waveguide

SUPPRESSING QUANTUM ERRORS BY FINITE GAP DISSIPATION IN CHIRAL RING WAVEGUIDES

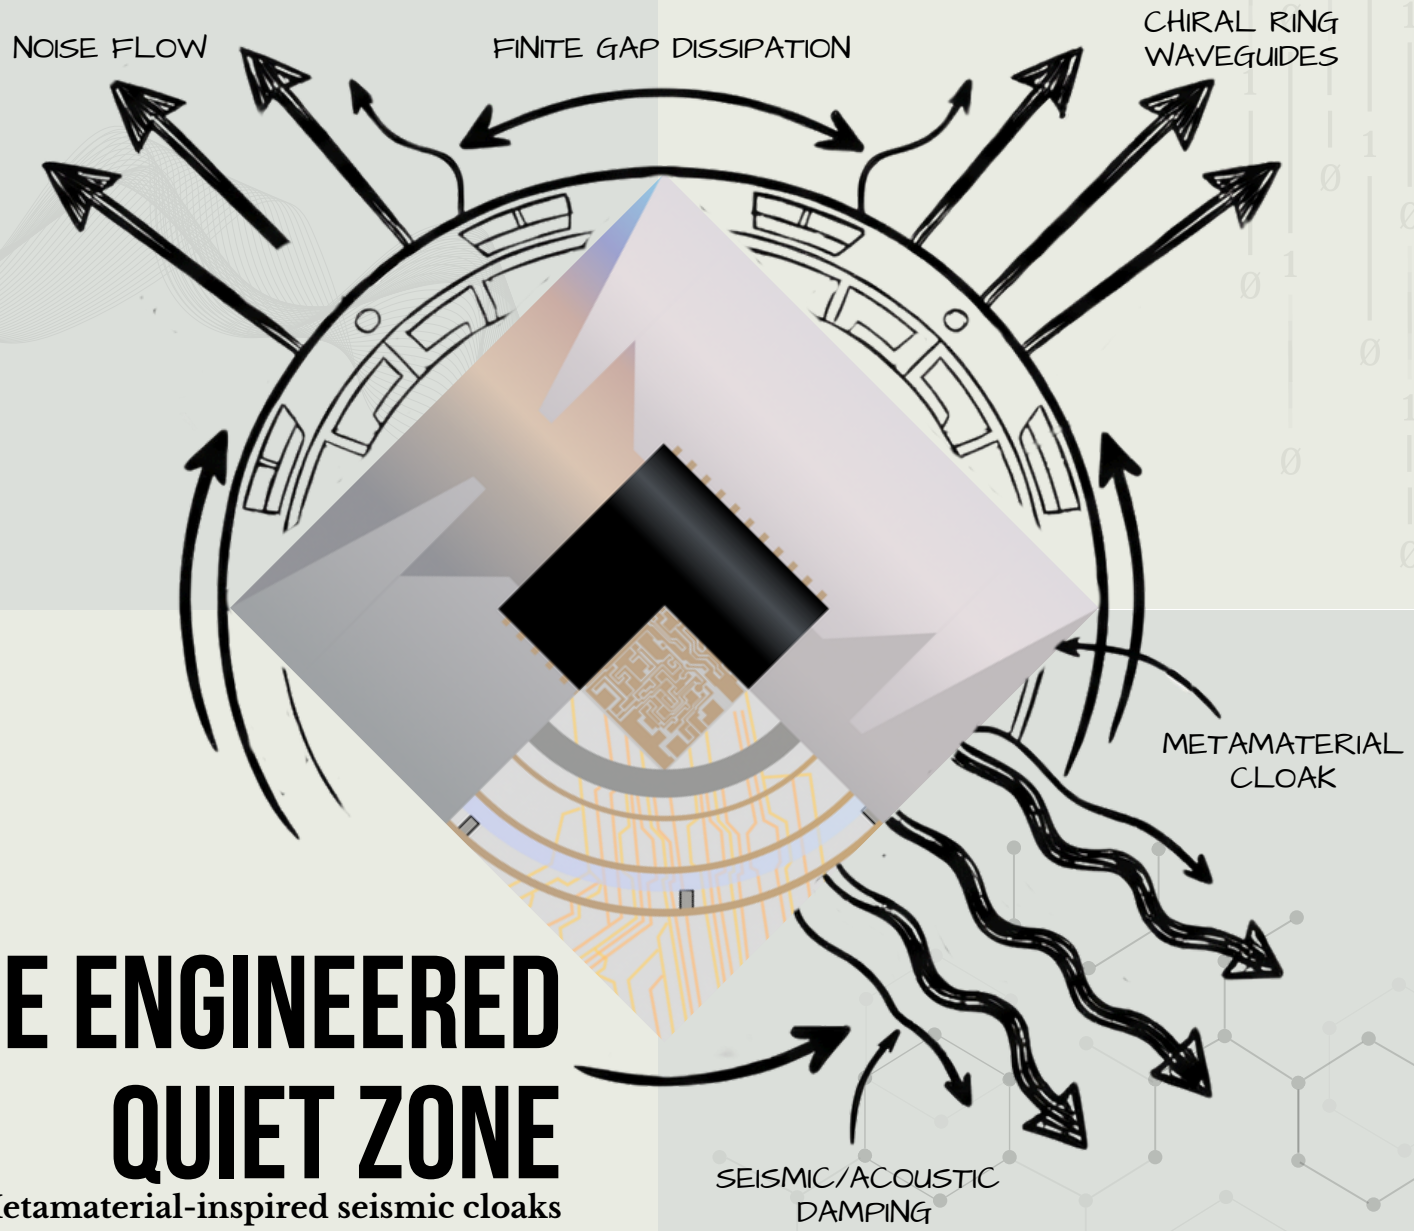

THE ENGINEERED QUIET ZONE

Metamaterial-inspired seismic cloaks steer noise outward, preserving the logic.

QUIET CERTIFICATE

A measurement grade safety stamp proving the qubit's surrounding package stays calm and cannot amplify or feed energy into the qubit during compute. It turns invisible microwave chaos into a clear pass or fail report at a declared plane.

PASSIVE EVAC.

A pump free "error drain" that steers stray excitations away from the qubit and into engineered sinks using only the package geometry. Inspired by seismic cloaking, it redirects the disturbance outward so the protected region stays quiet.

TAIL SAFTY

Proof that PACER does not trade rare catastrophic failures for better averages, by stress testing bursts and long tail events. It adds a reliability report that checks stability over time, not just a single headline improvement.

NEST ORIGINALS
11 November 2025
Vol .01, No. 001

CONTENTS

NTRODUCTION

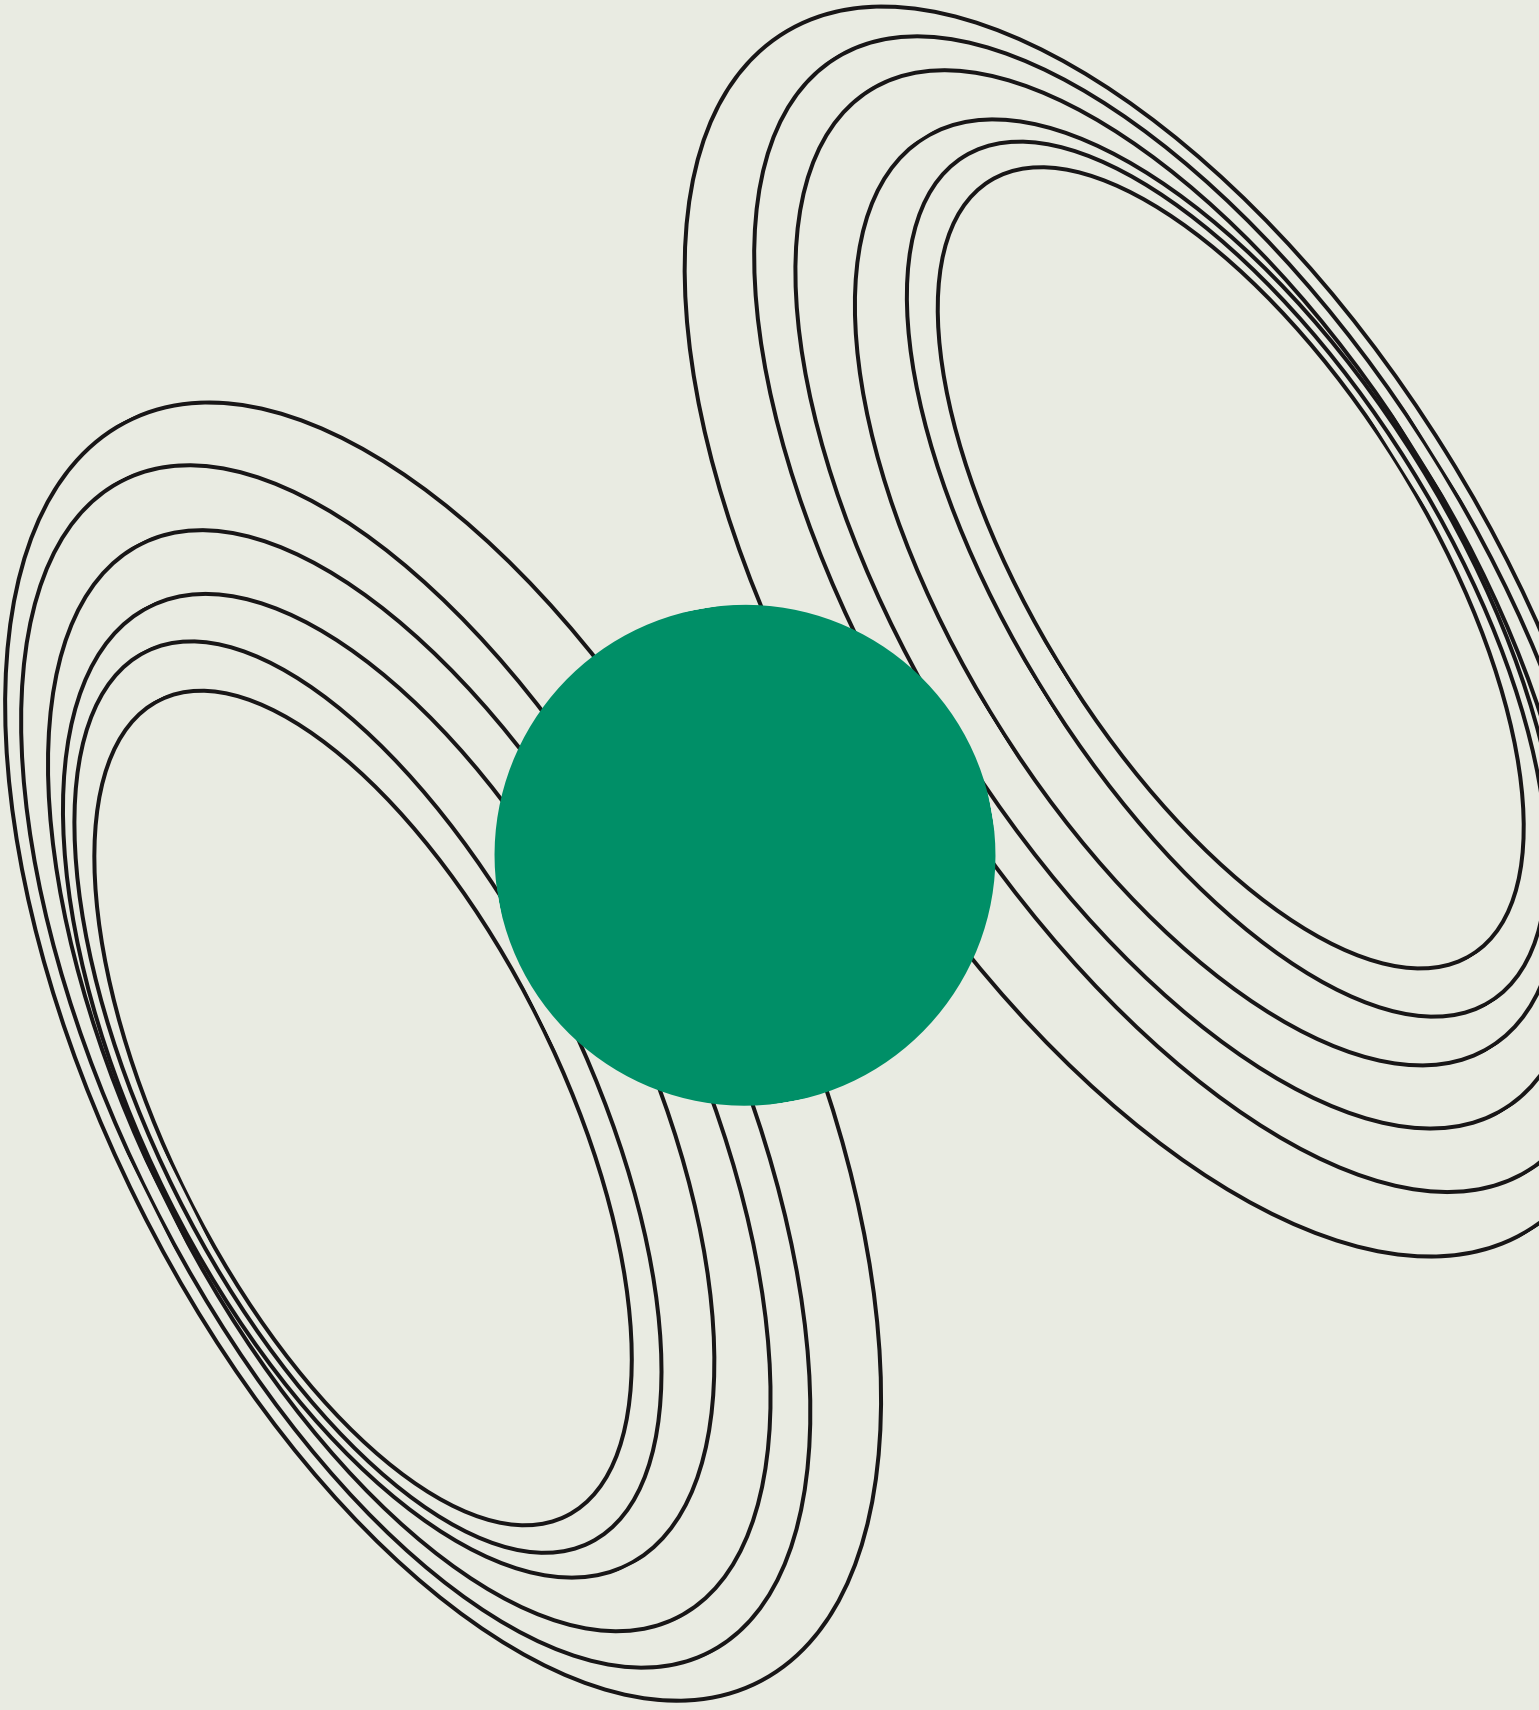

Abstract

What is PACER?

PACER turns the QPU environment into a versioned, testable boundary condition that reshapes noise at the stabilizer-cycle level, converting damage into located records and suppressing rare long-tail failures through engineered evacuation to witnessed drains.

Key Innovations

Boundary conditioned quiet zone

PACER makes the QPU environment a declared, testable boundary that redirects noise away from the protected patch.

Multi-ring routing plus drains

Engineered rings guide error energy outward into matched-loss drains, reducing backflow and reflections.

Witnessed, auditable performance

Drains produce cycle-synchronous located flags and tail-first metrics with explicit gates, so claims are falsifiable and reproducible.

Why Does it Matter?

Architectural leverage on scalability

By reducing the effective noise incident on the protected region, PACER can lower the redundancy and cycle overhead required to reach a target logical reliability, improving the scaling economics of fault tolerance.

Observability and inferential control

Routing a larger fraction of deleterious dynamics into monitored channels increases event localization, enabling evidence-based inference and reducing reliance on unobserved error assumptions.

Tail-risk mitigation as the governing criterion

In realistic devices, logical failure is frequently dominated by rare, correlated excursions rather than mean error rates; PACER explicitly targets suppression and qualification of these tail events.

Evolution of PACER

Waveguide V1

The initial idea behind PACER, treats errors as excitations that can propogate

Waveguide V2

A complete passive defense stack and presents aggressive forward projections as part of the simulation framework

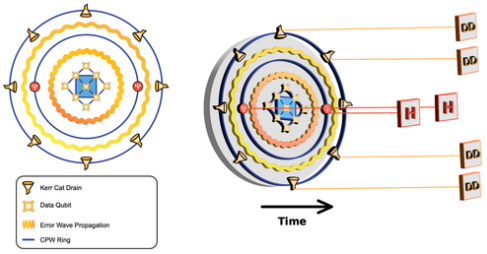

Waveguide V3

Waveguide v3 fixes rigor and adds chiral outward error flow compared to version 2

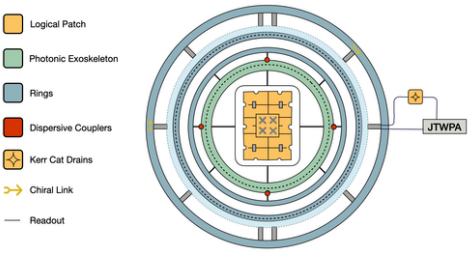

PACER

PACER Breakwater creates a quiet zone by steering errors into drains and proving the boundary works with clear tests.

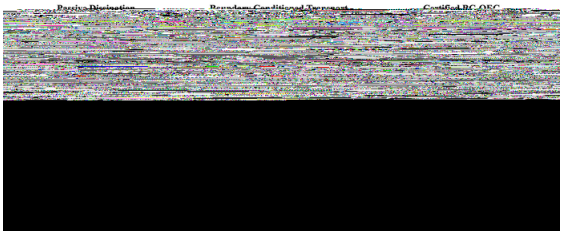

Starting from a simple “what if we could have a quantum version of seismic cloaking ” idea, PACER over the past year evolved into a disciplined, test driven architecture that treats the region around the qubits as an engineered pathway. Errors are encouraged to flow outward into designated sinks instead of lingering and re entering the protected area. As the concept matured, the focus shifted from drawing a compelling structure to specifying what must be measured, what counts as success, and what failure looks like, so the claimed “quiet zone” is not a metaphor but a repeatable, operational effect supported by clear evidence.

How Does PACER Work

PACER operates by reshaping how errors interact with the quantum processor, rather than by reacting to errors after they occur. The core idea is to treat deleterious dynamics as transportable excitations whose trajectories can be redirected by engineering the environment surrounding the protected qubits. Instead of asking how errors are detected and corrected inside the code, PACER asks a different question, what boundary conditions determine whether error energy enters, lingers, or escapes the protected region?

Errors as Transported Excitations

In realistic quantum hardware, noise is not a sequence of isolated, memoryless events. Energy injected by control imperfections, coupling to spurious modes, or correlated environmental fluctuations can propagate through the device, interact with multiple qubits, and persist across cycles. These processes give rise to rare but catastrophic failure events that dominate logical error rates.
PACER models these effects as excitations that move through an engineered environment. If left unmanaged, such excitations can scatter back into the logical region. If properly guided, they can be evacuated before causing uncorrectable damage.

Boundary Conditioning Instead of Internal Correction

Traditional quantum error correction focuses on repeatedly interrogating the code to infer internal faults. PACER instead conditions the boundary of the protected region so that harmful dynamics are discouraged from entering and encouraged to leave.
The region surrounding the logical qubits is treated as an engineered pathway with declared properties. Within this framework, the environment is no longer an uncontrolled source of noise but a structured interface that redirects error flow outward. This boundary conditioning reduces the effective noise incident on the code without requiring adaptive control or real-time feedback.

Rings, drains, and witnesses

PACER implements this idea using three physical roles:

Routing structures	Matched Drains	Witness Channels
Concentric engineered rings guide error energy away from the protected patch, suppressing backflow and reducing reflection.	Lossy elements placed beyond the routing region absorb outgoing excitations in a controlled manner, preventing their return while respecting physical constraints.	A subset of the evacuated dynamics is monitored, producing cycle-synchronous records that localize events and enable direct measurement of tail behavior.

Consequence for Logical Reliability

By reducing the effective noise seen by the code and increasing the fraction of errors that are both evacuated and witnessed, PACER changes the economics of fault tolerance. Fewer redundant resources and fewer cycles are required to achieve a target logical reliability, and rare correlated excursions are explicitly targeted rather than implicitly averaged away.
The result is a system in which performance claims can be tied to measurable boundary behavior and event statistics, rather than inferred solely from mean error rates or assumed noise models.

Cycle Level Operation and Declared Contract

PACER is defined and evaluated at the level of the stabilizer cycle. Its output is not a qualitative notion of reduced noise, but a cycle-synchronous effective noise contract exported to decoding and analysis.
During each cycle, the protected region executes the declared gate and measurement sequence while the surrounding boundary conditions shape the flow of deleterious dynamics. Outgoing excitations are either absorbed into matched drains or converted into witnessed records. The remaining unobserved damage is constrained by declared and testable small parameters associated with boundary performance.

What PACER does and does not claim

PACER does not assume that all errors are detected, nor that all noise is rendered benign. Instead, it guarantees that any claimed reduction in logical error rate is supported by measured boundary behavior and observed event statistics, rather than inferred solely from mean error rates or assumed noise models.
If the declared boundary conditions fail to hold, or if witnessed data does not satisfy acceptance criteria, the corresponding performance claims are invalidated. This separation between mechanism and claim is intentional and central to the framework.

