too slow, the logical state accumulates uncorrected T_1 , T_2 , leakage, crosstalk, and stochastic errors. If extraction is too aggressive, measurement photons, resonator depopulation tails, amplifier activity, control-line power, quasiparticles, chip heating, microwave crosstalk, and readout-induced dephasing can increase the physical error rate. Practical QEC therefore requires an operating point: enough useful information per unit time to support the decoder, but not so much physical information-extraction activity that the quantum substrate is damaged faster than it is protected.

This paper names that operating surface the **capacity–backaction frontier**: a compact coordinate system for the relation between useful information extraction and induced physical instability. The corresponding scaling question is:

As cryogenic quantum processors scale, does useful online correction capacity grow faster than the physical instability induced and revealed by obtaining that capacity?

Below-threshold operation establishes that the current operating point is favorable for the tested distances, circuits, noise model, and decoder. The frontier question concerns how the operating point moves when the syndrome volume, readout multiplexing, wiring density, cycle-time pressure, cryogenic controller activity, and correlated error channels scale.

Operational Statements

- Useful syndrome information is a physical resource. Its extraction and timely use can be quantified as an effective capacity C_{eff} .
- In superconducting QEC, increasing C_{eff} can also increase physical instability through a backaction function $g_N(C_{\text{eff}})$.
- Logical performance can be analyzed on the frontier between useful capacity and induced instability, rather than against nominal readout rate or raw controller power alone.
- Capacity–backaction measurements on cryogenic qubits are interpreted here within standard open-system QEC.

1.1 Structure of the Paper

Section 2 reviews the superconducting-QEC facts that motivate the frontier. Section 3 defines C_{eff} , h_{eff} , and the backaction function. Section 4 develops the frontier model and its local optimality condition. Section 5 states the scaling question. Section 6 proposes retrospective and prospective measurements. Section 7 lists falsification and weakening criteria. Section 8 discusses implications and limitations.

2 Known Engineering Frontier

The capacity–backaction frontier is built from standard facts about superconducting processors.

2.1 Syndrome Extraction Is Physical

Surface-code and related QEC experiments repeatedly measure stabilizers or detector events. Each round requires controlled gates, measurement pulses, resonator readout, amplification, digitization, classification, decoding, and frame update. The useful information is classical, but its acquisition is not abstract: photons and control pulses enter a cryogenic device whose qubits are sensitive to residual excitation and correlated disturbance.

Google’s 2023 surface-code scaling experiment demonstrated that increasing code size can improve logical performance when the physical error density is low enough [2]. Google’s Willow experiment reported below-threshold surface-code memories with real-time decoding [3]. These results constrain the frontier: current superconducting systems can sit on its favorable side for tested distances, circuits, and hardware.

2.2 Backaction Channels Are Standard

The physical channels entering $g_N(C_{\text{eff}})$ are not speculative. Examples include:

- **Photon shot-noise and residual-photon dephasing:** photons in a coupled readout mode shift qubit frequency and dephase the qubit [4–6].
- **Quasiparticle poisoning:** nonequilibrium quasiparticles produce relaxation and correlated errors in superconducting devices [7–9].
- **Radiation and burst errors:** ionizing radiation and high-energy impacts can create quasiparticle bursts and chip-wide correlated failure modes [10–12].
- **Readout crosstalk and leakage:** measurement and reset operations can affect neighboring qubits and later cycles if photons, leakage population, or resonator excitations persist.
- **Decoder and latency constraints:** real-time decoders must return useful decisions within the time budget of the code; late information is not equivalent to timely correction capacity [13].

The aggregate question is how these component-level channels compose into an empirical frontier between useful information rate and induced instability.

2.3 Why Aggregate Coordinates Are Useful

Component-level error budgets are necessary for engineering. They answer where an error came from and what knob fixes it. Aggregate coordinates answer a different question: how close is the whole system to a rate boundary where extra measurement/control effort stops buying logical protection?

An aggregate coordinate is useful when:

- different backaction mechanisms share a common cause, such as increased measurement cadence or readout power;
- different architectures trade information rate against physical disturbance in incomparable ways;
- the scaling question concerns exponents and frontiers rather than one component;
- published data report logical performance, cycle time, and error budgets but not a single capacity-like variable.

3 Operational Variables

3.1 Useful Correction Capacity

Definition 3.1 (Useful online correction capacity). The useful online correction capacity C_{eff} is the rate, in bits/s, at which information about the relevant error state is acquired, decoded, and incorporated into a timely Pauli-frame or feedback decision.

A practical estimator is

$$C_{\text{eff}} = \sum_j r_j I(E_j; S_j) a_j \ell_j, \quad (1)$$

where j indexes syndrome or detector channels, r_j is the accepted update rate, $I(E_j; S_j)$ is the mutual information between the relevant error variable E_j and reported syndrome feature S_j , a_j is the accepted fraction after rejection/filtering/classifier confidence gates, and $\ell_j \in [0, 1]$ is a timeliness factor representing whether the information reaches the correction policy before it loses operational value.

This definition excludes:

- raw ADC bandwidth that is not decoded;
- measurement data that wait in a queue past the cycle budget;
- dummy readout power that produces no useful syndrome information;
- offline information used only after the run is complete;
- controller clock speed not allocated to the QEC task.

For a surface-code memory, C_{eff} can be estimated per round and converted to bits/s by dividing by T_{cyc} . A simple upper bound is the entropy rate of accepted detector events, but the more relevant quantity is mutual information with the true error process as seen by the decoder. In experiments, the true error process is latent; simulation-calibrated error models, detector likelihoods, and decoder posterior entropy can be used as proxies.

3.2 Effective Instability Rate

Definition 3.2 (Effective instability rate). The effective instability rate h_{eff} is the rate, in nats/s, at which uncertainty relevant to logical protection is generated by the processor, the control stack, and the information-extraction process.

We decompose it as

$$h_{\text{eff}}(N, d, u) = h_0(N, d) + h_{\text{corr}}(N, d) + g_N(u), \quad (2)$$

where N is physical qubit count, d is code distance, and u denotes the measurement/control setting. Here:

- h_0 is the baseline independent-error instability: relaxation, dephasing, reset errors, idle errors, and gate errors under a low-backaction reference condition.
- h_{corr} captures correlated or structured instability not caused by the immediate capacity sweep: radiation bursts, substrate events, spatially correlated drift, fabrication inhomogeneity, or crosstalk fixed by the chosen circuit.
- $g_N(u)$ is the backaction caused by the physical information-extraction and processing setting: measurement strength, cadence, readout power, resonator photon population, amplifier activity, cryo-control heat, feedback switching, and related effects.

If the setting u maps monotonically to C_{eff} , we may write $g_N(C_{\text{eff}})$ as shorthand. In real experiments the mapping is multidimensional: two settings can have the same C_{eff} but different backaction because one uses stronger readout and another uses a better decoder.

3.3 Capacity–Backaction Ratio

The proposed aggregate coordinate is

$$\rho_{\text{CB}} = \frac{\epsilon_{\text{QEC}} C_{\text{eff}} \ln 2}{h_{\text{eff}}}. \quad (3)$$

Here ϵ_{QEC} is not universal. It encodes the code- and decoder-specific efficiency with which useful information rate converts into protection margin. For a loose tracking-boundary convention one may set $\epsilon_{\text{QEC}} = 1$. For a real code, $\epsilon_{\text{QEC}} < 1$ may be more appropriate because only a fraction of extracted information contributes to suppressing logical failure.

Remark 3.3. ρ_{CB} is a diagnostic coordinate rather than a replacement for logical error rate. Logical performance remains the observable; ρ_{CB} is useful to the extent that it predicts or collapses logical-performance trends across capacity and backaction sweeps.