Why this changes the fault-tolerance regime

By reducing the effective noise incident on the protected region and increasing the fraction of deleterious dynamics that are both evacuated and witnessed, PACER alters the scaling economics of fault tolerance. Redundancy and cycle overhead are lowered not by stronger internal correction, but by reshaping the environment in which the code operates.
Crucially, rare correlated excursions are treated as first-class failure modes rather than statistical outliers. PACER targets these events directly through engineered evacuation and explicit measurement, enabling logical reliability to be argued from evidence rather than assumption.

PACER Certificates

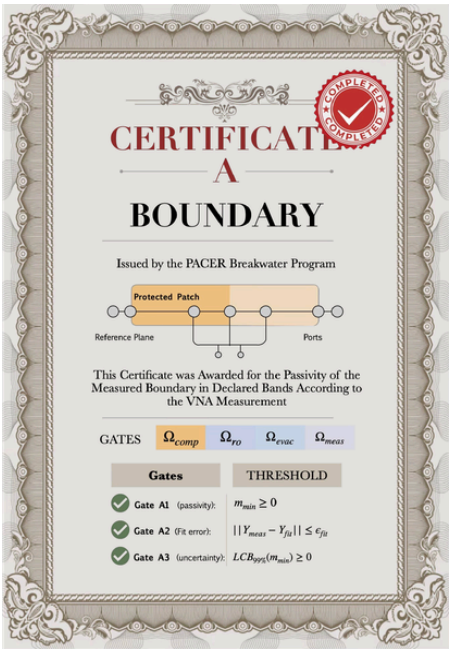

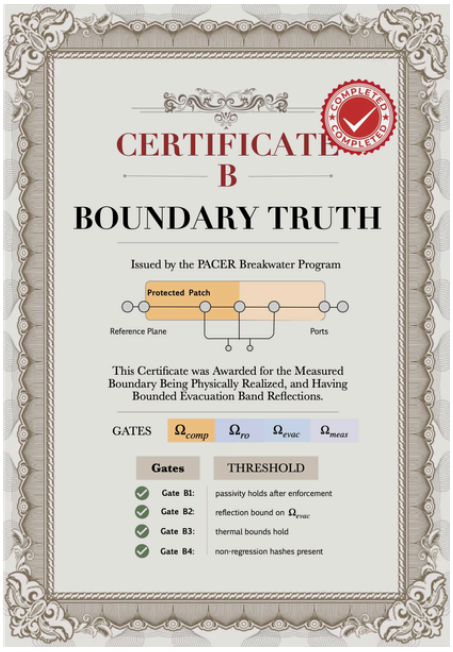

PACER uses certificates as formal, auditable acceptance tests that regulate what the program is permitted to claim about a device, package, or operating point. In a field where performance can appear improved by reporting only favorable metrics, tuning post hoc, or emphasizing median outcomes while rare correlated failures dominate reliability, certificates force every claim to be backed by a compact, reproducible set of gates. Each certificate specifies the measured object, the declared operating bands, the derived quantities computed from the measurement, and explicit pass thresholds. A certificate is therefore neither a legal guarantee nor a statement of perfection. If a gate fails or required provenance is missing, the claim is treated as unproven, regardless of how promising any single headline number appears.

Taken together, the certificate bundle is PACER's mechanism for converting qualitative architectural narratives into quantitative, checkable commitments. The program's claims are deliberately staged, Certificate A establishes that the measured boundary model is physically legitimate, Certificate B establishes that the boundary is realized and behaves as intended with bounded inward reflection, and Certificate C establishes that performance claims are not driven by a fragile improvement in typical cases while leaving rare catastrophic behavior unconstrained. This structure is designed to make PACER's evidence portable: a reviewer, collaborator, or future replication effort can identify exactly what was measured, how it was transformed into a boundary object and cycle-level evidence, which gates were applied, and which claims are justified by the resulting pass set.

Bespoke PACER QPU CHIP

PACER's "tailored QPU" is not a standard transmon processor with a nicer package. It is a co-designed qubit plus boundary system whose deliverable is a cycle-synchronous noise instrument on the protected patch, backed by a measured boundary object that is stable, passive, and versioned. If those two exports are missing, the design reduces to an unverified packaging concept with no auditable QEC meaning.

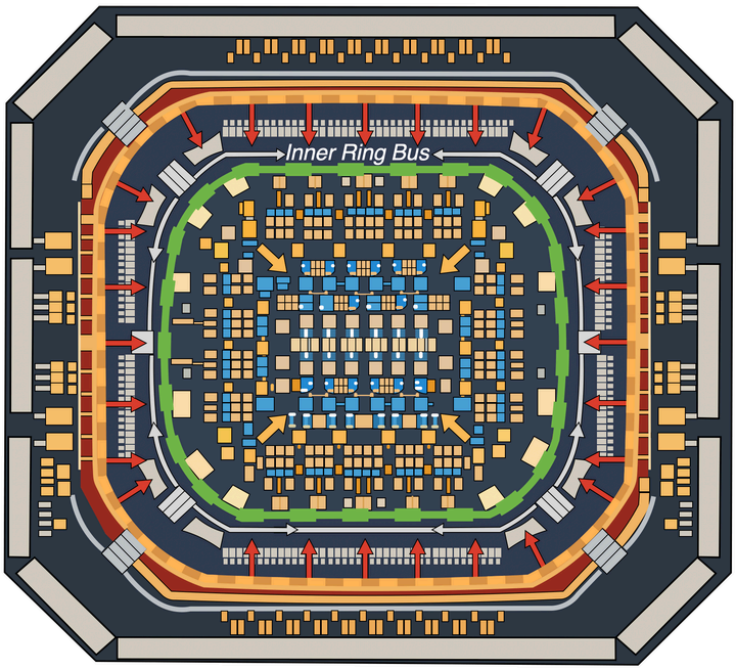

A standard superconducting QPU pipeline is: maximize coherence, reduce crosstalk, then run QEC on whatever noise remains. PACER inverts the emphasis:

1. Shape the boundary and transport so that a large fraction of damaging events become located (erasure flags, drain witness hits).
2. Export a cycle-synchronous instrument that decoders can consume immediately.
3. Veto wins that fail in the tail, using first-passage and hazard diagnostics rather than celebrating median lifetime improvements

RESEARCH

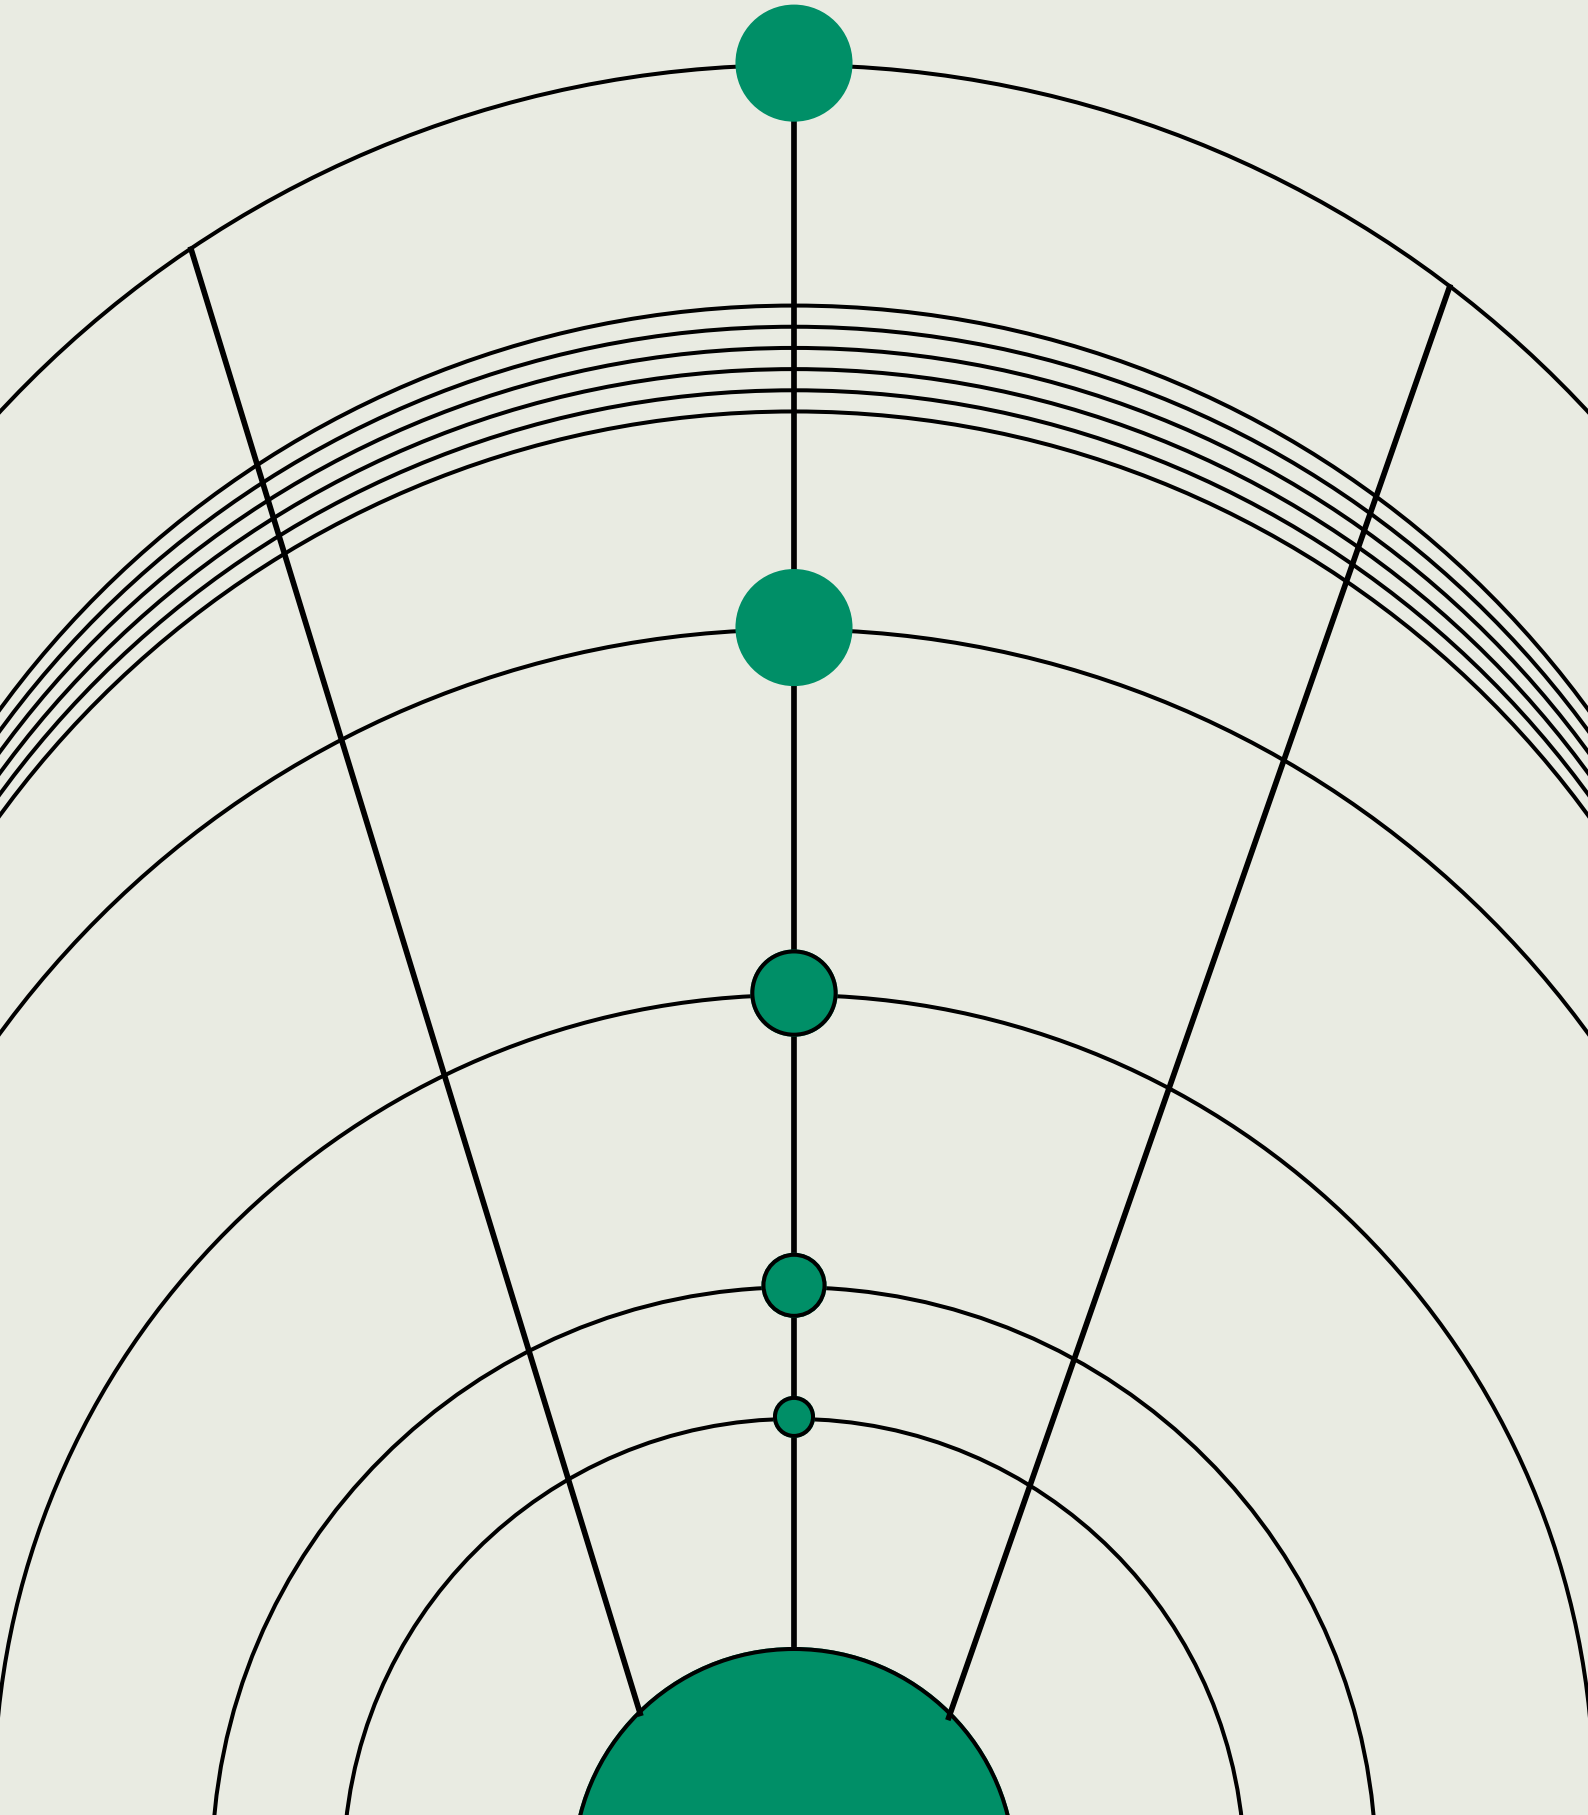

Auditable electromagnetic boundaries for below-threshold quantum error correction

Kim Jaebum[1]

[1]NEST, Hefei, China
erickim1492@gmail.com

Abstract

As quantum error correction enters a below-threshold regime, residual logical failures can be set by rare, device-specific excursions rather than by mean physical error rates alone. A recurrent driver of such excursions is the chip package microwave boundary, which is commonly treated as an implicit bath and therefore excluded from decoder inputs. Here we promote the boundary to an explicit, reproducible input by representing it at a declared reference plane as a calibrated multiport admittance and compiling it into a stable, passive rational macromodel. We then engineer transport and monitored drains that redirect a subset of boundary-driven excitations into cycle-resolved witness records, enabling decoders to use that subset as located side information. In validated simulation at the stated operating point, a distance-5 memory benchmark over $N_{\mathrm{cyc}} = 40$ cycles with $\tau_{\mathrm{cyc}} = 1.0\,\mu\mathrm{s}$ (2,000,000 shots) yields $p_L = 5.372 \times 10^{-3}$, corresponding to an effective logical failure rate $\Gamma_L = 134.7$ Hz (95% CI $[132.1, 137.2]$ Hz). A boundary-toggle protocol links controlled perturbations to exposure shifts. In a 10-condition sweep, toggle-induced shifts satisfy $\rho = 0.9965$ (bootstrap 95% CI $[0.9913, 0.9996]$) with perfect rank ordering ($\rho_s = 1.0$), and obey an approximately linear map between absorbing and reflecting variants. Across ring depth $n = 1 \ldots 12$, extracted transport proxies follow $\Delta_{\mathrm{trans}}(n) \propto n^{-\alpha}$ with $\alpha \approx 0.89$ at two couplings, consistent with scalable evacuation under the stated model class. Together these elements define a falsifiable interface between packaging physics and logical reliability and treat cycle-synchronous export and tail qualification as operating requirements rather than post hoc diagnostics.

Below-threshold quantum error correction can be limited by rare excursions that depend on a processor's electromagnetic boundary rather than on average error rates alone. We treat the chip package boundary as an experimental input by instantiating its multiport admittance at a declared reference plane and compiling it into a stable, passive macromodel. PACER couples this boundary-aware description to engineered transport and monitored drains that convert a subset of boundary-driven excitations into cycle-synchronous witness records usable by decoders as located side information, while bounding the residual unflagged contribution at the stated operating point.

Introduction

Fault-tolerant quantum computation is often framed by a threshold. Once physical error rates fall below a critical value, increasing redundancy should suppress logical failure. As experiments

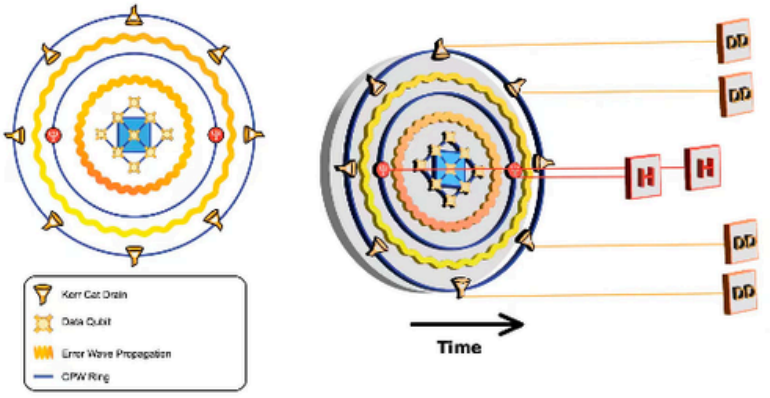

Figure 1: Conceptual PACER schematic (overview). A protected qubit patch is surrounded by concentric CPW ring transport stages that guide boundary-driven excitations outward. Engineered drains evacuate excitations and produce monitored click records that act as cycle-synchronous witness evidence. The operational objective is to convert a subset of boundary-driven damage into decoder-visible located information while maintaining a physically admissible boundary description at a declared reference plane.

approach and enter a below-threshold regime, the central question shifts from average suppression to mechanism. Which processes dominate the residual failures, where they occur, and how they cluster in time. Surface-code experiments now probe this regime directly, where structured mechanisms, mitigation choices, and rare excursions can set an effective floor [1, 2]. In parallel, alternative code families, including quantum low-density parity-check constructions, sharpen the same message from a different direction. Correlated noise and decoder architecture can dominate overhead even when mean physical error rates appear favorable [3].

A recurring and largely platform-independent constraint is the electromagnetic boundary between a device and its surroundings. Spurious package modes, impedance discontinuities, and burst-like disturbances can reshape the effective noise seen by a code, with disproportionate impact on the tails of the logical-failure distribution. Packaging has therefore become a primary design axis for superconducting quantum processors [4–6]. Below threshold, two practical gaps remain. First, boundary engineering is rarely expressed as a shareable interface-level object at a specified plane, with sufficient metadata to reproduce conversions, fits, and validation checks [7, 8]. Second, boundary information is typically not exported at the stabilizer-cycle timescale in a form a decoder can use, leaving no disciplined route to qualify tail behavior as a primary endpoint [9, 10].

Here we introduce boundary-conditioned quantum error correction (BC-QEC), in which the near-field electromagnetic environment is elevated from an implicit bath to a declared boundary object at a reference plane. In an input output description, this boundary supplies effective driving fields and provides measurable outputs, so it can be treated as an experimental input rather than background context [11, 12]. Our objective is operational. We export a cycle-level description that separates decoder-visible, located evidence from an explicitly bounded unobserved remainder, stated at the level of instruments rather than microscopic bath assumptions [13–15]. The boundary object is reconstructed from multiport network data and compiled into a stable, passive rational macromodel suitable for time-domain prediction and decoder integration [16–18]. Finally, the link between boundary changes and observed rate shifts is assessed using controlled boundary perturbations that

produce falsifiable differential predictions under documented changes [4, 19].

PACER (*Passive Architecture for Chiral Error Rejection*) instantiates BC-QEC by declaring the boundary as a multiport admittance $Y_{\mathrm{in}}(\omega)$ at a specified reference plane and wave convention, enforcing stability and passivity in the compiled macromodel $\widehat{Y}(s)$, and engineering transport and monitored drains that convert a subset of boundary-driven excitations into cycle-resolved witness records. Its central boundary object $Y_{\mathrm{in}}(\omega)$ is obtained from calibrated network measurements or from simulation instantiation and is expressed under a declared power-wave convention [20–22]. Directional and dissipative motifs connect to reservoir-engineering approaches for nonreciprocal interactions in superconducting microwave systems [23, 24], and structured resonator networks motivate realizations of engineered transport lattices [25]. The focus of this manuscript is the exported interface. It consists of a boundary object and a cycle-synchronous record with declared semantics, together with tail-targeted statistics that can be recomputed from archived artifacts.

Methods

Boundary record (Ξ_{BC}) for boundary objects used in quantitative analyses

We treat the chip package electromagnetic environment as an explicit input to the analyses reported here. Each boundary object used in quantitative results is accompanied by a complete boundary record, denoted Ξ_{BC}, that is sufficient to regenerate the boundary response, reproduce conversions and model fits, and re-evaluate the validation checks applied in this study. The primary claim path uses simulation-instantiated boundaries. Experimental metrology is discussed as an implementation route and is not required for the reported results [7, 8].

For PACER, the boundary record associated with $Y_{\mathrm{in}}(\omega)$ includes a reference-plane specification and port map, the raw boundary response as multiport $S(\omega)$ or $Y_{\mathrm{in}}(\omega)$ on a specified frequency grid together with the simulation method and settings required to reproduce it, the wave and characteristic-impedance conventions $Z_0(\omega)$ together with the exact $S \to Y$ mapping, a stable passive rational macromodel $\widehat{Y}(s)$ with deterministic fitting settings and residual diagnostics, and a numerical uncertainty envelope together with the stability, passivity, and conformance checks applied after any update to synthesis, conversion, conditioning, or fitting.

Performance statements that depend on boundary information are made only at operating points where Ξ_{BC} is complete and where the stated checks pass. If a required element is missing, internally inconsistent, or fails validation, the corresponding boundary object is reported as qualitative context and is not used for quantitative analyses.

Boundary synthesis, reference-plane declaration, and numerical uncertainty

All boundary analyses were referenced to an interface plane and port map recorded in Ξ_{BC}. The plane was chosen to separate the boundary network from the protected-patch model and to remove ambiguity in port ordering and in any reductions or de-embedding applied downstream [4].

The boundary response was supplied either as simulated multiport scattering parameters $S(\omega)$ or directly as $Y_{\mathrm{in}}(\omega)$ on a specified grid. When $S(\omega)$ was supplied, Ξ_{BC} recorded the simulation settings required to regenerate the response, including the grid definition, discretization controls, solver tolerances, and any conditioning performed at source. When a macromodel was used as the

primary representation, Ξ_{BC} also recorded the evaluation settings used to generate derived $S(\omega)$ or $Y_{\mathrm{in}}(\omega)$ tables.

Because boundaries were simulation-instantiated, uncertainty was treated as numerical and model-discrepancy uncertainty rather than measurement noise. We recorded one uncertainty envelope in Ξ_{BC}. It was a convergence envelope from discretization refinement, a solver-tolerance envelope, or a conservative envelope induced by conditioning in the $S \to Y$ inversion chain. Conformance tests and passivity tolerances were interpreted relative to the recorded envelope.

Conversion $S \to Y$ under specified conventions

When a boundary was provided as $S(\omega)$, we converted it to multiport admittance $Y_{\mathrm{in}}(\omega)$ using the wave convention and characteristic-impedance matrix $Z_0(\omega)$ recorded in Ξ_{BC} and held fixed for comparisons. For diagonal or full positive-definite $Z_0(\omega)$, we used the power-wave conversion

$$Z(\omega) = Z_0^{1/2}(\omega)\,(I + S(\omega))(I - S(\omega))^{-1}\,Z_0^{1/2}(\omega), \qquad Y_{\mathrm{in}}(\omega) = Z^{-1}(\omega). \tag{1}$$

When $(I - S)$ was ill-conditioned, we applied a recorded conditioning rule (optionally including Tikhonov regularization) and stored all hyperparameters in Ξ_{BC}. All conversions were recomputable from archived $S(\omega)$ and the recorded conventions [20, 22].

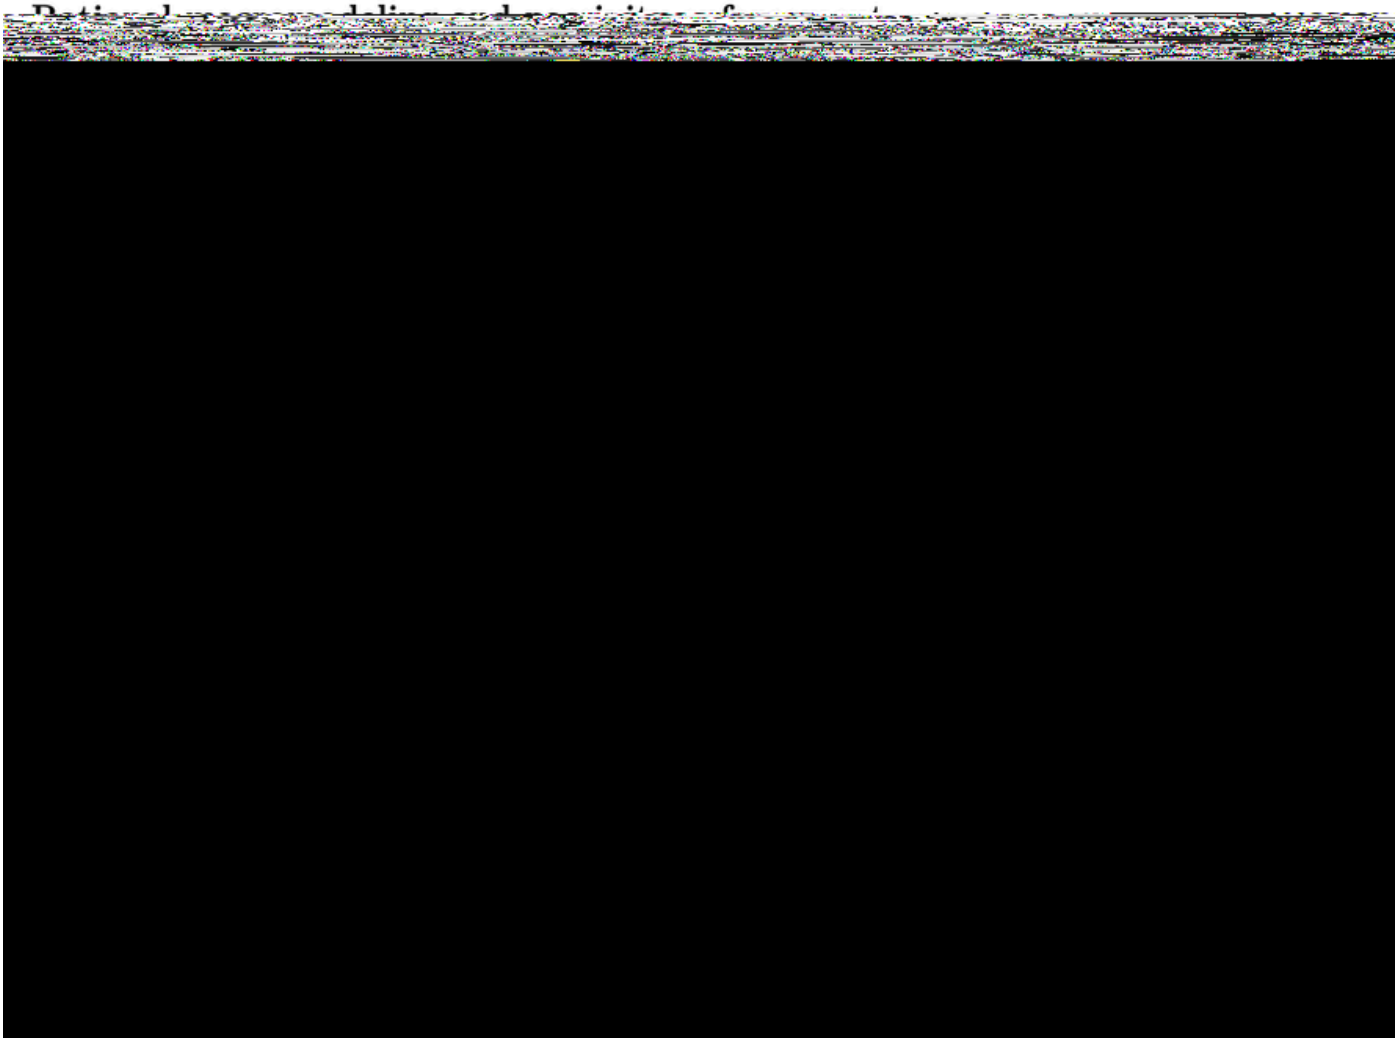

micro-sweeps, we did not use micro-sweep numerics as quantitative inputs to bounds or headline claims, and we instead treated them as a flag motivating a dedicated control sweep [15, 44, 45].

Acceptance criteria and BC-QEC bound

At each operating point, we evaluated criteria combining RF admissibility, witness calibration, tail diagnostics, and optional non-Markov diagnostics. RF passivity used the conditioned $S \to Y$ conversion under the recorded convention, with passivity margin defined as the minimum eigenvalue of $(Y + Y^{\dagger})/2$ [26, 46]. Reflection criteria included $\max_{\omega,i} |S_{ii}(\omega)|$ over the evacuation band and a reduced S_{11} criterion under a recorded reduction. The Bode Fano integral was computed by numerical integration of $\int \log(1/|S_{11}|)\, d\omega$ over $\omega = 2\pi f$, with clipping and load-model thresholds recorded in configuration [47].

Tail criteria included hazard-flatness diagnostics, tail-ratio statistics with bootstrap confidence intervals, extremal-index diagnostics when invoked, and minimum-event requirements [9, 10, 48]. Witness criteria enforced bounds on false negatives and false positives at the operating point [38, 39].

We upper-bounded the unlocated remainder using a BC-QEC bound written in terms of a boundary-conditioning defect δ_{cloak}, witness miss probability p_{fn}, and a non-Markov defect $\varepsilon_{\mathrm{nm}}$

$$\mathrm{bound} = c_1\, \delta_{\mathrm{cloak}}^2\, \tau + c_2\, p_{\mathrm{fn}} + c_3\, \varepsilon_{\mathrm{nm}} \;\leq\; \mathrm{unlocated}_{\max}, \tag{12}$$

with coefficients c_i specified (unity unless stated) and τ the chosen time or cycle scale. Quantitative use of this bound requires that each defect parameter entering the bound is itself supported by a validated artifact at that operating point.

Reproducibility, versioning, and consistency checks (simulation)

All quantitative results in this manuscript are derived from a deposited artifact bundle that is sufficient to regenerate boundary objects, reproduce macromodel fitting and diagnostics, recompute acceptance statistics, and re-render figure panels [7, 8, 33–35]. Changes that can affect reported numbers are tracked explicitly, and the bundle includes the configuration, hashes, and intermediate tables needed for independent recomputation.

Because the boundary object and the cycle export define the interface between boundary modeling and decoding, we versioned Ξ_{BC} and the cycle-export schema. Any change that could affect quantitative outcomes (including boundary synthesis, conversion, conditioning, fitting, decoding, censoring, criterion computation, or the HEOM micro settings) incremented the corresponding version and triggered the recorded consistency checks [49].

Runs and artifacts that fail declared gates (for example incomplete repeatability fields, transient-gate failures in boundary certificates, or internally inconsistent tail-export tables) are retained in the artifact bundle with failure flags, but they are not used for manuscript numbers. This policy prevents post hoc cherry-picking and ensures that the manuscript reports only quantities supported by self-consistent, recomputable exports.

Results

Boundary objects as interfaces between packaging physics and QEC

We represent the electromagnetic boundary surrounding a protected patch as a multiport object that can be measured, modeled, and perturbed under controlled changes. Consider a protected patch coupled through m ports to a surrounding electromagnetic environment. At a specified reference plane, the multiport input admittance $Y_{\mathrm{in}}(\omega) \in \mathbb{C}^{m\times m}$ relates port voltages and currents as

$$\boldsymbol{I}(\omega) = Y_{\mathrm{in}}(\omega)\,\boldsymbol{V}(\omega). \tag{13}$$

Because reference-plane choices, calibration procedures, and wave conventions can change the inferred boundary response, all analyses in this work require $Y_{\mathrm{in}}(\omega)$ to be accompanied by Ξ_{BC} [7, 8]. This record fixes the reference plane, port ordering, wave convention, uncertainty model, and the checks used to validate the fitted macromodel.

A practical consequence is that boundary engineering can be evaluated using constraints that are both physical and exportable. In PACER, the quiet-zone boundary is designed to suppress leakage into the protected band while remaining consistent with a passive network description, so that time-domain predictions do not rely on unphysical gain. Figure 2 summarizes these requirements in the operating band, alongside reflection structure that defines the usable window for evacuation and monitoring.

Deep logical-failure benchmark at fixed cycle budget (distance 5)

We next quantify end-to-end logical reliability under a fixed cycle budget using a high-statistics Monte Carlo experiment. The benchmark is stated in three equivalent normalizations. They are the logical failure probability after a declared number of cycles, an implied per-cycle logical failure probability under a constant-hazard reduction, and an effective logical failure rate in physical units.

We run a distance-5 memory experiment for $N_{\mathrm{cyc}} = 40$ QEC cycles at a declared cycle time $\tau_{\mathrm{cyc}} = 1.0\,\mu\mathrm{s}$, so the exposure time per shot is

$$T \;=\; N_{\mathrm{cyc}}\,\tau_{\mathrm{cyc}} \;=\; 4.0\times 10^{-5}\ \mathrm{s}. \tag{14}$$

Across $n = 2{,}000{,}000$ shots, we observe $k = 10{,}743$ logical failures, yielding

$$\hat{p}_L \;=\; \frac{k}{n} \;=\; 5.372\times 10^{-3}, \tag{15}$$

with 95% confidence interval $[5.27\times 10^{-3},\ 5.473\times 10^{-3}]$.

We map the 40-cycle failure probability to an exponential-hazard equivalent

$$\Gamma_L \;=\; -\frac{\ln(1-\hat{p}_L)}{T}, \tag{16}$$

giving

$$\Gamma_L \;=\; 134.7\ \mathrm{Hz} \quad (95\%\ \mathrm{CI}\ [132.1,\ 137.2]\ \mathrm{Hz}). \tag{17}$$

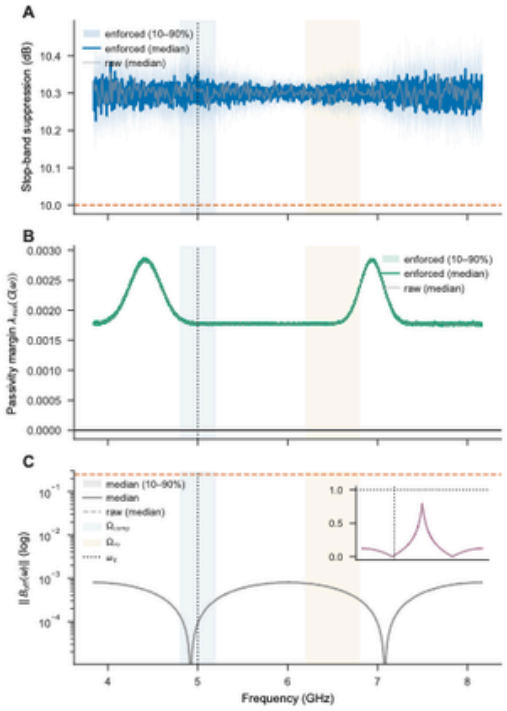

Figure 2: Quiet-zone boundary diagnostics and admissibility checks. (A) Stop-band suppression in the protected band. (B) Passivity margin of the enforced macromodel relative to the raw fit over the operating range. (C) Representative reflection magnitude over frequency, with the protected band indicated. Shaded regions indicate the 10 to 90% interval across the reported ensemble where applicable.

The corresponding per-cycle probability under the same constant-hazard reduction is

$$p_{L,\mathrm{cyc}} \;=\; 1-(1-\hat{p}_L)^{1/N_{\mathrm{cyc}}} \;\approx\; 1.35\times10^{-4} \quad (95\%\ \mathrm{CI}\ [1.32\times10^{-4},\ 1.37\times10^{-4}]). \tag{18}$$

These conversions are convenience normalizations for comparison at fixed $(N_{\mathrm{cyc}}, \tau_{\mathrm{cyc}})$ and do not assume that the underlying process is memoryless.

This benchmark summarizes the net logical vulnerability that remains after boundary conditioning and any decoder-visible evidence exported by witnessed drains. It is a single operating-point statement, meaning a fixed distance, cycle budget, cycle time, and configuration. By itself it does not establish scaling, nor does it separate which fraction of the reduction arises from modified boundary admission, guided transport, or improved located information. Those mechanisms are isolated in the ablations, boundary toggles, and tail analyses reported below.

A separate high-shot cross-check compares the same cycle budget with and without witness flags enabled. Without witness flags we obtain $p_L = 5.322\times10^{-3}$ (95% CI $[5.292\times10^{-3}, 5.353\times10^{-3}]$) across 22,000,000 shots. With witness flags we obtain $p_L = 5.372\times10^{-3}$ (95% CI $[5.317\times10^{-3}, 5.429\times10^{-3}]$) across 8,000,000 shots. These intervals overlap substantially, so this cross-check does not support a claim that the current witness implementation improves end-to-end logical failure at this operating point.

An importance-sampling variant was executed as a variance-reduction diagnostic. Because its estimate did not reconcile with the brute-force confidence interval under nominally matching

condition identifiers, we do not use importance-sampling numerics for quantitative claims in this manuscript.

The ring-depth sweep provides a compact cross-check that logical vulnerability decreases as the transport depth increases under the stated model class. At gamma fraction 1.0, the exported brute-force estimates decrease from $p_L = 8.367 \times 10^{-7}$ (95% CI $[7.361 \times 10^{-7}, 9.497 \times 10^{-7}]$) at $N_r = 1$ to $p_L = 1.655 \times 10^{-7}$ (95% CI $[1.249 \times 10^{-7}, 2.143 \times 10^{-7}]$) at $N_r = 3$. At $N_r = 5$ we observe zero failures in the reported run and therefore report a 95% upper bound $p_L \leq 4.993 \times 10^{-8}$. At $N_r = 9$ we obtain $p_L = 3.333 \times 10^{-8}$ (95% CI $[6.877 \times 10^{-9}, 9.738 \times 10^{-8}]$). These values are not used as headline performance numbers, but they are consistent with improved evacuation as the transport depth increases.

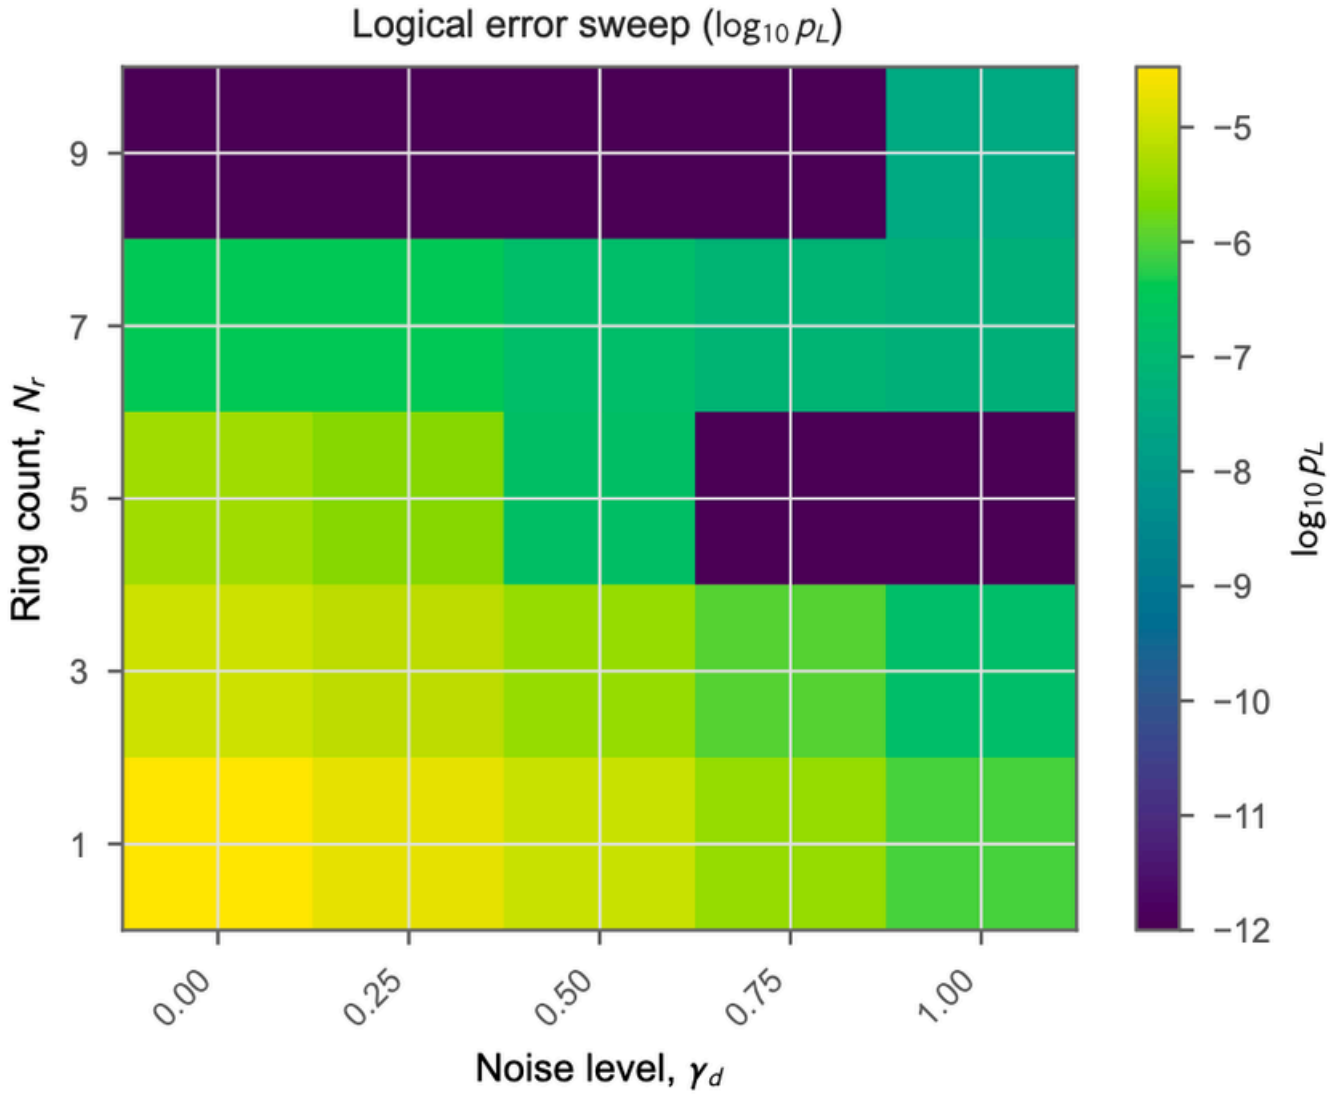

Figure 3: **Logical error landscape across boundary-noise strength and ring depth.** Heatmap of $\log_{10} p_L$ as a function of noise strength (reported as a gamma fraction on the horizontal axis) and ring count N_r on the vertical axis. Each cell aggregates brute-force estimates from the exported run summaries.