4 Frontier Model

4.1 A Minimal Rate Model

Let u denote an experimentally chosen measurement/control operating point. The logical error per round can be written abstractly as

$$p_{\text{L}} = \mathcal{F}(N, d, h_{\text{eff}}(N, d, u), C_{\text{eff}}(N, d, u), \mathcal{D}), \quad (4)$$

where \mathcal{D} denotes the decoder and code implementation. This function is hardware-specific. The capacity–backaction proposal is not to derive \mathcal{F} from first principles, but to test whether much of its variation under information-extraction sweeps is organized by ρ_{CB} .

A minimal phenomenological form is

$$\log p_L = A_0(N, d) + A_h h_{\text{eff}}(N, d, u) - A_C \epsilon_{\text{QEC}} C_{\text{eff}}(N, d, u) \ln 2 + \eta, \quad (5)$$

where η contains residual component-level errors. This equation is a regression target: if $A_h > 0$ and $A_C > 0$ after standard nuisance controls, then the capacity and backaction coordinates carry explanatory weight.

4.2 Operating-Point Condition

Suppose a one-dimensional knob increases useful capacity, so that C_{eff} itself can serve as the independent variable. Define a protection margin

$$M(C_{\text{eff}}) = \epsilon_{\text{QEC}} C_{\text{eff}} \ln 2 - h_{\text{eff}}(C_{\text{eff}}) = \epsilon_{\text{QEC}} C_{\text{eff}} \ln 2 - h_0 - h_{\text{corr}} - g_N(C_{\text{eff}}). \quad (6)$$

Capacity helps locally when

$$\frac{dM}{dC_{\text{eff}}} = \epsilon_{\text{QEC}} \ln 2 - g'_N(C_{\text{eff}}) > 0. \quad (7)$$

It hurts locally when

$$g'_N(C_{\text{eff}}) > \epsilon_{\text{QEC}} \ln 2. \quad (8)$$

An optimum occurs near

$$g'_N(C_{\text{eff}}^*) = \epsilon_{\text{QEC}} \ln 2, \quad (9)$$

provided g_N is sufficiently convex over the measured range.

Equations (7)–(9) give the local frontier condition: the marginal value of another timely syndrome bit is compared with the marginal physical instability caused by obtaining and using it.

4.3 Backaction Functional Forms

Several qualitative regimes are possible:

Backaction law	Behavior	Interpretation
$g_N(C_{\text{eff}}) \approx 0$	Capacity monotonically improves margin	Measurement/control overhead negligible in tested regime
$g_N(C_{\text{eff}}) = aC_{\text{eff}}$	Net benefit depends on $a < \epsilon_{\text{QEC}} \ln 2$	Linear cost per useful bit; no interior optimum unless slope exceeds benefit
$g_N(C_{\text{eff}}) = aC_{\text{eff}}^q$, $q > 1$	Interior optimum possible	Aggressive readout/control becomes increasingly damaging
$g_N(C_{\text{eff}})$ threshold-like	Sudden degradation after critical drive/cadence	Resonator depopulation, heating, leakage, or quasiparticle onset
Architecture-changing g_N	Different methods give different frontiers at same C_{eff}	Better engineering shifts the Pareto curve, not merely the operating point

4.4 Relation to Data-Rate Language

The Data-Rate Theorem states that stabilizing an unstable plant over a finite-capacity channel requires a capacity exceeding the relevant instability rate, under specified control assumptions [14–16]. The present paper applies rate accounting to QEC engineering without reducing surface-code dynamics to a scalar unstable mode. The empirical question is whether aggregate rates inferred from hardware obey a useful frontier structure.

5 Scaling Question

5.1 Local Frontier Versus Scaling Frontier

The local frontier is measured at fixed N , d , circuit family, hardware generation, and decoder. The scaling frontier asks how the entire Pareto curve changes as the processor grows.