Transport cloak sets a relaxation floor with scalable evacuation

The transport region surrounding the quiet zone is designed to evacuate boundary-driven excitations into monitored drains rather than allowing them to accumulate near the protected patch. In the ring-and-drain model, increasing the number of transport stages N_r increases the effective relaxation rate while revealing a design-relevant floor set by the model class and coupling constraints. Figure 4 quantifies this scaling and shows the transient-fit procedure used to extract the reported proxies in the same parametrization used by the first-passage simulations.

Over $n = 1 \dots 12$, fitting $\Delta_{\mathrm{trans}}(n) \propto n^{-\alpha}$ yields $\alpha = 0.878$ ($R^2 = 0.9969$, approximate 95% CI $[0.843, 0.912]$) at $\kappa^2 = 180$ kHz and $\alpha = 0.893$ ($R^2 = 0.9968$, approximate 95% CI $[0.857, 0.928]$) at $\kappa^2 = 220$ kHz. These exponents summarize the observed transport-law scaling for the stated model class and do not, by themselves, establish a universal asymptotic form.

Boundary-toggle test (differential agreement under controlled perturbations)

To test whether boundary changes translate into consistent rate changes in the modeled system, we use a differential protocol in which controlled toggles perturb $Y_{\mathrm{in}}(\omega)$ while holding confounds fixed. For each toggle, we compare the predicted change in boundary exposure to an independently computed relaxation-rate shift, and we summarize agreement using pre-specified statistics.

We compute a dissipative exposure at the operating frequency ω_{01} from the boundary object. A convenient scalar proxy is

$$g_{\mathrm{eff}}(\omega) = \mathrm{Re}\left[Y_{\mathrm{in}}^{(\mathrm{red})}(\omega)\right], \tag{19}$$

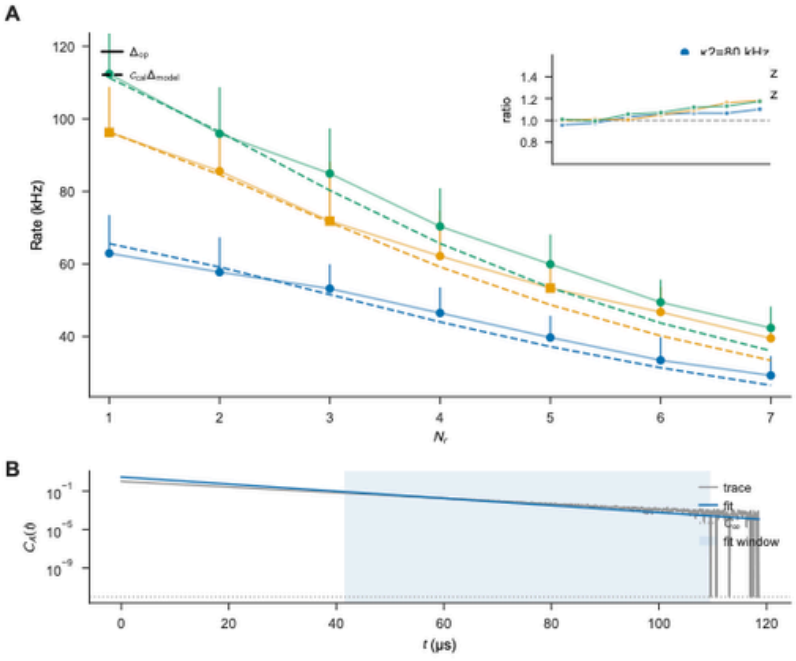

Figure 4: Transport scaling and floor diagnostics in the ring model. (A) Relaxation-rate proxy as a function of transport depth N_r, with model-based reference and ratios shown inset. (B) Example transient trace and fit window used to extract the reported rate proxy. Error bars denote the reported variability under the declared ensemble for the corresponding condition.

or, more generally,

$$g_{\mathrm{eff}}(\omega) = \mathbf{v}^{\dagger}\mathbf{G}(\omega)\mathbf{v}, \qquad \mathbf{G}(\omega) = \frac{\mathbf{Y}_{\mathrm{in}}(\omega) + \mathbf{Y}_{\mathrm{in}}^{\dagger}(\omega)}{2}, \tag{20}$$

where the reduction choice and coupling vector $\mathbf{v}$ are documented in Ξ_{BC} [12, 19].

We summarize toggle agreement using sign agreement, Pearson correlation ρ, and a robust regression slope ratio, evaluated over a documented toggle set. As a target for a mature predictivity test, we use $\rho \geq 0.7$, slope ratio in $[0.5, 2]$, and a recommended minimum of 10 toggle states to stabilize correlation and slope estimates. In this manuscript, Figure 5 reports a small pilot toggle set with 4 conditions, interpreted only as a directionality and internal-consistency check.

In a separate 10-condition toggle sweep, Pearson correlation between toggle-induced shifts was $\rho = 0.9965$ (bootstrap 95% CI $[0.9913, 0.9996]$), with perfect rank ordering ($\rho_s = 1.0$). A linear fit gives

$$\Delta_{\mathrm{ref}} \approx 0.7386\,\Delta_{\mathrm{abs}} + 167 \text{ Hz}. \tag{21}$$

The ratio satisfies

$$\Delta_{\mathrm{ref}}/\Delta_{\mathrm{abs}} = 0.7422 \pm 0.0263 \quad \text{(mean and s.d.; bootstrap 95\% CI for the mean } [0.7270, 0.7578]). \tag{22}$$

These statistics summarize a differential mapping between absorbing and reflecting variants under controlled toggles and support treating the toggle family as a reproducible perturbation set rather than uncontrolled scatter.

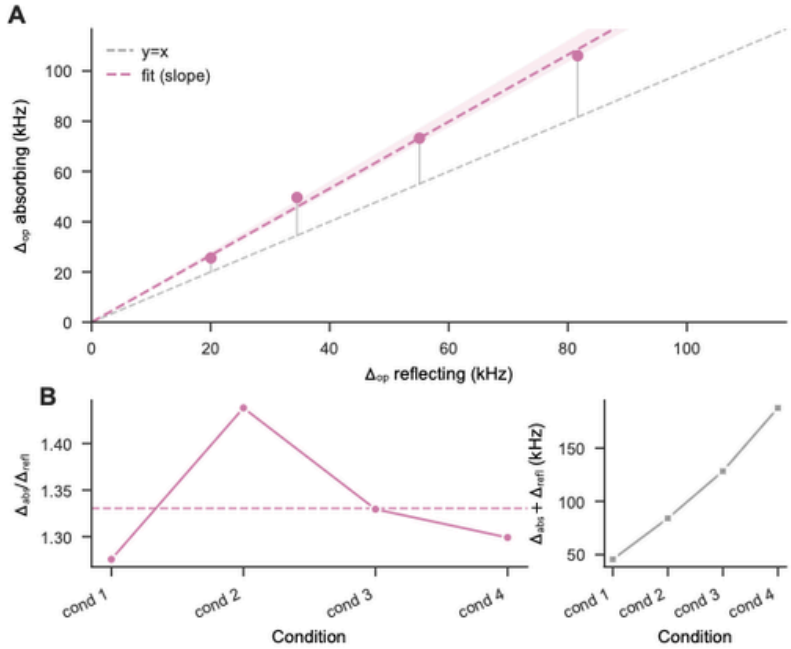

Figure 5: Differential boundary-toggle agreement (pilot visualization). (A) Agreement between exposure-predicted and independently observed rate shifts across toggle conditions, with the identity line shown for reference. (B) Summary statistics across conditions, including $\rho = 0.998$ for the pilot set ($n = 4$) and a fitted slope ratio. These tests are differential by design. They assess whether controlled boundary changes translate into consistent rate changes under a fixed coupling and decoding context.

Cycle-synchronous export as an effective instrument for decoding

A boundary description is operationally useful for QEC only if it constrains dynamics at the stabilizer-cycle timescale seen by the code and decoder. We therefore express boundary-driven dynamics as a cycle-synchronous quantum instrument indexed by a finite event alphabet aligned with record-based decoder inputs [13–15].

Let τ_{cyc} denote the stabilizer cycle duration and let ρ_k be the protected-patch state at the start of cycle k. We define the exported cycle instrument

$$\mathcal{I}_{\tau_{\mathrm{cyc}}} \equiv \{M_e\}_{e\in\mathcal{E}}, \qquad p_k(e) = \mathrm{Tr}[M_e(\rho_k)], \qquad \rho_{k+1}^{(e)} = \frac{M_e(\rho_k)}{p_k(e)}, \tag{23}$$

where each completely positive map M_e corresponds to an observable event e during the cycle, including drain clicks, erasure flags, or other witnessed boundary records.

The event alphabet $\mathcal{E}$ partitions into a located subset $\mathcal{E}_{\mathrm{loc}}$ and a residual, unlocated branch $e = \varnothing$. This yields the operational decomposition

$$\mathcal{N}_{\mathrm{BC}} = \underbrace{\sum_{e\in\mathcal{E}_{\mathrm{loc}}} M_e}_{\text{located evidence (decoder-visible)}} + \underbrace{M_{\varnothing}}_{\text{unlocated branch}}, \tag{24}$$

and motivates an erasure-dominant cycle-level description of the form

$$\mathcal{N}_{\mathrm{BC}} \approx (1 - p_E - p_P)\,\mathcal{I} \;+\; p_E\,\mathcal{E}_{\mathrm{loc}} \;+\; p_P\,\mathcal{P}_{\mathrm{unloc}} \;+\; \Delta. \tag{25}$$

Here p_E is the per-cycle weight of located, decoder-addressable damage, $\mathcal{E}_{\mathrm{loc}}$ is a located channel family indexed by the exported record e, p_P collects unlocated faults, and Δ is a defect term that captures boundary-model uncertainty, convention drift, and residual non-Markovianity [15, 38, 39].

All bounds are tied to a boundary artifact $Y_{\mathrm{in}}(\omega)$ and its record Ξ_{BC}. We parameterize deviations from the idealized cycle model using three quantities. They are a boundary-conditioning defect δ_{cloak} that captures mismatch between intended and realized boundary response over the operating band, a witness miss probability p_{fn}, and a non-Markov defect $\varepsilon_{\mathrm{nm}}$ that quantifies memory beyond the cycle abstraction at the chosen τ_{cyc}. Under passivity and realizability of the fitted macromodel $\widehat{Y}(s)$, and given validated bounds on $\delta_{\mathrm{cloak}}, p_{\mathrm{fn}}, \varepsilon_{\mathrm{nm}}$, the resulting defect Δ is bounded in the operational sense relevant to decoding and tail statistics (Methods and Supplementary Information) [26, 44, 46].

The witness-vs-no-witness cross-check above is consistent with limited effective located information under the current witness construction. In a negative-control predictivity export, the witness stream is dominated by identity behavior with identity fraction 0.9790, and the calibrated miss probability is $p_{\mathrm{fn}} = 0.5997$ while the false-positive probability is $p_{\mathrm{fp}} = 5.215 \times 10^{-3}$. These values are not used to claim erasure-grade sensing. They instead quantify why the current witness stream does not yet produce a strong improvement in the end-to-end logical benchmark and they define a concrete calibration target for hardware-relevant iterations.

Tail-risk criteria and boundary-toggle cross-check

Median improvements are not sufficient when failure is governed by rare, correlated excursions. We therefore treat tail behavior as a primary acceptance criterion and condition tail-safety statements

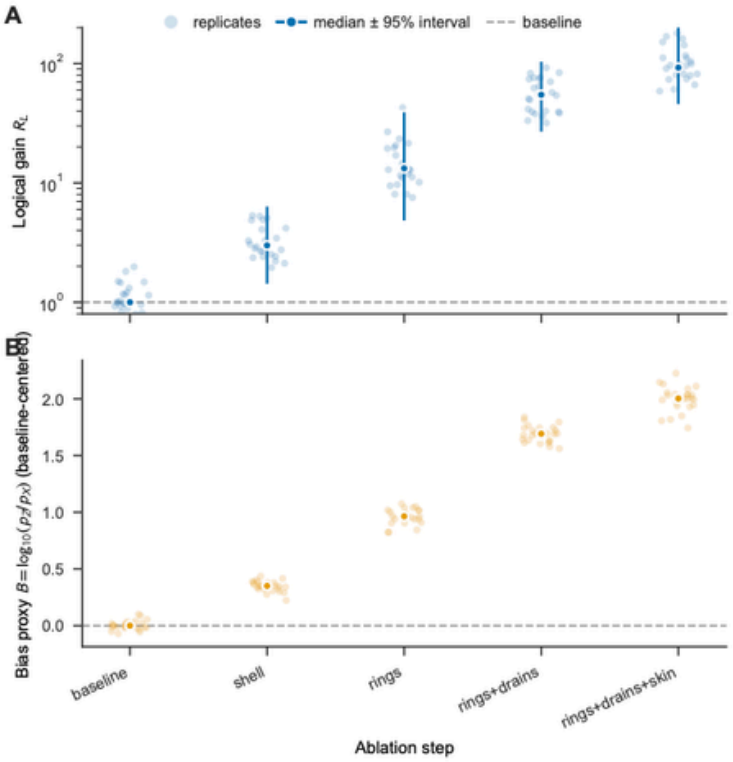

Figure 6: Decoder-level consequences of boundary conditioning and monitored evacuation. (A) Logical gain across ablation steps relative to a baseline configuration. (B) A bias proxy that tracks baseline-centered shifts in the inferred effective rate. Replicates and summary intervals are shown for each step, illustrating how transport, drains, and perimeter elements contribute in combination.

on statistics computed from first-passage and survival datasets [9, 10]. These criteria are defined to be recomputable from archived artifacts and to limit degrees of freedom in reporting [33, 35].

For a stated operating point and decoder of record, each run i yields a cycle index $N_i \in \mathbb{N}$ that is either the cycle to event (logical failure or a stated terminal event) or the cycle to censor. We denote by $\delta_i \in \{0, 1\}$ the event indicator and include with each record a configuration hash, decoder hash, boundary artifact hash, and random seed. This representation supports both first-passage analyses and right-censored survival estimation without reweighting.

Let $S(n) = \Pr(N > n)$ denote survival over stabilizer cycles. We estimate survival using the Kaplan Meier product-limit estimator [9]

$$\widehat{S}_{\mathrm{KM}}(n) = \prod_{t_j \le n} \left(1 - \frac{d_j}{r_j}\right), \tag{26}$$

where t_j are event times, d_j are event counts at t_j, and r_j is the risk set. We estimate cumulative hazard using the Nelson Aalen estimator [10, 36]

$$\widehat{H}_{\mathrm{NA}}(n) = \sum_{t_j \le n} \frac{d_j}{r_j}. \tag{27}$$

All tail criteria below are computed from the same first-passage dataset and are accompanied by the intermediate tables required for recomputation.

Tail statements at a given operating point are conditioned on four criteria. They are tail coverage, hazard flatness, tail ratio with bootstrap uncertainty, and clustering (extremal index). Failure of any criterion prevents tail qualification at that operating point.

Let $N_{\mathrm{evt}} = \sum_i \delta_i$ be the number of observed events. Let $n_{0.80}$ denote the 0.80-quantile of the event-time distribution, computed as an inverse Kaplan Meier quantile when censoring is non-trivial (stated rule), and otherwise as an empirical event-time quantile. Let N_{tail} be the number of events with $N_i \geq n_{0.80}$. The coverage criterion requires

$$N_{\mathrm{evt}} \geq 500, \qquad \frac{N_{\mathrm{tail}}}{N_{\mathrm{evt}}} \geq 0.10. \tag{28}$$

When this criterion fails, we report central-tendency metrics but do not interpret the dataset as supporting tail-qualified statements.

We define a region $[n_a, n_b]$, by default $[Q_{0.20}(N), Q_{0.80}(N)]$, using inverse Kaplan Meier quantiles when censoring is non-trivial, and we state the choice. Let $N_{\mathrm{evt,reg}}$ be the number of events with $N_i \in [n_a, n_b]$. We define the target event count per window as

$$w = \max(50,\ 0.02\, N_{\mathrm{evt,reg}})\,. \tag{29}$$

We partition $[n_a, n_b]$ into consecutive windows by accumulating events until each window contains at least w events (the final window may contain more). For window b, let d_b be the event count and define the exposure

$$E_b = \sum_{n \in \text{window } b} r_n, \qquad \widehat{h}_b = \frac{d_b}{E_b}, \tag{30}$$

where r_n is the Kaplan Meier risk set at cycle n. The hazard-flatness criterion requires at least 10 windows and

$$\mathrm{CV}\left(\widehat{h}_b\right) \leq 0.25. \tag{31}$$

The archived artifact bundle includes (n_a, n_b), w, the per-window table $(d_b, E_b, \widehat{h}_b)$, and a sensitivity report recomputed with w multiplied by 0.5 and 2.0.

We define the tail ratio

$$\mathrm{TR} = \frac{Q_{0.95}(N)}{Q_{0.50}(N)}. \tag{32}$$

Let the censoring fraction be $f_{\mathrm{cens}} = 1 - \frac{N_{\mathrm{evt}}}{N}$, where N is the number of runs. When $f_{\mathrm{cens}} > 0.10$, we compute quantiles via inverse Kaplan Meier. Otherwise, empirical event-time quantiles are permitted but must be stated. We compute a percentile bootstrap confidence interval using 2000 resamples, yielding $(\mathrm{TR}_{\mathrm{lo}}, \mathrm{TR}_{\mathrm{hi}})$. The tail-ratio criterion requires

$$\widehat{\mathrm{TR}} \leq 6, \qquad \mathrm{TR}_{\mathrm{hi}} \leq 6. \tag{33}$$

To test for burst-dominated failure, we estimate the extremal index θ on an extreme-event indicator stream derived from failures or stated high-severity proxies. Using the Ferro Segers

estimator [48], the clustering criterion requires

$$\widehat{\theta} \geq 0.60. \tag{34}$$

When $\widehat{\theta} < 0.60$, we do not interpret the operating point as tail-qualified because extreme events are too strongly clustered.

The boundary-toggle protocol, exposure definition, and agreement statistics are defined in Section . Predictivity statements are restricted to differential agreement under controlled toggles because $Y_{\mathrm{in}}(\omega)$ does not, by itself, determine absolute qubit rates without a validated platform mapping [12, 19].

Reduction statement

The decomposition in Eq. (25) is useful only if the defect term Δ can be bounded by measurable quantities and validated modeling assumptions. In BC-QEC, this bound is expressed in terms of a boundary artifact $Y_{\mathrm{in}}(\omega)$, its stable passive macromodel $\widehat{Y}(s)$, and a small set of parameters that quantify what the boundary model and the cycle abstraction do not capture [15, 26, 44].

The boundary object $Y_{\mathrm{in}}(\omega)$ supports quantitative claims only when it is consistent with a passive physical network at the specified reference plane. Operationally, this requires the Hermitian

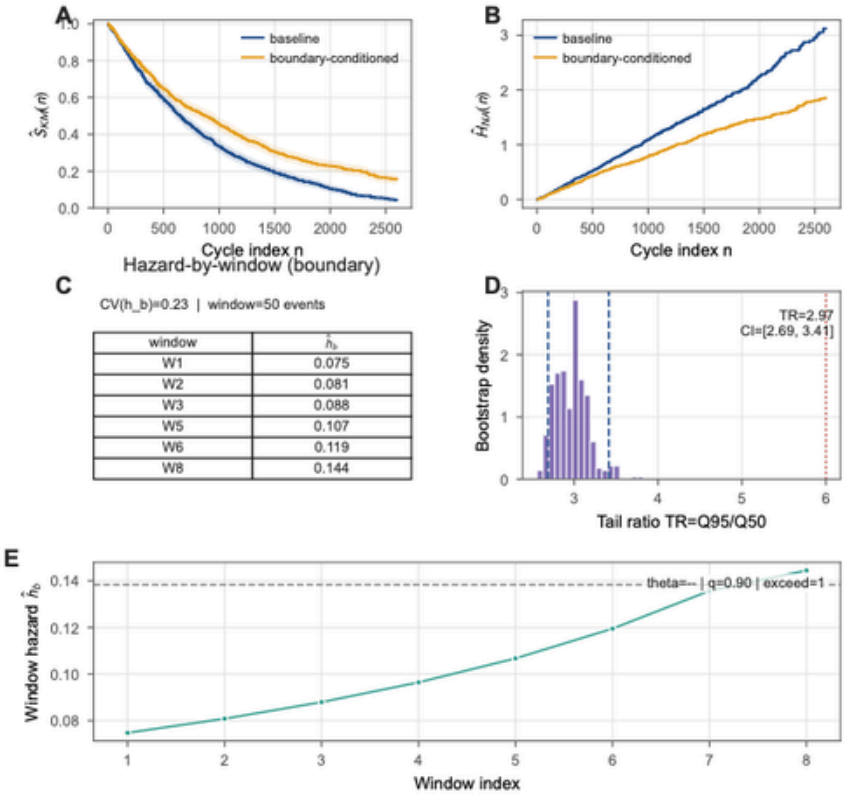

Figure 7: Tail criteria computed from first-passage datasets. PACER conditions tail statements on survival and hazard statistics and excludes operating points with insufficient tail coverage, non-flat hazard in event-count windows, excessive tail ratio under bootstrap uncertainty, or clustered extremes.

part

$$\boldsymbol{G}(\omega) = \frac{Y_{\mathrm{in}}(\omega) + Y_{\mathrm{in}}(\omega)^{\dagger}}{2} \tag{35}$$

to be positive semidefinite over the operating band, up to stated uncertainty, so that the boundary cannot inject net energy into the ports under the adopted wave convention. For rational models, passivity can be expressed through positive-real conditions [26, 46, 50]. In practice, we fit a stable rational macromodel $\widehat{Y}(s)$ and apply passivity enforcement when required, while reporting diagnostics that make any enforcement step reproducible [17, 18].

Let $\mathcal{N}_{\mathrm{BC}}$ denote the true per-cycle map on the protected patch induced by the physical boundary, including monitored drains and witness channels, and let $\mathcal{N}_{\mathrm{model}}$ denote the per-cycle map implied by the stated macromodel $\widehat{Y}(s)$ together with the stated coupling and cycle-time assumptions. Define the effective cycle model

$$\mathcal{N}_{\mathrm{eff}} \equiv (1 - p_E - p_P)\,\mathcal{I} \;+\; p_E\,\mathcal{E}_{\mathrm{loc}} \;+\; p_P\,\mathcal{P}_{\mathrm{unloc}}, \tag{36}$$

where $\mathcal{E}_{\mathrm{loc}}$ is indexed by the exported record $e \in \mathcal{E}_{\mathrm{loc}}$ and $\mathcal{P}_{\mathrm{unloc}}$ summarizes decoder-blind faults. The reduction statement asserts that, for a stated operating point and cycle time τ_{cyc}, there exists a choice of $\mathcal{E}_{\mathrm{loc}}$ and p_E consistent with the exported instrument $\mathcal{I}_{\tau_{\mathrm{cyc}}}$ such that

$$\mathcal{N}_{\mathrm{BC}} = \mathcal{N}_{\mathrm{eff}} + \Delta, \qquad \|\Delta\|_{\diamond} \le \epsilon_{\Delta}(\delta_{\mathrm{cloak}}, p_{\mathrm{fn}}, \varepsilon_{\mathrm{nm}}, \eta_{\mathrm{fit}})\,, \tag{37}$$

where $\|\cdot\|_{\diamond}$ is the diamond norm, η_{fit} is a stated bound on macromodel fit and enforcement residuals, and the principal defect parameters are the boundary-conditioning defect δ_{cloak}, the witness miss probability p_{fn}, and the non-Markov defect $\varepsilon_{\mathrm{nm}}$ [15, 44].

Equation (37) does not imply that $Y_{\mathrm{in}}(\omega)$ alone determines absolute qubit error rates. Translating a boundary admittance into transition rates requires platform-dependent participation factors and a validated coupling specification [12, 19]. For this reason, we restrict quantitative statements to exported cycle-level objects and their stability, differential agreement under boundary toggles, and tail qualification based on first-passage and survival criteria. Absolute-rate statements are made only when the coupling model and its uncertainty are explicitly stated and validated.

The derivation proceeds in four steps, each tied to an archived artifact. First, Ξ_{BC} yields a stable, passive macromodel $\widehat{Y}(s)$ that implies a dissipativity inequality for port variables [26, 50]. Second, the stated coupling maps this inequality to bounds on spectral exposure of relevant patch operators over a cycle, with residual uncertainty captured by η_{fit} and δ_{cloak} [12]. Third, witnessed drains provide an unravelling that partitions evolution into located records and an unrecorded branch, yielding the instrument in Eq. (23) and the split in Eq. (24) [30, 31]. Fourth, memory beyond the cycle abstraction is bounded by $\varepsilon_{\mathrm{nm}}$, constrained by dedicated diagnostics and required to remain below a stated threshold for the cycle-level description to apply [15, 42]. Together these steps produce an explicit function $\epsilon_{\Delta}(\cdot)$ and provide failure modes when its assumptions do not hold.

Reference architecture and future work toward a PACER-native processor

The quantitative statements in this manuscript are conditioned on boundary artifacts, exported cycle instruments, and the criteria defined above. A chip-level instantiation is not required to define

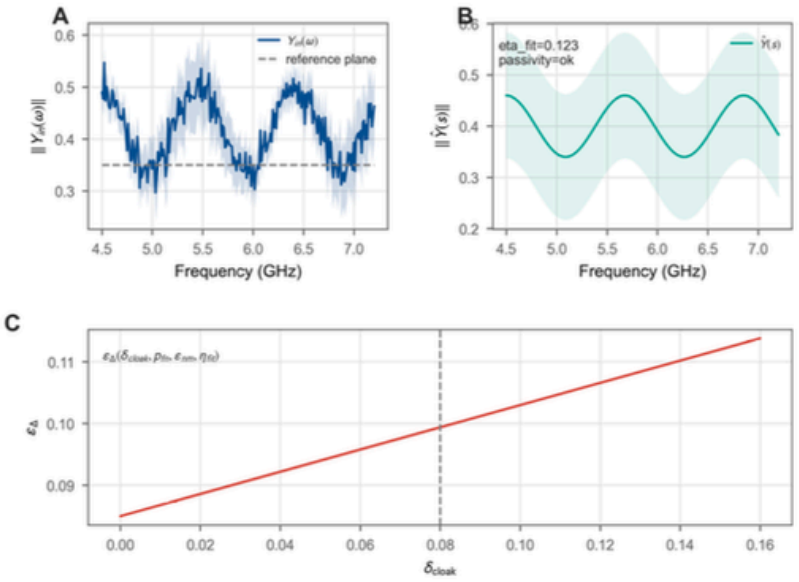

Figure 8: Main-text reduction statement. PACER conditions the cycle-level description on a stable, passive boundary macromodel and a stated coupling model, yielding a bounded unlocated remainder. Each arrow corresponds to an artifact included in the reproducibility bundle.

BC-QEC, but it is useful because it fixes the port map, witness semantics, and latency budget that give operational meaning to the exported records [4–6]. Here we outline a PACER-native pathway as a reference architecture and validation program. This section is non-claiming. It describes an implementation plan whose success would be evaluated using the same boundary and tail criteria already defined.

The reference design decomposes a PACER-native processor into three elements aligned with the cycle-level description in Eq. (25). The quiet-zone core implements the logical degrees of freedom and stabilizer-measurement circuitry, with coupling to the boundary mediated through specified ports. A perimeter layer converts boundary-adjacent faults into located records, such as dual-rail erasure flags and drain events, increasing the located fraction p_E while bounding the unlocated remainder [38, 39]. A surrounding transport region implements matched-loss evacuation over a specified frequency band and routes damaging excitations into monitored drains, making the instrument in Eq. (23) operational because records are produced with calibrated semantics and bounded latency [30, 31].

A central objective of a PACER-native design is to make the boundary object and cycle export a stable interface across hardware iterations. We therefore define an explicit boundary-to-decoder interface at a specified reference plane

$$\mathsf{ABI} \equiv \left(Y_{\mathrm{in}}(\omega),\ \Xi_{\mathrm{BC}},\ v_{\mathrm{BAP}},\ v_{\mathrm{EXP}},\ \mathcal{I}_{\tau_{\mathrm{cyc}}},\ \mathcal{E}_{\mathrm{sem}},\ \mathcal{L}_{\mathrm{lat}}\right), \tag{38}$$

where $\mathcal{E}_{\mathrm{sem}}$ is the event-semantics dictionary and $\mathcal{L}_{\mathrm{lat}}$ is the latency budget under which records are treated as cycle-synchronous. The purpose of Eq. (38) is to specify what is held fixed as hardware iterates. A reviewer should be able to recompute reported statistics and verify that performance changes track changes in boundary artifacts, exported event streams, and tail metrics rather than convention drift [7, 8, 49].

The pathway is naturally staged, with each stage producing a complete artifact bundle sufficient

for independent reanalysis [7, 8]. Stage 0 fabricates a boundary test coupon with the full port map and ring network but without an active qubit patch. Stage 1 integrates monitored drains and the erasure-conversion perimeter to establish event semantics and the latency budget. Stage 2 attaches a minimal protected patch and runs cycle-synchronous experiments sufficient to generate first-passage datasets.

Overhead and feasibility of witnessed drains and cycle-synchronous export

PACER introduces deliberate hardware overhead in exchange for a stricter, auditable cycle-level description. This overhead consists of additional evacuation hardware (rings and engineered drains) and additional observability hardware (witness sensors and per-cycle export), evaluated under a do-no-harm constraint on compute and readout operation.

Hardware overhead in simulation

In simulation we evaluate configurations with $P = 4$ monitored drains as a minimal instance and $P = 8$ as an expanded instance, alongside ring depth $N_r = 3$ to 5. This range fixes the routing footprint and readout-channel count needed to realize cycle-synchronous witness records at the

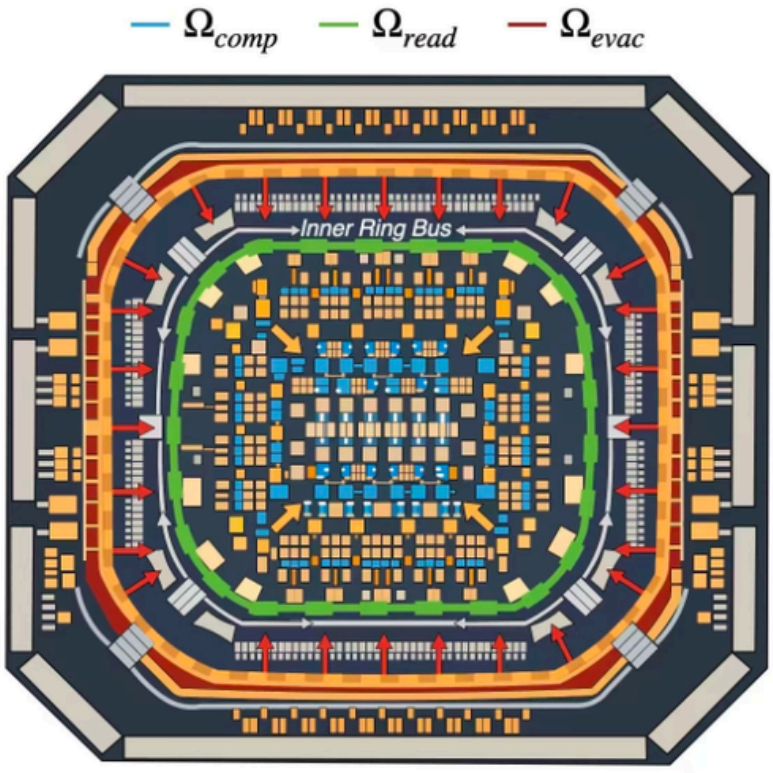

Figure 9: Reference pathway (not a demonstrated result). A PACER-native processor co-design fixes the port map, witness semantics, and latency budget required to instantiate the interface in Eq. (38). Validation is staged and is evaluated using the same boundary, toggle, and tail criteria that condition quantitative statements elsewhere in the manuscript.

interface plane.

Each monitored drain contributes at least one witness channel whose output is time-stamped and calibrated. Boundary erasure modules contribute additional flag channels. To first order the channel count scales linearly

$$N_{\text{witness}} \approx P + N_{\partial},$$

where N_{∂} is the number of independently flagged boundary sectors (or dual-rail modules) exported to the decoder.

We use the term *erasure* only when located information is calibrated with quantified false-negative and false-positive rates at a declared granularity. Otherwise the correct term is conditioning. Calibration therefore requires per-sector confusion-matrix estimates and drift monitoring at the stated operating point.

Sensors and readout feasibility

Witness sensing is feasible when each witness channel yields an event decision with controlled error rates and bounded latency at the stabilizer-cycle timescale. Practical realizations include drain-event sensing that detects evacuation into a monitored drain line via a dedicated receiver that declares an event when a calibrated statistic crosses threshold, and boundary erasure flags implemented by perimeter modules that map loss outside the codespace to a measurable signature producing a classical flag for decoding. The feasibility requirement is defined by the exported interface, meaning whether the decision is cycle-synchronous and whether $(p_{\text{fn}}, p_{\text{fp}})$ are bounded in the same conditions used for QEC evaluation.

Latency budget

We treat witness flags as cycle-synchronous only when they are available within the same stabilizer round

$$\tau_{\text{flag},k} \leq \tau_{\text{cyc}}, \qquad \forall k,$$

where the declared design instance pins $\tau_{\text{cyc}} = 1.0\,\mu\text{s}$. Time zero is defined as the controller trigger that begins the stabilizer round, and $\tau_{\text{flag},k}$ is measured to the time-stamp at which the flag becomes available to the decoding pipeline. Flags that arrive later are excluded from cycle-synchronous decoding for that round and are treated as delayed side information only if explicitly modeled as such.

Added loss budget (do-no-harm constraint)

The cloak, drains, and additional couplers can increase relaxation and dephasing in the compute and readout bands. We therefore declare an incremental-loss budget

$$\Delta\Gamma_{1,\text{PACER}} \equiv \Gamma_{1,\text{with PACER}} - \Gamma_{1,\text{baseline}}, \qquad \Delta\Gamma_{1,\text{PACER}} \leq 0.10\,\Gamma_{1,\text{baseline}} \quad \text{on } \Omega_{\text{comp}} \cup \Omega_{\text{ro}}.$$

Interpretation is that the PACER stack is allowed to degrade T_1 by at most 10% at the operating point. Beyond that, the configuration is not considered viable for the stated objective.

Cross-talk and spurious-mode risks and how they are bounded

Additional drain ports and witness lines increase routing density and opportunities for unintended coupling. They also enlarge the package electromagnetic complexity, which can create cavity or slotline modes that couple into qubits and readout buses. We therefore evaluate feasibility using two packaging constraints. First, no high-Q package eigenmode is allowed inside the compute or readout bands $\Omega_{\mathrm{comp}} \cup \Omega_{\mathrm{ro}}$. Second, out-of-band modes must have low participation on qubit electrodes and readout elements, enforced by geometry, bond-fence strategy, and absorber features. A concrete target is to ensure identified in-band package modes are absent or sufficiently damped (for example $Q_{\mathrm{pkg}} < 100$ in-band), and that out-of-band participation at qubit islands is at least 30 dB below the dominant intended mode (device-specific definition declared with the CAD report).

Discussion

Below threshold, quantum error correction is increasingly shaped by electromagnetic structure that originates at the boundary of a protected patch but is often treated in practice as an implicit bath. Boundary-mediated mechanisms can dominate device-specific behavior and the rare correlated excursions that set the tail of the logical-failure distribution. BC-QEC makes this boundary explicit by treating it as a specified, testable object and by pairing QEC-level statements with a cycle-synchronous exported instrument and validation criteria [13–15]. For this reason, boundary-to-rate claims are restricted to differential agreement under controlled toggles and to operating points that pass realizability, stability, and conformance checks [7, 8].

The central contribution is not a new decoder or a new stabilizer family. It is a reduction layer that translates packaging and near-field engineering into an effective, decoder-usable description. Equation (25) separates effective cycle noise into located evidence, an unlocated remainder, and a defect term bounded by stated uncertainty and modeling discrepancy [26, 44, 46]. This separation addresses a practical conflation. Improvements in coherence do not necessarily imply improved tail behavior, and architectural changes that increase located evidence can improve logical reliability even when mean error rates change only modestly [38, 39].

A second design choice is to treat tail behavior as a primary endpoint rather than a secondary diagnostic. In below-threshold experiments, logical performance can become sensitive to structured and temporally correlated mechanisms, making survival and hazard statistics more informative than a single mean rate [1, 2]. Accordingly, PACER conditions tail-qualified statements on first-passage or survival datasets without reweighting and on criteria that can be recomputed from archived artifacts [9, 10, 35].

The boundary object provides a disciplined route from microwave engineering to open-system dynamics. Passive realizability of a macromodel $\widehat{Y}(s)$ is a condition that the boundary description is physically meaningful under the stated conventions and does not implicitly encode an unphysical energy source [26]. Rate-level statements are then restricted to what can be validated. When platform-specific coupling declarations are incomplete, differential agreement under boundary toggles is the appropriate standard for linking boundary changes to rate changes [4, 19]. This guards against treating a network response as a unique predictor of qubit rates without stating the coupling assumptions that are required.

BC-QEC is adjacent to monitored-dissipation and reservoir-engineering approaches, where

engineered coupling and dissipation reshape dynamics and can yield directional transport or autonomous stabilization [23]. The distinguishing element here is the exported interface. Rather than treating boundary engineering as an implementation detail, PACER externalizes it as a versioned object with conformance tests and decoder-facing semantics [7, 49]. This aligns with best practice for reproducible computational and experimental science, where conventions and interfaces must be explicit for conclusions to be comparable across implementations [33, 34].

The framework has limitations. First, mapping a boundary admittance to absolute transition rates is platform dependent, and BC-QEC does not remove the need for stated coupling models and quantified uncertainty when absolute-rate statements are made [12]. Second, the cycle abstraction can fail when non-Markovian memory competes with the stabilizer-cycle timescale, and such memory must be treated as a first-class defect with explicit diagnostics [15, 42]. Third, witness channels introduce their own failure modes, including false negatives, false positives, and latency violations, which must be bounded for the located and unlocated partition to remain meaningful [38, 39].

A convincing hardware demonstration is therefore a staged program in which boundary declaration, differential toggle agreement, witness calibration, feasibility criteria, and tail criteria can all be evaluated at a stated operating point and decoder of record [7, 8].

Data availability

The data supporting the findings of this study will be deposited in a public repository upon publication, with persistent identifiers provided in the published article. The deposited release will include the minimum dataset required to interpret, reproduce, and extend the reported analyses. It will include raw multiport $S(\omega)$ measurements together with calibration metadata and the stated uncertainty model, the converted $Y_{\mathrm{in}}(\omega)$ at the specified reference plane and the complete boundary record Ξ_{BC}, the fitted macromodel $\widehat{Y}(s)$ with stability, passivity, and fit-residual diagnostics, boundary-toggle datasets linking exposure changes to measured rate shifts, and first-passage tables $\{(N_i, \delta_i)\}$ with configuration hashes and provenance fields sufficient to recompute survival, hazard, and clustering statistics [7, 8].

Code availability

The code supporting the findings of this study will be made available in a public repository upon publication and archived in a long-term record. The release will include the workflow used to process boundary measurements, fit passive macromodels, compute acceptance statistics, and generate manuscript figures, together with dependency lockfiles and scripts that reproduce the main analyses and figure panels from the deposited data [33–35]. Any restrictions will be stated explicitly. For any restricted components, we will provide a runnable substitute workflow for the analyses that underpin the central conclusions.

References

[1] Google Quantum AI et al. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 2023. doi: 10.1038/s41586-022-05434-1. URL `https://www.nature.com/articles/s41586-022-05434-1`.

[2] Rajeev Acharya et al. Quantum error correction below the surface code threshold. *Nature*, 2025. doi: 10.1038/s41586-024-08449-y. URL `https://www.nature.com/articles/s41586-024-08449-y`.

[3] Sergey Bravyi et al. High-threshold and low-overhead fault-tolerant quantum memory. *Nature*, 2024. doi: 10.1038/s41586-024-07107-7. URL `https://www.nature.com/articles/s41586-024-07107-7`.