Let the best achievable useful capacity and induced instability at size N be

$$C_{\text{eff}}^*(N) = \max_u C_{\text{eff}}(N, u) \quad (10)$$

subject to a tolerated backaction budget, and let the best achievable margin be

$$M^*(N, d) = \max_u [\epsilon_{\text{QEC}} C_{\text{eff}}(N, d, u) \ln 2 - h_{\text{eff}}(N, d, u)]. \quad (11)$$

The favorable-scaling condition is

$$M^*(N, d) > 0 \quad (12)$$

with a margin that does not collapse as N and d increase.

5.2 Exponent Accounting

For asymptotic intuition, suppose

$$h_0(N) + h_{\text{corr}}(N) = \alpha N^p, \quad C_{\text{eff}}^*(N) = \beta N^q, \quad g_N(C_{\text{eff}}^*) = \gamma N^r. \quad (13)$$

Then the margin scales as

$$M^*(N) = \epsilon_{\text{QEC}} \beta N^q \ln 2 - \alpha N^p - \gamma N^r. \quad (14)$$

If $\max(p, r) > q$, a finite crossing is possible unless coefficients or architecture changes push it beyond relevant scales. If $q \geq \max(p, r)$ with sufficient coefficient margin, the capacity–backaction frontier remains favorable.

This exponent accounting identifies the quantities to be measured: the effective scaling exponents and coefficients of useful capacity and induced instability.

5.3 Why Below-Threshold QEC Matters

Below-threshold surface-code results are evidence that, for tested distances and operating points, the capacity–backaction ratio is favorable. Such results are constraints on the frontier model.

The open question is longer-range. As systems move toward longer-lived logical memories, larger patches, more logical qubits, lattice surgery, dynamic circuits, cryogenic control integration, and lower-latency decoders, the frontier may shift. Dynamic surface-code experiments are particularly relevant because they alter circuit layout, coupler use, leakage handling, and correlated-error structure [17]. Such architecture changes can reduce h_{corr} or g_N at fixed C_{eff} , thereby moving the frontier.

5.4 Architecture as Frontier Engineering

Different architectural improvements act on different terms:

- Better materials and shielding reduce h_0 and h_{corr} .
- Faster, lower-power readout can increase C_{eff} while reducing g_N .
- Better resonator depopulation reduces residual-photon components of g_N .
- Leakage-reduction units and dynamic circuits reduce persistence and correlated memory in h_{eff} .
- Parallel and local decoding increase the timeliness factor ℓ_j in Eq. (1).
- Cryo-CMOS can reduce latency but may increase local heat; its value depends on net movement of M^* .

The sign of an engineering move is determined by the net frontier displacement, not by the capacity increase alone.

6 Experimental Program

6.1 Retrospective Audit

Existing datasets may already contain enough variation to estimate parts of the frontier. A useful retrospective audit would collect, per run or calibration block:

- code distance, patch geometry, cycle time, and circuit family;
- readout power, measurement duration, resonator depopulation wait, reset policy, and measurement cadence;
- detector-event rates, leakage estimates, erasure or flag rates, and spatial/temporal correlation metrics;
- decoder latency, queue depth, throughput, rejected samples, and missed-deadline fraction;
- logical error per round or memory lifetime;
- chip-temperature proxies such as qubit frequency drift, T_1/T_2 drift, resonator shifts, or calibrated thermometry;
- radiation or burst-event vetoes where available.

The audit should avoid treating raw readout power or raw readout rate as C_{eff} . The central quantity is useful timely information. A high-power readout that increases assignment fidelity but creates leakage and residual photons may have a different frontier position from a lower-power readout paired with a better decoder.

6.2 Prospective Capacity–Backaction Sweep

A prospective experiment should vary measurement/control settings while holding the logical task fixed.

Design.