[4] Sihao Huang et al. Microwave package design for superconducting quantum processors. *PRX Quantum*, 2:020306, 2021. doi: 10.1103/PRXQuantum.2.020306.

[5] T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf. Multilayer microwave integrated quantum circuits for scalable quantum computing. *npj Quantum Information*, 2:16002, 2016. doi: 10.1038/npjqi.2016.2.

[6] D. Rosenberg et al. 3d integrated superconducting qubits. *npj Quantum Information*, 3:42, 2017. doi: 10.1038/s41534-017-0044-0.

[7] Mark D. Wilkinson et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3:160018, 2016. doi: 10.1038/sdata.2016.18.

[8] Nature Portfolio. Reporting standards and availability of data, materials, code and protocols (nature portfolio editorial policy), 2025. URL `https://www.nature.com/nature-portfolio/editorial-policies/reporting-standards`. Accessed 31 Dec 2025.

[9] E. L. Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481, 1958. doi: 10.1080/01621459.1958.10501452.

[10] Wayne Nelson. Theory and applications of hazard plotting for censored failure data. *Technometrics*, 14(4):945–966, 1972. doi: 10.1080/00401706.1972.10488991.

[11] C. W. Gardiner and M. J. Collett. Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation. *Physical Review A*, 31(6):3761–3774, 1985. doi: 10.1103/PhysRevA.31.3761.

[12] Alexandre Blais, Arne L. Grimsmo, Steven M. Girvin, and Andreas Wallraff. Circuit quantum electrodynamics. *Reviews of Modern Physics*, 93(2):025005, 2021. doi: 10.1103/RevModPhys.93.025005.

[13] E. B. Davies and J. T. Lewis. An operational approach to quantum probability. *Communications in Mathematical Physics*, 17:239–260, 1970. doi: 10.1007/BF01647093.

[14] Karl Kraus. *States, Effects, and Operations: Fundamental Notions of Quantum Theory.* Springer, 1983.

[15] Felix A. Pollock, César Rodríguez-Rosario, Thomas Frauenheim, Mauro Paternostro, and Kavan Modi. Non-markovian quantum processes: Complete framework and efficient characterization. *Physical Review A*, 97(1):012127, 2018. doi: 10.1103/PhysRevA.97.012127.

[16] R. B. Marks. A multiline method of network analyzer calibration. *IEEE Transactions on Microwave Theory and Techniques*, 39(7):1205–1215, 1991. doi: 10.1109/22.85388.

[17] Bjørn Gustavsen and Adam Semlyen. Rational approximation of frequency domain responses by vector fitting. *IEEE Transactions on Power Delivery*, 14(3):1052–1061, 1999. doi: 10.1109/61.772353.

[18] Bjørn Gustavsen. Fast passivity enforcement for pole-residue models by perturbation of residue matrix eigenvalues. *IEEE Transactions on Power Delivery*, 23(4):2278–2285, 2008. doi: 10.1109/TPWRD.2008.2002659.

[19] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155–1208, 2010. doi: 10.1103/RevModPhys.82.1155.

[20] Kaneyuki Kurokawa. Power waves and the scattering matrix. *IEEE Transactions on Microwave Theory and Techniques*, 13(2):194–202, 1965. doi: 10.1109/TMTT.1965.1125964.

[21] Glenn F. Engen and Clarence A. Hoer. Thru-reflect-line: An improved technique for calibrating the dual six-port automatic network analyzer. *IEEE Transactions on Microwave Theory and Techniques*, 27(12):987–993, 1979. doi: 10.1109/TMTT.1979.1129778.

[22] Allan Lewandowski, Dylan F. Williams, Alex Dienstfrey, and Junfang Wang. Covariance-based uncertainties for S-parameters measured with a vector network analyzer. Technical report, National Institute of Standards and Technology (NIST), 2010. URL `https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=905041`.

[23] A. Metelmann and A. A. Clerk. Nonreciprocal photon transmission and amplification via reservoir engineering. *Physical Review X*, 5:021025, 2015. doi: 10.1103/PhysRevX.5.021025.

[24] Leonardo Ranzani and Jose Aumentado. Graph-based analysis of nonreciprocity in coupled-mode systems. *New Journal of Physics*, 17(2):023024, 2015. doi: 10.1088/1367-2630/17/2/023024.

[25] Alicia J. Kollár, Mattias Fitzpatrick, and Andrew A. Houck. Hyperbolic lattices in circuit quantum electrodynamics. *Nature*, 571:45–50, 2019. doi: 10.1038/s41586-019-1348-3.

[26] B. D. O. Anderson. A system theory criterion for positive real matrices. *SIAM Journal on Control*, 5(2):171–182, 1967. doi: 10.1137/0305011.

[27] Naomichi Hatano and David R. Nelson. Localization transitions in non-hermitian quantum mechanics. *Physical Review Letters*, 77(3):570–573, 1996. doi: 10.1103/PhysRevLett.77.570.

[28] Naomichi Hatano and David R. Nelson. Vortex pinning and non-hermitian quantum mechanics. *Physical Review B*, 56(14):8651–8673, 1997. doi: 10.1103/PhysRevB.56.8651.

[29] Emil J. Bergholtz, Jan Carl Budich, and Flore K. Kunst. Exceptional topology of non-hermitian systems. *Reviews of Modern Physics*, 93(1):015005, 2021. doi: 10.1103/RevModPhys.93.015005.

[30] Jean Dalibard, Yvan Castin, and Klaus Mølmer. Wave-function approach to dissipative processes in quantum optics. *Physical Review Letters*, 68(5):580–583, 1992. doi: 10.1103/PhysRevLett.68.580.

[31] Martin B. Plenio and Peter L. Knight. The quantum-jump approach to dissipative dynamics in quantum optics. *Reviews of Modern Physics*, 70(1):101–144, 1998. doi: 10.1103/RevModPhys.70.101.

[32] Jason P. Fine and Robert J. Gray. A proportional hazards model for the subdistribution of a competing risk. *Journal of the American Statistical Association*, 94(446):496–509, 1999. doi: 10.1080/01621459.1999.10474144.

[33] Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. Ten simple rules for reproducible computational research. *PLOS Computational Biology*, 9(10):e1003285, 2013. doi: 10.1371/journal.pcbi.1003285.

[34] Greg Wilson et al. Best practices for scientific computing. *PLOS Biology*, 12(1):e1001745, 2014. doi: 10.1371/journal.pbio.1001745.

[35] Roger D. Peng. Reproducible research in computational science. *Science*, 334(6060):1226–1227, 2011. doi: 10.1126/science.1213847.

[36] Odd O. Aalen. Nonparametric inference for a family of counting processes. *The Annals of Statistics*, 6(4):701–726, 1978. doi: 10.1214/aos/1176344247.

[37] D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–220, 1972. doi: 10.1111/j.2517-6161.1972.tb00899.x.

[38] Thomas M. Stace, Sean D. Barrett, and Andrew C. Doherty. Thresholds for topological codes in the presence of loss. *Physical Review Letters*, 102(20):200501, 2009. doi: 10.1103/PhysRevLett.102.200501.

[39] Henry Levine et al. Demonstrating a long-coherence dual-rail erasure qubit. *Physical Review X*, 14:011051, 2024. doi: 10.1103/PhysRevX.14.011051.

[40] Craig Gidney. Stim: a fast stabilizer circuit simulator. 2021. URL `https://arxiv.org/pdf/2103.02202`.

[41] J. Pablo Bonilla Ataides, David K. Tuckett, Stephen D. Bartlett, Steven T. Flammia, and Benjamin J. Brown. The XZZX surface code. *Nature Communications*, 12:2172, 2021. doi: 10.1038/s41467-021-22274-1.

[42] Yoshitaka Tanimura. Numerically "exact" approach to open quantum dynamics: The hierarchical equations of motion (heom). *The Journal of Chemical Physics*, 153(2):020901, 2020. doi: 10.1063/5.0011599.

[43] Akihito Ishizaki and Yoshitaka Tanimura. Quantum dynamics of system strongly coupled to low-temperature colored noise bath: Reduced hierarchy equations approach. *Journal of the Physical Society of Japan*, 74(12):3131–3134, 2005. doi: 10.1143/JPSJ.74.3131.

[44] Heinz-Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 2002.

PROPOSAL

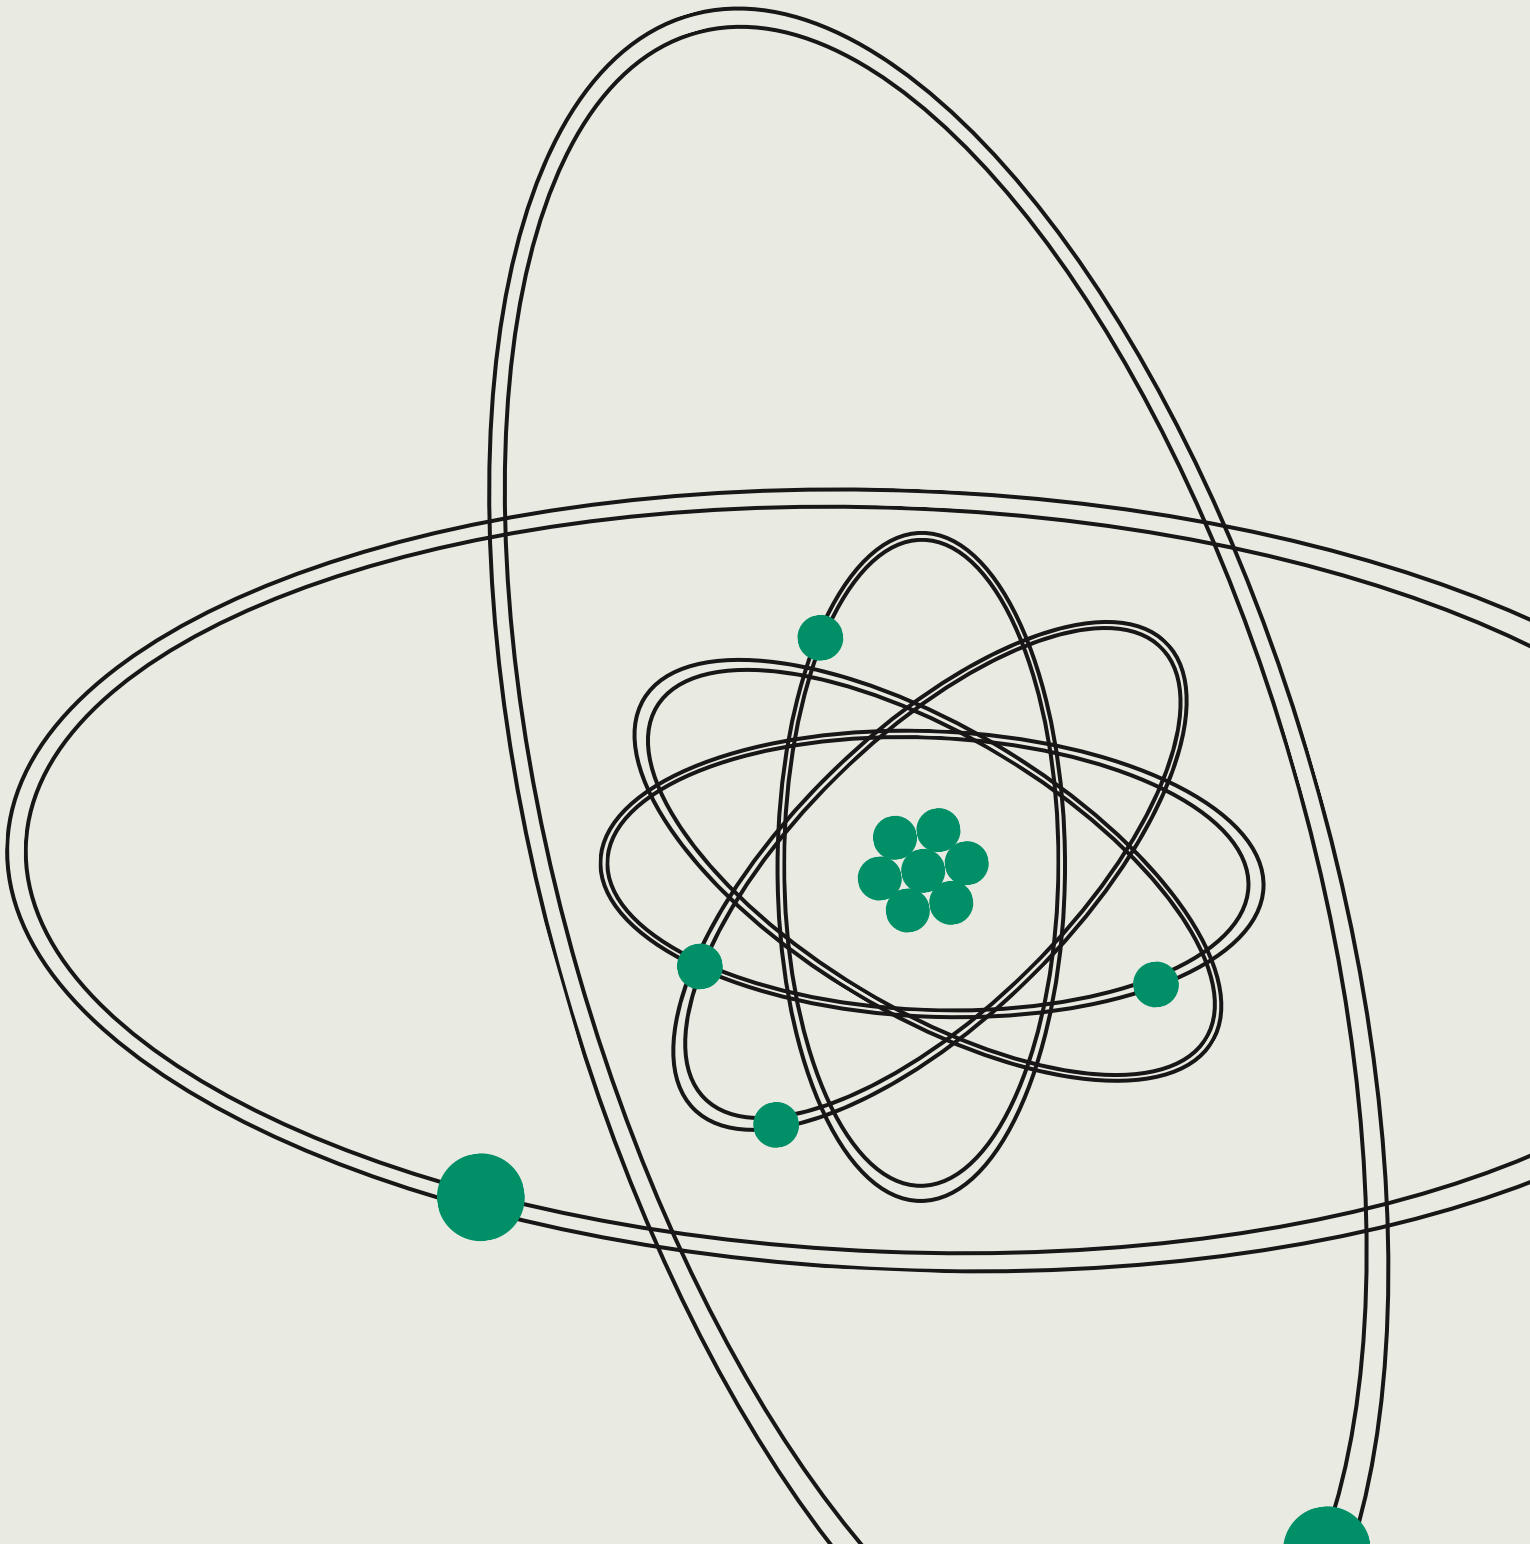

PACER Winter Research Proposal

PACER Quiet Certificate (21-day internship plan)

Eric Kim
Student researcher erickim1492@gmail.com CISH/NEST

For discussion with Dr. Hays

Abstract

PACER is a modular, measurement-first add-on that certifies a superconducting-qubit system's electromagnetic boundary condition. For WACQT at Chalmers, it provides fast, auditable screening of environment-facing microwave variants before committing cryogenic time. Over 21 days, I will deliver a PACER Quiet Certificate bundle: calibrated multiport $S(\omega)$ at a declared plane (Z_0, wave convention), converted to $Y_{\text{in}}(\omega)$, fit with a stable passive $\widehat{Y}(s)$ over $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$, and checked with non-regression gates (dissipation, reactive cross-coupling) plus sanity checks (passivity margin, fidelity, conditioning, grid stability, linearity). Optional cryogenic witness checks (T_1, T_2^*, readout) may be included if scheduling permits.

Day-21 outcome

Over 21 days, I will deliver a lab-reviewable certificate pipeline and validate it on one WACQT-relevant hardware variant: a small, connectorized boundary-condition module that can be swapped into an existing WACQT microwave path (examples include an inline filtering insert, a connectorized launch-to-launch coupon, or a shielding and return-path insert in a fixture segment). The resulting certificate bundle produces $\widehat{Y}(s)$ at a declared reference plane with repeat-derived uncertainty and a `manifest.json` that fully regenerates figures from raw VNA exports.

Day-4 deliverable. A full end-to-end run of the pipeline on one archived WACQT-style dataset (or a same-day lab capture), producing a draft certificate bundle with provisional thresholds and plots.

Decision rule. Given a reference configuration and a candidate variant, I will treat the variant as a pass only when the enforced model remains passive across the declared bands $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$ on a dense feature-capturing check grid, passivity enforcement does not hide modeling error beyond what repeat measurements justify, compute-band dissipative loading does not regress under conservative uncertainty bounds, reactive cross-coupling remains within conservative bounds, and the response remains small-signal linear over the probed regime.

Validity scope. The certificate is valid only for the declared reference plane, calibration, wave convention, port mapping, and instrumented coupling paths. If a coupling path is not represented by a port or fixture, it is out of scope by default and must not be implicitly assumed certified.

1. Proposed role and operating assumptions

During a 21-day on-site visit, my contributions are going to be: implement and document the certificate pipeline, assist with bench measurements under lab supervision, and prepare a small module compatible with the lab's existing measurement path.

In order to mimize risk, we must apply the following. Baseline operation is 2-port; anything beyond 2 is treated as stretch only. The certificate follows the lab's calibration, reference-plane, de-embedding, and repository conventions.Cryogenic access is not required for baseline success. Each issued certificate includes a short mapping from each port to a physical interface and states what is out of scope.

2. Objectives and scope

2.1. Primary objective

Produce a lab-reviewable and reproducible PACER Quiet Certificate for a small module or package variant, measured at room temperature. Concretely: measure calibrated $S(\omega)$ for a reference and test configuration with repeats; convert $S(\omega) \to Y_{\text{in}}(\omega)$ at a stated plane under a stated wave convention and Z_0; fit $\widehat{Y}(s)$, enforce passivity on $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$, and evaluate gates on a dense check grid that captures narrow features; then issue a certificate bundle that includes uncertainty, diagnostics, and a manifest that regenerates all plots.

2.2. Declared bands

The certificate declares bands, grids, and thresholds in `manifest.json`. Typical placeholders are shown here for concreteness:

$$\Omega_{\text{comp}}/2\pi \approx 4.5\,\text{GHz to } 5.5\,\text{GHz}, \qquad \Omega_{\text{ro}}/2\pi \approx 6.0\,\text{GHz to } 8.0\,\text{GHz}.$$

These values are replaced by the WACQT platform bands at issuance.

2.3. WACQT relevance (what this enables)

This workflow is intended to support WACQT superconducting-quantum-hardware iteration by turning environment-facing microwave variants into auditable models and conservative accept or reject gates. In practice, it aims to reduce iteration risk for components that influence gate fidelity and measurement stability, such as frequency-dependent loss channels, spurious resonances, and environment-mediated cross-coupling that presents as line-to-line crosstalk or unintended hybridization. Where relevant, the same certificate structure is compatible with quantifying coupling paths that are critical when integrating multi-physics modules (for example microwave-to-acoustic interfaces), because it makes the environment model explicit at a declared plane.

2.4. Non-goals

This plan does not claim that the certificate alone will improve gate fidelity during the visit, propose a new QEC code, assume a full design-to-cooldown cycle within 21 days, or require new switch matrices and bespoke calibration tooling.

3. Module definition and bench-first ablation plan

3.1. What the hardware is

The baseline hardware is a connectorized bench module that perturbs local microwave boundary conditions in a controlled way while minimizing mechanical iteration. It is compatible with WACQT's existing 2-port measurement path; the reference plane is declared at a lab-approved plane with de-embedding recorded; the form factor is a PCB coupon, fixture insert, filtering insert, or small shield and return-path insert that swaps without rework; and there is one physical knob with a small sweep range whose setting IDs are recorded in the manifest.

3.2. Coverage statement

Each certificate contains a short coverage statement: what each port corresponds to physically; what segment is being certified (for example launch-to-launch across an insert); and what the certificate is not asserting (for example couplings that require additional ports, global enclosure modes, or cryogenic-only mechanisms).

3.3. Minimum ablation matrix

Row	Definition
A (reference)	Variant removed, detuned, or set to WACQT baseline; measure $Y_{\text{in}}^{(0)}(\omega)$ with repeats.
B (test)	Variant engaged; measure $Y_{\text{in}}(\omega)$ with repeats; issue certificate and compute gates.
C (optional)	Repeat Row B for 2 to 3 knob settings to assess stability of margins versus setting.