1. Choose a stable QEC memory, repetition-code memory, or echo-like circuit with enough repetitions to estimate logical performance and detector correlations.
2. Define a baseline operating point u_0 .
3. Sweep one capacity-producing knob at a time: measurement duration, readout amplitude, cadence, accepted-update rate, decoder time budget, feedback cadence, or beacon/reference allocation.
4. Include dummy-load controls: same microwave or digital activity where possible, but with data discarded or delayed, to separate physical backaction from useful information benefit.
5. Randomize run order and interleave repeated baseline blocks to detect drift.
6. Record high-resolution physical diagnostics and decoder timing logs.

Primary observables.

- logical error per round p_L or memory lifetime;
- useful capacity estimate C_{eff} from Eq. (1);
- detector-event entropy and correlation length/time;
- leakage and reset residuals;
- readout-induced dephasing and assignment error;
- quasiparticle-sensitive proxies;
- decoder missed-deadline fraction and latency distribution.

Primary analysis. Fit models of the form

$$\log p_L = \beta_0 + \beta_C C_{\text{eff}} \ln 2 + \beta_h \hat{h}_{\text{phys}} + \beta_g \hat{g} + \beta_T \Delta T + \beta_{\text{lat}} L_{\text{dec}} + \eta, \quad (15)$$

where \hat{h}_{phys} and \hat{g} are measured instability/backaction proxies, ΔT captures thermal drift, and L_{dec} captures latency. Capacity–backaction structure is indicated when useful C_{eff} and backaction proxies retain explanatory power after standard nuisance variables are included.

6.3 Decision Outcomes

Observed result	Interpretation
p_L improves monotonically with C_{eff} ; backaction proxies stable	Frontier is favorable in tested range; capacity increase is net beneficial
p_L improves then saturates; backaction proxies rise	Interior operating optimum; capacity and backaction both matter
p_L degrades with stronger readout/control despite higher assignment fidelity	Backaction dominates; useful capacity is not the limiting variable or is overpaid physically
Dummy load and active readout degrade equally	Physical backaction, not useful information, explains trend
Active readout improves over dummy at matched load	Useful information benefit is measurable after physical cost is controlled
No collapse under ρ_{CB} after nuisance controls	Aggregate coordinate not useful for this regime

7 Falsification and Weakening Criteria

The capacity–backaction coordinate is weakened or falsified as an aggregate scaling coordinate if:

1. C_{eff} estimates do not predict logical performance better than raw readout fidelity, raw readout power, or conventional error-budget terms.
2. Backaction proxies g_N add no predictive value after standard component-level error channels are modeled.
3. Performance across measurement/readout/decoder sweeps fails to collapse against ρ_{CB} or the margin M .
4. Dummy-load controls fully explain the apparent capacity trend, leaving no separable useful-information benefit.
5. Scaling data show that $M^*(N, d)$ remains favorable with increasing N and d and no unfavorable trend in g_N .
6. Architecture changes repeatedly shift the frontier favorably, showing that the proposed aggregate constraint is not structurally limiting in the relevant regime.

The coordinate is supported as an engineering diagnostic if:

1. useful C_{eff} and induced h_{eff} can be independently estimated with uncertainty bounds;
2. logical performance exhibits an optimum or saturation explained by $g'_N(C_{\text{eff}})$ approaching the marginal information benefit;
3. the same aggregate coordinate predicts trends across multiple capacity-producing knobs;

4. the fitted scaling of $M^*(N, d)$ becomes less favorable with N or d after known nuisance variables are controlled.

Such support would establish the usefulness of a shared rate-accounting language for superconducting QEC.

8 Discussion

8.1 Aggregate Rate Coordinates

Capacity, backaction, latency, and correlated instability are often reported in separate engineering coordinates. A common rate coordinate enables comparison of operating points and scaling frontiers. If the frontier can be measured, then interventions such as faster readout, stronger measurement, local decoding, or cryo-CMOS integration can be evaluated by their net movement of M^* , rather than by one component metric.