This is an environment-first measurement: the boundary condition is the measured object, enabling fast iteration decisions without needing device-level attribution during the visit.

4. Certificate specifications

Certificate Spec (baseline 2-port)

Claim: Certify a configuration by a passive multiport admittance model $\widehat{Y}(s)$ and compute-band gate plots over declared bands $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$, referenced to a stated plane and supported by calibration metadata, repeatability, a feature-capturing check grid, and a model-fidelity screen.

Inputs: Calibrated $S(\omega)$ exports (Touchstone or WACQT standard) for reference and test, calibration and VNA settings summary, reference-plane and de-embedding record, repeats $N_{\text{rep}} \geq 3$, and a port-to-hardware coverage statement (with out-of-scope paths stated).

Method: Confirm small-signal linearity (power sweeps at two Ω_{comp} points), convert $S(\omega) \to Y_{\text{in}}(\omega)$ under declared wave convention and Z_0, fit a stable rational model, enforce passivity on $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$, and evaluate all gates on a dense check grid that includes measured points, $\geq 5\times$ densification per band, pole and narrow-feature insertion, and adaptive refinement until extrema stabilize within $\tau_{\text{grid}} = 0.05$ dB (logged in the manifest).

Outputs: Passivity and fidelity plots ($\lambda_{\min}(\mathbf{G})$, $\mathcal{D}_{\text{fit}}$, enforced vs. measured overlays), conversion and fit diagnostics, compute-band gate plots with repeat-derived uncertainty, and a `manifest.json` (hashes, calibration/plane metadata, wave convention and Z_0, grids, thresholds, uncertainty summaries, refinement log, pass/fail), plus exports of $\widehat{Y}(s)$ and reconstructed $Y_{\text{in}}(\omega)$ on the check grid.

4.1. Measurement recipe block

The certificate includes a short measurement recipe table. Defaults below are conservative and are expected to be overridden to match WACQT standards.

Item	Default used unless overridden in manifest
Band coverage	Cover $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$ with margin (example: 1 GHz to 12 GHz)
Frequency points	≥ 1601 points across band (or segmented sweeps)
IF bandwidth	≤ 3 kHz (or WACQT standard)
Averaging	≥ 4 (or enough to make repeatability stable)
Source power	Set to remain linear by Gate 3 power sweep (documented)
Repeats	$N_{\text{rep}} \geq 3$ per configuration, same reconnect procedure
Calibration	WACQT standard SOLT or TRL with date, kit, and plane recorded
De-embedding	WACQT standard, with reference plane named and stored

5. Certificate contract: waves, admittance, passivity, uncertainty, and gates

5.1. Wave convention and Z_0 contract

The certificate treats wave definition and reference impedance as part of the measurement contract. The manifest states the wave convention implied by the VNA export and records the per-port Z_0 used in export (ingesting Touchstone `[Reference]` if present, otherwise the option-line R). If renormalization is required, the certificate records the transformation and stores both as-measured and as-issued forms. If any of these items cannot be stated unambiguously, the certificate is not issued and the pipeline stops after prechecks.

5.2. Multiport admittance

$$Y_{\text{in}}(\omega) \in \mathbb{C}^{P\times P}, \qquad \omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}.$$

Split into conductance and susceptance:

$$\mathbf{G}(\omega) = \frac{Y_{\text{in}}(\omega) + Y_{\text{in}}^{\dagger}(\omega)}{2}, \qquad \mathbf{B}(\omega) = \frac{Y_{\text{in}}(\omega) - Y_{\text{in}}^{\dagger}(\omega)}{2\mathrm{i}}.$$

5.3. Conversion $S(\omega) \to Y_{\text{in}}(\omega)$

For diagonal $Z_0\ with Y_0 = Z_0^{-1}$,

$$Y_{\text{in}}(\omega) = Y_0^{1/2}\big(\mathbf{I} - S(\omega)\big)\big(\mathbf{I} + S(\omega)\big)^{-1}Y_0^{1/2}.$$

The pipeline logs conditioning diagnostics for $(\mathbf{I}+S)^{-1}$, flags near-singular frequencies, and records any regularization used. All acceptance checks are performed on the check grid (not only at measured points).

5.4. Uncertainty from repeats

For any quantity $Q(\omega)$ computed from repeats $Q_r(\omega)$, define pointwise median and dispersion:

$$\widetilde{Q}(\omega) = \text{median}_r\ Q_r(\omega), \qquad \sigma_Q(\omega) = 1.4826 \cdot \text{median}_r \left|Q_r(\omega) - \widetilde{Q}(\omega)\right|.$$

The certificate uses $\sigma_Q(\omega)$ to justify tolerances and to draw uncertainty bands.

5.5. Defaults and why they are conservative

Parameter	Default (override requires manifest justification)
Repeats per configuration	$N_{\text{rep}} \geq 3$
k_G (passivity tolerance multiplier)	4
k_S (Gate 1 conservatism multiplier)	4
k_B (Gate 2 regularizer multiplier)	6
k_F, k_A (Gate 2 uncertainty multipliers)	4, 4
$\mathcal{D}_{\text{fitmax}}$ (Gate 0b cap)	6
$\mathcal{S}_{\max\star}$ (Gate 1 target)	0 dB
Check-grid stability tolerance	$\tau_{\text{grid}} = 0.05$ dB

Defaults implement a safety-factor style conservatism because repeatability, reconnect variability, and calibration drift are typically non-Gaussian. Overrides are permitted, but the manifest must record the rationale and a minimal sensitivity note (whether pass/fail changes under default versus override).

5.6. Gate 0: Passivity (hard safety)

Quiet certification requires passivity of the enforced model on band:

$$\mathbf{G}_{\text{model}}(\omega) \succeq 0 \quad \forall \omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}.$$

Measured VNA data can show slight apparent non-passivity due to repeatability limits, so the certificate reports both measured and enforced margins:

$$m_{\text{meas}} = \inf_{\omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}} \lambda_{\min}\big(\mathbf{G}_{\text{meas}}(\omega)\big), \qquad m_{\text{model}} = \inf_{\omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}} \lambda_{\min}\big(\mathbf{G}_{\text{model}}(\omega)\big).$$

Define a measured-data tolerance τ_G from repeats:

$$\tau_G = k_G \cdot \text{median}_{\omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}}\ \sigma_\lambda(\omega), \qquad k_G = 4,$$

where $\sigma_\lambda(\omega)$ is the dispersion of $\lambda_{\min}(\mathbf{G}_{\text{meas},r}(\omega))$ across repeats. Checks:

$$m_{\text{meas}} \geq -\tau_G, \qquad m_{\text{model}} \geq 0$$

evaluated on the check grid.

5.7. Gate 0b: Model fidelity and enforcement distortion (fail fast)

Passivity enforcement must not hide large modeling error. Define a distortion metric on the check grid:

$$\mathcal{D}_{\text{fit}} = \sup_{\omega \in \Omega_{\text{comp}} \cup \Omega_{\text{ro}}} \frac{\|Y_{\text{inmodel}}(\omega) - Y_{\text{inmeas}}(\omega)\|_F}{\sigma_Y(\omega) + \epsilon},$$

where $\sigma_Y(\omega)$ is derived from repeat sweeps (dispersion definition declared in the manifest). Quiet certification requires:

$$\mathcal{D}_{\text{fit}} \leq \mathcal{D}_{\text{fitmax}}, \qquad \mathcal{D}_{\text{fitmax}} = 6.$$

If $\mathcal{D}_{\text{fit}}$ fails, the certificate is not issued and the pipeline reports the frequencies and matrix entries driving the failure.

5.8. Gate 1: Compute-band dissipative loading (worst-case non-regression)

Define the worst-case dissipative channel:

$$g_{\max}(\omega) = \lambda_{\max}\big(\mathbf{G}(\omega)\big), \qquad g^{(0)}_{\max}(\omega) = \lambda_{\max}\big(\mathbf{G}^{(0)}(\omega)\big).$$

Use a conservative lower-confidence bound:

$$\mathcal{S}^{-}_{\max}(\omega) = 10\log_{10}\left(\frac{\widetilde{g}^{(0)}_{\max}(\omega) + k_S\,\sigma_{g^{(0)}}(\omega)}{\widetilde{g}_{\max}(\omega) + k_S\,\sigma_{g}(\omega)}\right), \qquad k_S = 4.$$

Quiet certification requires:

$$\inf_{\omega\in\Omega_{\text{comp}}} \mathcal{S}^{-}_{\max}(\omega) \geq \mathcal{S}_{\max\star}, \qquad \mathcal{S}_{\max\star} = 0\ \text{dB}.$$

Gate 1 is a conservative screen against new dissipative eigenchannels presented by the environment. It bounds environmental loading and does not claim a specific device-limited T_1.

5.9. Gate 2: Reactive cross-coupling firewall (bounded coupling)

Define diagonal and off-diagonal parts of $\mathbf{B}$:

$$\mathbf{B}_{\text{diag}}(\omega) = \text{diag}\big(B_{11}(\omega), \ldots, B_{PP}(\omega)\big), \qquad \mathbf{B}_{\text{off}}(\omega) = \mathbf{B}(\omega) - \mathbf{B}_{\text{diag}}(\omega).$$

Use a regularized ratio plus an absolute cap:

$$\mathcal{F}_B(\omega) = \frac{\|\mathbf{B}_{\text{off}}(\omega)\|_2}{\|\mathbf{B}_{\text{diag}}(\omega)\|_2 + \delta_B}, \qquad \mathcal{A}_B(\omega) = \|\mathbf{B}_{\text{off}}(\omega)\|_2\,.$$

Regularizer from repeats:

$$\delta_B = k_B \cdot \text{median}_{\omega\in\Omega_{\text{comp}}}\ \sigma_A(\omega), \qquad k_B = 6,$$

$$\sigma_A(\omega) = 1.4826 \cdot \text{median}_r\,\big|\mathcal{A}_{Br}(\omega) - \widetilde{\mathcal{A}_B}(\omega)\big|.$$

Default bounds (reference plus repeatability, auditable).. Define reference-derived caps:

$$\epsilon_B = \sup_{\omega\in\Omega_{\text{comp}}} \widetilde{\mathcal{F}_B}^{(0)}(\omega) + k_F \cdot \text{median}_{\omega\in\Omega_{\text{comp}}}\,\sigma^{(0)}_{\mathcal{F}_B}(\omega), \qquad k_F = 4,$$

$$B_{\text{off,max}} = \sup_{\omega\in\Omega_{\text{comp}}} \widetilde{\mathcal{A}_B}^{(0)}(\omega) + k_A \cdot \text{median}_{\omega\in\Omega_{\text{comp}}}\,\sigma^{(0)}_{\mathcal{A}_B}(\omega), \qquad k_A = 4.$$

Then compare the test configuration using conservative upper bounds:

$$\mathcal{F}^{+}_B(\omega) = \widetilde{\mathcal{F}_B}(\omega) + k_F\,\sigma_{\mathcal{F}_B}(\omega), \qquad \mathcal{A}^{+}_B(\omega) = \widetilde{\mathcal{A}_B}(\omega) + k_A\,\sigma_{\mathcal{A}_B}(\omega).$$

Quiet certification requires:

$$\sup_{\omega\in\Omega_{\text{comp}}} \mathcal{F}^{+}_B(\omega) \leq \epsilon_B, \qquad \sup_{\omega\in\Omega_{\text{comp}}} \mathcal{A}^{+}_B(\omega) \leq B_{\text{off,max}}.$$

If the lab prefers absolute targets, ϵ_B and $B_{\text{off,max}}$ can be set directly and recorded in the manifest.

Interpretation.. Gate 2 bounds environment-mediated reactive coupling between ports within the certified segment. It is intentionally conservative and does not claim to isolate device-internal coupling.

5.10. Gate 3: Small-signal linearity

Gate 3 is a compact power-sweep at two representative frequencies in Ω_{comp}: $|S(\omega)|$ must remain within repeatability bounds and must not show gain onset. If strong power dependence appears, the configuration is declared out of scope for Quiet certification in this 21-day plan.

5.11. Failure-action mapping

Gate 0 fail. Recheck wave convention and Z_0 ingestion; verify calibration and reference plane; rerun conversion conditioning diagnostics; confirm check-grid feature insertion.

Gate 0b fail. Adjust model order or weighting; improve repeat quality; inspect near-singular $(\mathbf{I} + S)$; consider segmented fitting and pole insertion.

Gate 1 fail. Treat as dissipative regression; adjust the physical knob or boundary condition; repeat measurement with identical reconnect procedure; verify extrema stability under grid refinement.

Gate 2 fail. Treat as coupling regression; adjust shielding or return-path geometry; repeat measurement; compare delta plots and uncertainty envelopes.

Gate 3 fail. Reduce drive level, fix compression, or declare out of scope for this certificate version.

6. Modeling workflow

6.1. Rational fitting and passivity enforcement

The pipeline fits a stable rational model $\widehat{Y}(s)$ that reproduces measured $Y_{\text{in}}(\omega)$ over $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$ and exports a state-space realization. It (i) fits $Y_{\text{in}}(\omega)$ using a Vector Fitting style routine with a common stable pole set across the multiport admittance, (ii) selects the smallest order that satisfies Gate 0b without fitting repeatability noise, (iii) enforces passivity over $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$ via positive-real enforcement with minimal correction, and (iv) re-evaluates all gates on the check grid while logging refinement and extremum stability against τ_{grid}.

6.2. Fit quality metrics

The certificate reports relative Frobenius error on the check grid,

$$e(\omega) = \frac{\|Y_{\text{inmeas}}(\omega) - Y_{\text{inmodel}}(\omega)\|_F}{\|Y_{\text{inmeas}}(\omega)\|_F + \epsilon},$$

and may additionally report entrywise residual summaries for dominant entries (percentiles and worst-case). Acceptance checks are evaluated on the check grid, not only on measured points.

7. Methods and deliverables

7.1. Goal 1: certificate extraction on a 2-port path

Confirm WACQT calibration, wave convention, Z_0, and reference-plane conventions; measure calibrated $S(\omega)$ for reference and test with $N_{\text{rep}} \geq 3$ repeats; run Gate 3 power sweeps; convert

$S \to Y_{\text{in}}$ with conditioning diagnostics; construct and refine the check grid until extrema stabilize within τ_{grid}; fit $\widehat{Y}(s)$, enforce passivity on $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$, compute $\mathcal{D}_{\text{fit}}$, and evaluate all gates with conservative uncertainty bounds; then generate the certificate bundle and one-command regeneration recipe.

Goal 1 is successful if the enforced model is passive on $\Omega_{\text{comp}} \cup \Omega_{\text{ro}}$ on the check grid, $\mathcal{D}_{\text{fit}}$ meets declared bounds with auditable enforcement distortion, compute-band gates are reported with uncertainty and declared thresholds, and a WACQT lab member can regenerate the certificate from raw exports plus the manifest.

7.2. Goal 2: 4-port extraction, only if routine

If WACQT already has a routine 4-port path (for example via an established switch matrix and calibration), extend the certificate to 4 ports. Otherwise skip to avoid consuming lab time.

7.3. Goal 3: optional cryogenic witness checks

If a compatible device and schedule exist, run compact witness tests (for example T_1, T_2^*, and standard readout checks) under reference and test configurations and report them adjacent to certificate gates. These are explicitly non-critical to baseline success.

8. Software, data handling, and workflow fit

8.1. Pipeline outputs

The pipeline produces: standard review plots (passivity, $\mathcal{D}_{\text{fit}}$, fit error, compute-band gates); a manifest with hashes, calibration metadata, wave convention and Z_0, planes, grids, thresholds, uncertainty summaries, and pass/fail; a check-grid refinement log with extremum stability evidence; and exports of $\widehat{Y}(s)$ plus check-grid reconstructions for reuse in simulation or system models.

8.2. Workflow fit for WACQT

Ingestion follows WACQT's existing export formats, repository layout, and naming. The bundle is kept small enough for ordinary internal review. This workflow is aligned with WACQT iteration because it converts each environment-facing microwave variant into an auditable environment model plus consistent pass or fail gates that can be compared across versions.

8.3. Data handling

All raw and processed measurement data remain under WACQT's control and policies. Any external use of plots or notebooks is subject to PI approval and sanitization.

9. Risk management and fallback plan

9.1. Key risks and mitigations

The primary risks are reference-plane ambiguity and wave-convention ambiguity; both are mitigated by adopting WACQT standards and failing fast if the manifest cannot state the contract unambiguously. Multiport overhead is controlled by a 2-port baseline with higher port counts treated as stretch only. Missed narrow features are mitigated by the check-grid feature insertion and

adaptive refinement with logged extremum stability. Weak effect size still leaves value because a reproducible certificate constrains the design space and leaves behind reusable tooling. Overfitting or passivity enforcement masking error is controlled by Gate 0b and check-grid evaluation. Fridge scheduling risk is reduced by design since cryogenic work is optional.

9.2. Fallback

If only Goal 1 completes, WACQT still receives a working certificate pipeline plus a documented benchmark dataset and a clear method for comparing variants under passivity, model fidelity, dissipative non-regression, and cross-coupling constraints.

10. Timeline (21 days) and deliverables

10.1. Representative timeline

Pre-arrival (remote). Pipeline scaffolding, synthetic-data tests, alignment with WACQT file formats and calibration conventions.

Days 1 to 4. Confirm fixture, wave convention, Z_0, and reference plane; perform initial calibrated measurements with repeats; generate Day-4 draft certificate.

Days 5 to 10. Conversion, uncertainty estimation, conditioning checks, check-grid construction and refinement, and first stable plots.

Days 11 to 16. Vector fitting, passivity enforcement, Gate 0b fidelity screen, and gate reporting; optional knob sweep if appropriate.

Days 17 to 21. Lock certificate bundle, internal documentation, and optional cryogenic witness checks if scheduling permits.

10.2. Deliverables

1. Reference versus test dataset: calibrated $S(\omega)$ exports with repeats and recorded metadata.
2. PACER Quiet Certificate bundle: $\widehat{Y}(s)$, check-grid passivity check, Gate 0b fidelity screen, conversion conditioning diagnostics, fit diagnostics, gate plots with conservative uncertainty handling, check-grid refinement log, and a manifest.
3. Minimal documented pipeline that reproduces the certificate from raw exports plus the manifest.
4. (Stretch) 4-port extension if routine in WACQT.
5. (Stretch) Optional cryogenic witness summary adjacent to the certificate.

11. Summary

This proposal is conservative on WACQT lab overhead. The core deliverable is a measured, passive, and reproducible environment model expressed as multiport admittance at a stated reference plane, under a stated wave convention and Z_0 normalization, with compute-band non-regression gates that bound worst-case dissipation and reactive cross-coupling using conservative uncertainty bounds, plus sanity checks for passivity, model fidelity, conversion conditioning, check-grid stability, and small-signal behavior. Baseline success does not depend on cryogenic time or on demonstrating qubit improvements during the visit.

A. Context figures (illustrative only)

These plots are included only to make gate concepts and ablation logic concrete. They should be treated as illustrative and may be simulated or schematic unless explicitly labeled as measured with a certificate manifest. They are not required for Day-21 baseline success.

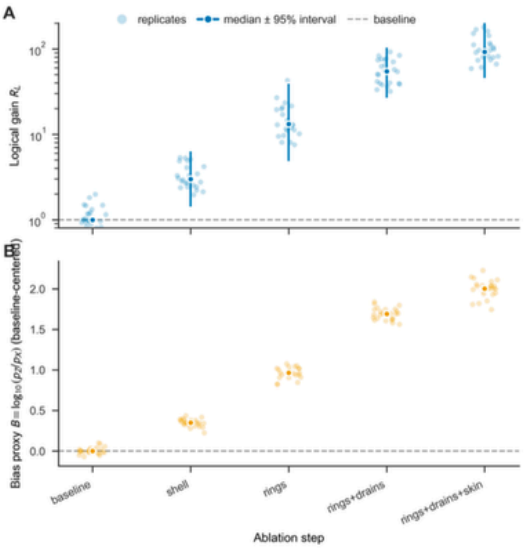

Figure 1: Ablation sweep across packaging steps (illustrative). Top: logical gain R_L (replicates with median and 95% interval; baseline shown). Bottom: bias proxy $B = \log_{10}(p_Z/p_X)$ (baseline-centered) across the same steps.

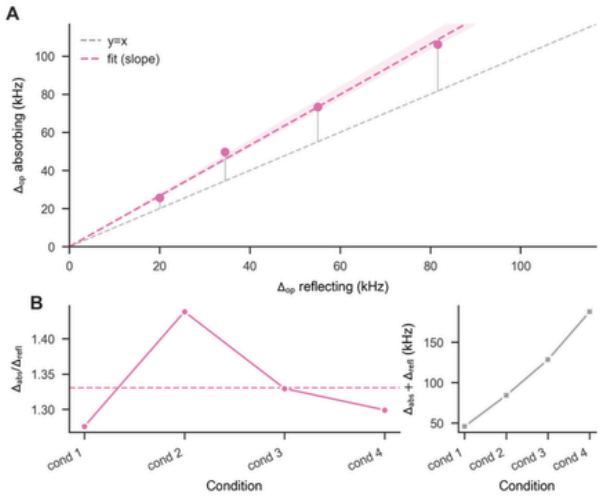

Figure 2: Absorbing versus reflecting decomposition sanity checks (illustrative). Panel A: Δ_{op} absorbing versus reflecting with a slope fit and $y = x$ reference. Panel B: condition summaries including $\Delta_{abs}/\Delta_{ref}$ and $\Delta_{abs} + \Delta_{ref}$.