8.2 Relation to Standard Error Budgets

Component-level error budgets remain necessary for diagnosing specific degradation channels: photon shot noise, resonator ringdown, leakage, quasiparticles, amplifier saturation, crosstalk, TLS drift, and decoder backlog. The capacity–backaction coordinate sits above those details and summarizes their net effect when the amount of useful information extracted per unit time is changed.

8.3 Relation to Standard Quantum Dynamics

The capacity–backaction frontier is compatible with Lindblad dynamics, circuit QED, Bayesian estimation, standard decoding theory, and ordinary thermodynamics. Its value is organizational and predictive: the aggregate coordinate is tested by whether it forecasts where logical performance improves, saturates, or degrades.

8.4 Implications for Architectures

Architectural progress can defeat an unfavorable frontier in at least four ways:

- reduce h_0 through better qubits, materials, shielding, and gates;
- reduce h_{corr} through layout, dynamic circuits, leakage removal, and correlated-error mitigation;
- increase C_{eff} through faster readout, better assignment, parallel decoding, and lower latency;
- reduce g_N by making each useful bit cheaper in heat, photons, crosstalk, and quasiparticle production.

The most valuable improvements are those that move more than one term favorably. For example, a circuit redesign that reduces correlated errors and allows lower readout power at the same logical performance shifts both h_{corr} and g_N .

8.5 Implications for Reporting

Future QEC reports could make the frontier more visible by reporting:

- useful syndrome mutual information per cycle;
- decoder latency distribution and missed-deadline rate;
- readout power/cadence together with leakage and residual-photon proxies;
- spatial and temporal detector-correlation metrics;
- logical performance under matched dummy-load controls;
- performance under deliberate cycle-time/readout-strength sweeps.

These quantities are often available internally. Publishing even partial versions would allow cross-platform comparison without requiring proprietary low-level logs.

9 Conclusion

This paper formulates a capacity–backaction frontier for superconducting quantum error correction. QEC is treated as finite information extraction under physical backaction.

The proposed contribution is a capacity–backaction coordinate:

$$h_{\text{eff}}(N, C_{\text{eff}}) = h_0(N) + h_{\text{corr}}(N) + g_N(C_{\text{eff}}), \quad \rho_{\text{CB}} = \frac{\epsilon_{\text{QEC}} C_{\text{eff}} \ln 2}{h_{\text{eff}}}.$$

The coordinate is useful to the extent that it predicts logical-performance trends across capacity and backaction sweeps better than raw engineering variables alone. Its central experimental question is whether useful information capacity and induced physical instability can be scaled independently, or whether an unfavorable Pareto frontier appears as processors grow.

A Appendix: Minimal Preregistered Analysis Template

Primary hypothesis. At fixed circuit family, code distance, run order randomization, and monitored thermal conditions, logical error per round is better predicted by a model containing both useful capacity C_{eff} and measured backaction proxies than by a model containing nominal readout rate or raw readout power alone.

Primary endpoint. Logical error per round, memory lifetime, or equivalent task-level failure probability.

Primary predictors. Useful capacity C_{eff} , backaction proxy vector \hat{g} , decoder latency L_{dec} , detector-correlation metrics, leakage, readout-induced dephasing, and thermal proxies.

Controls. Matched dummy-load blocks, interleaved baseline blocks, fixed randomization seed for run order, and exclusion criteria for large drift or burst events.

Model comparison. Compare:

$$\mathcal{M}_0 : \log p_L \sim \text{temperature} + \text{readout fidelity} + \text{cycle time}, \quad (16)$$

$$\mathcal{M}_1 : \log p_L \sim \mathcal{M}_0 + C_{\text{eff}}, \quad (17)$$

$$\mathcal{M}_2 : \log p_L \sim \mathcal{M}_0 + C_{\text{eff}} + \hat{g}, \quad (18)$$

$$\mathcal{M}_3 : \log p_L \sim \mathcal{M}_0 + \rho_{\text{CB}} \text{ or } M. \quad (19)$$

Use cross-validated predictive likelihood, AICc/BIC as secondary summaries, and bootstrap confidence intervals for marginal slopes.

Positive engineering result. \mathcal{M}_2 or \mathcal{M}_3 improves predictive performance over \mathcal{M}_0 and \mathcal{M}_1 , with signs consistent with useful capacity improving performance and backaction worsening it.

Negative result. No predictive gain from C_{eff} , \hat{g} , ρ_{CB} , or M after conventional error-budget variables are included.

B Appendix: Glossary

C_{eff}	Useful online correction capacity, bits/s. Information must be measured, decoded, and timely.
h_{eff}	Effective QEC instability rate, nats/s, including physical backaction.
$g_N(C_{\text{eff}})$	Backaction function: extra instability induced by obtaining and using capacity.
ϵ_{QEC}	Code/decoder-specific efficiency or margin parameter.
ρ_{CB}	Capacity–backaction ratio.
M	Protection margin, $\epsilon_{\text{QEC}} C_{\text{eff}} \ln 2 - h_{\text{eff}}$.

References

- [1] Aernoud Dekker. Bandwidth-limited quantum control: A finite-rate phase-reference control law and its operational benchmark. OSF Preprints, 2026. Preprint.
- [2] Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614:676–681, 2023. doi: 10.1038/s41586-022-05434-1.
- [3] Google Quantum AI and Collaborators. Quantum error correction below the surface code threshold. *Nature*, 638:920–926, 2025. doi: 10.1038/s41586-024-08449-y.
- [4] P. Bertet et al. Dephasing of a superconducting qubit induced by photon noise. *Physical Review Letters*, 95(25):257002, 2005.
- [5] J. Gambetta et al. Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting. *Physical Review A*, 74(4):042318, 2006.
- [6] G. Zhang et al. Suppression of photon shot noise dephasing in a tunable coupling superconducting qubit. *npj Quantum Information*, 3:1, 2017.
- [7] M. Lenander et al. Measurement of energy decay in superconducting qubits from nonequilibrium quasiparticles. *Physical Review B*, 84(2):024501, 2011.
- [8] C. Wang et al. Measurement and control of quasiparticle dynamics in a superconducting qubit. *Nature Communications*, 5:5836, 2014.
- [9] C. H. Liu et al. Quasiparticle poisoning of superconducting qubits from resonant absorption of pair-breaking photons. *Physical Review Letters*, 132(1):017001, 2024.
- [10] Antti P. Vepsäläinen et al. Impact of ionizing radiation on superconducting qubit coherence. *Nature*, 584:551–556, 2020.
- [11] C. D. Wilen et al. Correlated charge noise and relaxation errors in superconducting qubits. *Nature*, 594:369–373, 2021.
- [12] Matt McEwen et al. Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits. *Nature Physics*, 18:107–111, 2022. doi: 10.1038/s41567-021-01432-8.
- [13] Ben Barber et al. A real-time, scalable, fast and resource-efficient decoder for a quantum computer. *Nature Electronics*, 8:84–91, 2025. doi: 10.1038/s41928-024-01319-5.
- [14] Sekhar Tatikonda and Sanjoy Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.
- [15] Girish N. Nair and Robin J. Evans. Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM Journal on Control and Optimization*, 43(2):413–436, 2004.
- [16] Christoph Kawan. *Invariance Entropy for Deterministic Control Systems*, volume 2089 of *Lecture Notes in Mathematics*. Springer, 2013.
- [17] Alec Eickbusch et al. Demonstration of dynamic surface codes. *Nature Physics*, 2025. doi: 10.1038/s41567-025-03070-w.
- [18] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002.

- [19] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012.
- [20] Barbara M. Terhal. Quantum error correction for quantum memories. *Reviews of Modern Physics*, 87(2):307–346, 2015.
- [21] Alexandre Blais, Arne L. Grimsmo, S. M. Girvin, and Andreas Wallraff. Circuit quantum electrodynamics. *Reviews of Modern Physics*, 93(2):025005, 2021.