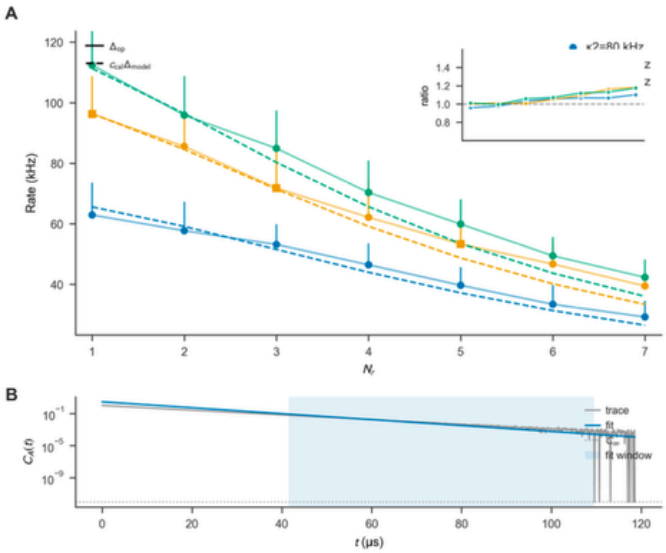

Figure 3: Scaling and memory-kernel diagnostics (illustrative). Panel A: rate versus ring count N_r. Panel B: correlation trace $C_\Delta(t)$ and fit window used to diagnose non-Markovian memory or effective decay structure.

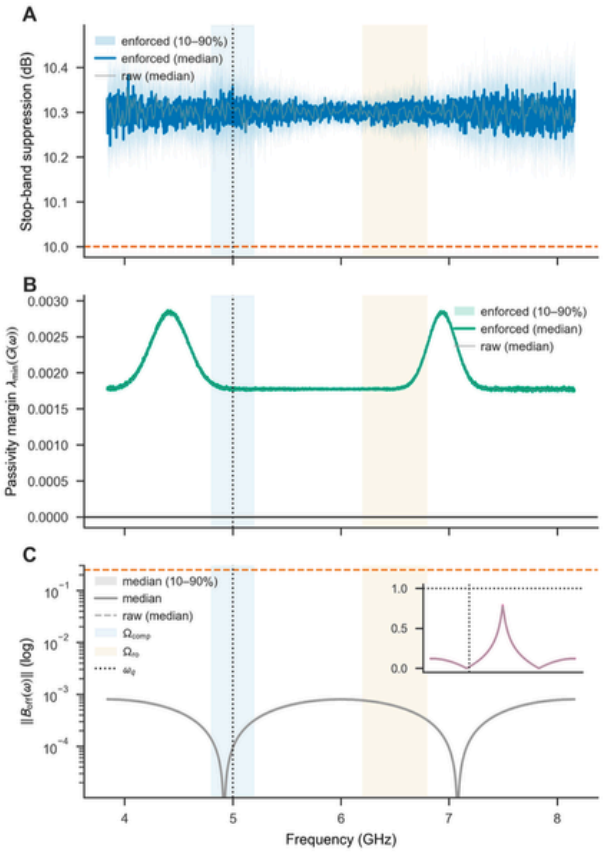

Figure 4: Certificate-style gate plots versus frequency (illustrative). Panel A: stop-band suppression (raw versus enforced, with uncertainty band). Panel B: passivity margin $\lambda_{\min}(\mathbf{G}(\omega))$. Panel C: reactive cross-coupling magnitude $\|\mathbf{B}_{\text{off}}(\omega)\|_2$ with band annotations and a qubit-frequency marker.

NOTEBOOK

2
Abstract/ Day 21 outcome

Understanding Linear Algebra

What is a vector space?

It is a set of vectors that can do 2 essential things: Addition & Scaling. If u & v are vectors in the space & k is a scalar, then $u+v$ is also in the space & ku / kv is also in the space

What are Eigenvectors & Eigenvalues?

Eigenvectors are lines where a matrix operation does not change the direction of the vector, only stretching or compressing the vector. Eigenvalues are the amount the vector is stretched or compressed in the Eigenvector.

A concrete introduction to Tensor Products

Z: vector space with basis $\{e_{v,w}\}$ for all $(v,w) \in V \times W$

$\{v,w\}$

↳ $\sum c(v,w)$ ⇐ where c is a constant

$V \oplus W$: $(v_1, w_1) + (v_2, w_2) = (v_1+v_2, w_1+w_2)$

$\lambda(v,w) = (\lambda v, \lambda w)$

Tensors has to:

1. Linear with respect to scalar multiplication
2. Distributive with respect to vector addition

Orthogonal vectors & subspaces

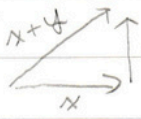

⇒ when dot product of $x \cdot y = 0$ ⇒ $x \cdot x + y \cdot y = (x+y) \cdot (x+y)$

$\|x\|^2 + \|y\|^2 = \|x+y\|^2$

$x^T x + y^T y = (x+y)^T (x+y) \Rightarrow 0 = x^T y$

Subspace S is orthogonal to Subspace T

means: Every vector in S is orthogonal to every vector in T

Basis for Subspace = Independent vectors that "span" the subspace

What is linear independence?

2 vectors a_1, a_2 span a plane ; spans: all different combinations

Independent : The only solution to $Av = 0$ is $v = 0$ (There is only 1 solution)

Dimensions : Number of basis vectors that "span" the subspace

Inner Products & Inner Product Spaces

Definition: An inner product on a real vector space V is a function that associates a real number $\langle u, v \rangle$ with each pair of vectors in V & satisfies the following axioms for all vectors u, v, & w in V & all scalars k

1) $\langle u, v \rangle = \langle v, u \rangle$ [symmetry axiom]
2) $\langle u+v, w \rangle = \langle u, w \rangle + \langle v, w \rangle$ [additivity axiom]
3) $\langle ku, v \rangle = k\langle u, v \rangle$ [homogenity axiom]
4) $\langle v, v \rangle \geq 0$ & $\langle v, v \rangle = 0$ if $v = 0$ [positivity axiom]

A real vector space V with an inner product is called a real inner product space

If V is a real inner product space, the length of vector v in V is denoted $\|v\|$

Hermitian Inner Product

$$\|u\| = \sqrt{a_1 \cdot \bar{a}_1 + a_2 \cdot \bar{a}_2 + a_3 \cdot \bar{a}_3} = \sqrt{\|a_1\|^2 + \|a_2\|^2 + \|a_3\|^2}$$

Properties of Hermitian Inner Product:

a) $\langle u, v \rangle = \langle v, u \rangle$ (conjugate communitiv)
b) $\langle av + bv, w \rangle = a\langle u, w \rangle + b\langle v, w \rangle$

3 blue 1 brown Essense of Linear Algebra

Chapter 1 : What is a vector?

There are 3 ways to define a vector,

- Physics : Magnitude & direction, location does not matter
- Computer Science : list of numbers
- Math : Simplification and Generalization of the 2 other definitions.

In linear algebra : vectors are rooted by the origin

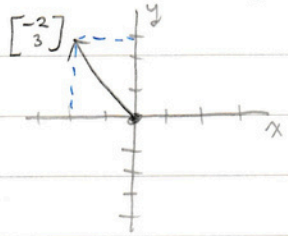

Vector addition : $\vec{v}$ ↗, $\vec{w}$ ↘ $\Rightarrow$ $\vec{v} + \vec{w}$ = 

Think that each vector is a set of instruction

↳ $\begin{bmatrix} x_1 \\ y_1 \end{bmatrix} + \begin{bmatrix} x_2 \\ y_2 \end{bmatrix} = \begin{bmatrix} x_1 + x_2 \\ y_1 + y_2 \end{bmatrix}$

Multiplication in vectors are called scaling

Chapter 2 : Linear Combinations, Span, and basis vectors

Unit vectors : $\hat{i}$ & $\hat{j}$; these are the vectors that are the "1" in the basis vector.

Adding 2 vectors are refered to linear combinations : $(a\hat{x} + b\hat{y})$

↳ The span is the set of all their linear combinations

Points are used to represent multiple vectors

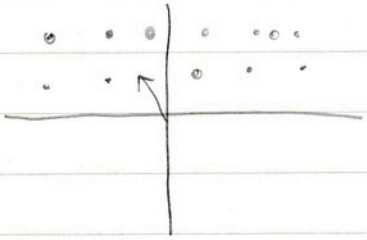

Challenge & Error Log

1.1 readout_guard.jsonl shows guard failing & simulation crashes
↳ Solution: Idos_ratio = 1 > max_Idos_ratio = 0.2, so the code thinks there is an error. Loosen the max boundaries

1.2 Wiring bug in _run_guard:
↳ Solution: the CLI looks for outdir/diagnostics so the guard always misses it, & falls back to notch defaults.

2.1 dependency error; mps not found (1)
↳ Solution: use WSL linux cloud not local device.

2.2 dependency error; mpx not found (2)
↳ solution: change from using pip to mamba

2.3 dependency error; mps not found (3)
↳ Solution: change from mamba to conda

2.4 dependency error; collision with dependency versions. (1)
↳ Solution: use a venv using python

2.5 dependency error; collision with dependency versions (2)
↳ Solution: use conda venv.

3.1 Local codebase not connected to github repo
↳ solution: make vscode linked to Gh

3.2 Authentication failed
↳ solution: use ssh key to allow login

4.1 Large files not connected
↳ Solution: Download Github LLF

4.2 LLF failed
↳ Solution: restarted application / redownload

Technical definition of Basis:

The basis of a vector space is a set of linearly independent vectors that span the full space.

Linear Transformations & Matrices – Chapter 3

Linear Transformation

function

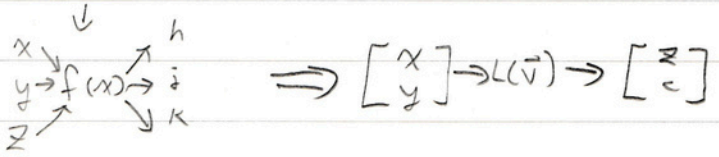

Linear transformation can be interpreted as the input vector moving to the output vector.

↳ You are moving the grid

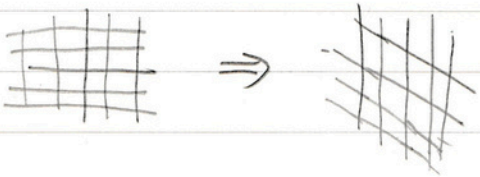

To calculate, you need to understand how the basis vectors $\hat{i}$ & $\hat{j}$ is moved.

{ if $\hat{i}$ & $\hat{j}$ is changed such that $\hat{i} \rightarrow \begin{bmatrix} 1 \\ -2 \end{bmatrix}$ and $\hat{j} \rightarrow \begin{bmatrix} 3 \\ 0 \end{bmatrix}$

$\begin{bmatrix} x \\ y \end{bmatrix} \rightarrow x\begin{bmatrix} 1 \\ -2 \end{bmatrix} + y\begin{bmatrix} 3 \\ 0 \end{bmatrix} = \begin{bmatrix} 1x+3y \\ -2x+0y \end{bmatrix}$ }

The 2×2 matrix:

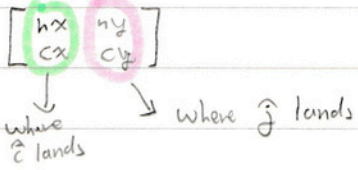

Linear Transformation is about looking at the change of the basis vector and multiplying the values we have to find.

5.1 Logical Reduction Failed during simulation
↳ missing modules / dependencies, can be fixed through using conda

6.1 Memory blowup: failed because matrix uses 5 PB of memory
↳ Brute simulation will not be viable for this research

7.1 Incorrect implementation of Kerr Cat Drain: Mathematically incorrect
↳ Reimplementation with 2 photon pump

8.1 Makefile error in base yaml
↳ rewrote base.yaml as a valid yaml document.
mirrored the correct structure the base.yaml in snapshot

9.1 Simulation fall back to BEST EFFORT from AUTHENTIC
↳ due to stubs, & missing implementation
ratio gate details "continue on fail" in reality this is a failed run despite results coming out.

10.1 Logical Reduction Failed:
↳ Added an explicit arm label argument.
Added a standard module guard
Overrided the BLAs helpers that were blocking LER
Reworked LER to properly work with Numpy & Scipy
Fixed / Added such that the code Outputted detailed errors

* Month 1 Progress *

[+24,194 -7265] lines of code
77 mb of code files
162 code files
20 recurring bugs

NEST Journal #1

This is my vision for NEST: a centralized society / club that generates the next talent in any industry imaginable, whether it is STEM related, arts & crafts related, or even related to sports, I want NEST to fast track people's dreams into actually happening. Personally, I will be teaching about academic research & the academic process, with the goal of becoming the highest ISEF sending school in China. Its not about the talent, its about the enviroment the kids are in, the environment that cements them into becoming giants in their respective fields. What if I make this google like workspace & have everyday meetings, projects, & classes in their?

This can be the system of NEST:

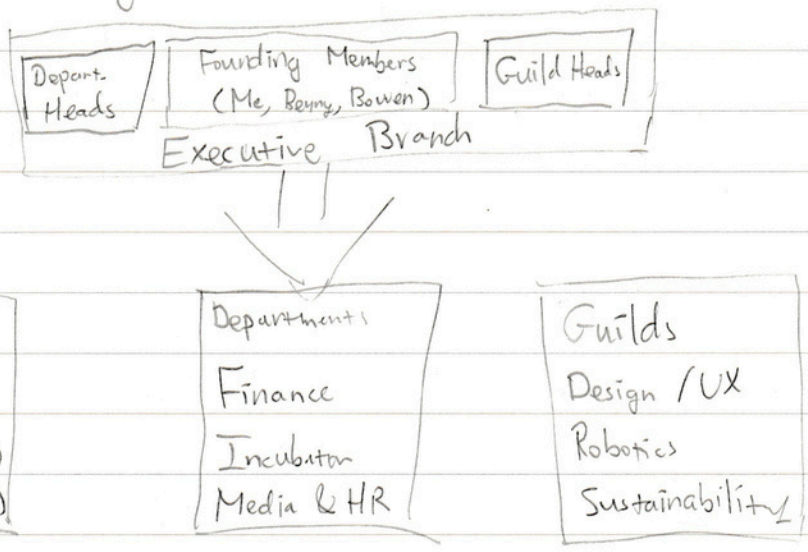

Finance Department Fundraises & Manages Funds : eg. Raised over ¥500,000 per year

Design / UX : Car Design, Concepts, Animations for Marketing

Media : Social Media, Newspaper, Podcast

NEST Journal #2 : Framework for Waveguide

Whilst browsing articles one day, I stumbled across a paper about seismic cloacking & metamaterials. That got me thinking, what if instead of reducing errors through systematic coding; equivalent to reinforcing the building to prepare for an earthquake; I create something to put on top of a pre-existing quantum that similar to a seismic cloak repels errors to benign places.

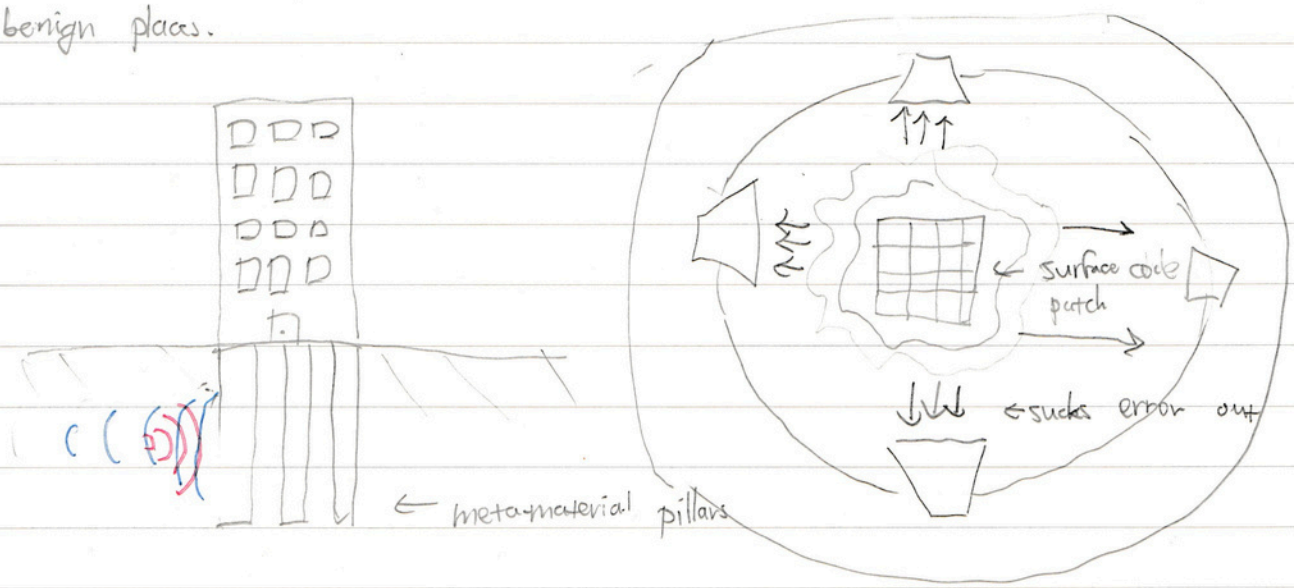

Challenges:

- I have to throw away my past expertise of QECC & I have to research materials, engineering, QM 1-3 and advanced techniques.
- What do I even know about making a QPU chip?
- Risk of all or nothing; possibility that this might not even work
- Feasibility: does civil engineering even apply to QC?

Next steps:

- Start literature review

NEST Journal #3 : Waveguide Literature

Whilst going through literature in preparation for waveguide, it seems there are multiple techniques that are out of the context of quantum error correction. While Im not sure about the feasibility of a QEC waveguide, it is evident that the seismic cloaking analogy is indeed possible.

CPW : coplanar waveguide rings are going to be the central piece of my research these are thin strips of superconducting film.

Going through Quantum Mechanics 3 topics is a struggle even after writing a paper on QECC. I seems non-hermitian quantum mechanics is required to support waveguide.

Multiple rings could block errors of certain "frequencies" or "types" and initial research suggests that from 7 rings we see diminishing returns

Because waveguide is the outer addon, implementation of waveguide seems that there will be minimal friction.

Kerr Cat Drains : these seem like the "drains", the seismic cloaking equivalent in quantum computing.

Maybe making a prototype "cypress" chip that demonstrates waveguide could be potentially made through an internship at a fab research center.

"Kerr Cat" seems to be a type of Qubit, maybe it has something to do with Cat Qubits?

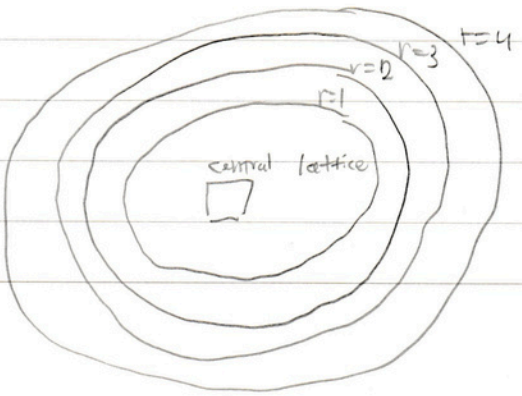

THE FIRST ISSUE OF NEST PUBLISHING

INNOVATING THE FIELD OF QPUS

Reliable quantum computers will not be limited by average noise but by rare, correlated events that slip past a decoder and silently accumulate into logical failure. PACER Breakwater targets that failure mode by treating the chip's near field environment as an engineered structure, like a seismic breakwater around a protected district, that redirects and dissipates damaging excitations outward into monitored drains instead of letting them reverberate in the core. The key shift is that the package boundary becomes a declared, measured object with passivity and stability guarantees, and the hardware exports cycle by cycle evidence that decoders can actually use, including calibrated located events and tail diagnostics rather than just improved median lifetimes. If this works, it changes how progress is judged, from unverified coherence gains to auditable reductions in unlocated faults and controlled long tail risk, which is the regime that matters for scaling error correction from demonstrations to dependable machines.

NEST ORIGINALS
11 November 2025
Vol .01, No. 001

CERTIFICATE C

TAIL AND BURST SAFETY

Issued by the PACER Breakwater Program

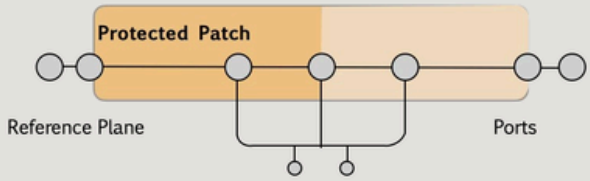

This Certificate was Awarded for the Rare Event Behavior Being Sampled, Bounded and Non Clustering Enough to Support a Tail Safety Statement

GATES Ω_{comp} Ω_{ro} Ω_{evac} Ω_{meas}

Gates	THRESHOLD
Gate C1: hazard flatness:	$CV(h_o) \le \eta_{haz}$
Gate C2: tail ratio:	$TR = \frac{Q_{0.95}(N)}{Q_{0.5}(N)}$
Gate C3: clustering:	external index $\theta \ge 0.6$