The Quantum Control Architecture: From a Fail-Closed Estimator to a Governed Measurement-and-Control Layer for Cryogenic Quantum Systems

Majid Hussain
Horos Labs Ltd, United Kingdom
info@horosengineering.com

Preprint — 2026 · v2

Abstract

A companion paper reported a fail-closed estimator — the Quantum Control Kernel (QCK) — that detects the precursor of systematic bias drift in cryogenic quantum control before actuator saturation, validated in simulation and against live, publicly accessible IBM hardware-calibration data. This paper reports what the estimator is *for*: the Quantum Control Architecture (QCA), a governed measurement-and-control layer in which the demonstrated kernel is the sensing floor of a larger construction. The construction is not free design. It is the quantum instantiation of a result established generically in the Finite Cognition series: a finite observer serving both a committing execution demand and a wide-field vigilance demand is forced to partition the two and re-couple them at the point of finalisation, so that no output is committed that vigilance has not verified. We show that quantum control information forces exactly this construction — telemetry is consumed by every layer of the control stack yet governed by none of them; observation of a quantum plant must be non-perturbative by discipline, not by hope; and an estimator that is fluent but ungoverned reproduces, on the control substrate, the same confident-and-blind failure the series identified generically. We then state the architecture that discharges the constraint, in its public units: a non-perturbative observer (NPO); the QCK detection floor; fail-closed gates that allow, abstain, or refuse; a state-coherence and equilibrium discipline over commits; provenance-checked memory that refuses known-failure patterns; a lifecycle and loop supervisor; a measurement plane of stability, coupling, margin, and cost features; and an integrity plane that audits the architecture's own calibration model and gates its portability to new hardware on provably leak-free relocation. We are explicit about maturity throughout: the kernel is demonstrated at Technology Readiness Level 4 under published bounds; the architecture beyond it is compile-ready geometry — derived, specified, and internally consistent, not validated as live middleware on partner hardware. The architecture is advisory by construction: the operator owns the plant at every depth.

1 Introduction: the kernel is not enough

The companion estimator paper [4] closed with a deliberate restraint: it reported a Technology Readiness Level 4 (TRL-4) detection layer, advisory and not a deployed controller, and it confined every claim to what simulation and read-only public calibration data had actually shown. That restraint invites a question this paper answers. A validated estimator, standing alone, is an instrument without a discipline of use. It reads control telemetry and emits verdicts; nothing in the estimator itself governs what is done with a verdict, remembers what became of the last one, notices when its own calibration assumptions have drifted, or refuses to carry its thresholds onto hardware where they were never established. Left ungoverned, a fluent estimator on a quantum control stack is a small instance of a failure mode identified generically elsewhere in this body of work: a committing channel producing well-formed output while structurally blind to its own error [3].

This paper reports the construction that closes that gap: the Quantum Control Architecture (QCA). The report is a compilation, not an invention. Every element is either (i) inherited from the generic results of the Finite Cognition series — the physical limits of finite observation [1], the conversion of the limit's shortfall into a located diagnostic signal [2], and the forced partition-and-coupling of contradictory computations [3] — or (ii) inherited from the demonstrated kernel [4], or (iii) stated explicitly as architecture: derived and specified, not yet exercised on hardware. The paper's contribution is the instantiation: showing that the generic constraints, applied to the specific substrate of cryogenic quantum control information, force this construction and not another.

2 The inherited discipline

Four results are taken as given.

The finite-observation bounds. No finite observer can represent a continuous environment without loss, estimate a continuous quantity exactly, retain its full history, be optimal across all environments, and certify its own consistency [1]. Under the modest correctness criterion adopted there — bounded error on every input the system acts upon — the bounds force a declared error budget, an explicit boundary of competence, and abstention outside it. The result is substrate-independent; the quantum estimator was named in that paper as one of its instances.

Deviation as located missingness. The shortfall the bounds impose is not merely a ceiling; under a controlled comparison of matched artefacts it becomes a localised signal — a direction of inquiry toward the uncaptured interaction [2]. In the present architecture this is the logic of every residual the integrity plane reads: a deviation between the calibration model and live telemetry is not noise to suppress but an address to investigate.

The forced partition and coupling. A finite system serving a committing, low-entropy execution demand and a non-committing, high-entropy vigilance demand cannot serve both on one undivided plane; it is forced to partition into parallel channels and then — because an isolated committing channel is fluent but blind — to couple them so that commitment is withheld until verification passes [3].

The demonstrated floor. A fail-closed estimator for precursor detection of systematic bias drift exists and has been characterised: in simulation, 96.7% of drift-collapse scenarios detected before saturation at a 12 s median lead with 0% false positives across 210 healthy runs over seven sensor-environment conditions; on real data, applied read-only to a public 156-qubit IBM Heron calibration history, it flagged the genuinely degraded qubits and one real relaxation-time decline (a single qubit's T_1 falling $88.0 \rightarrow 19.1\,\mu$s over 43 days) while raising no alarm on a chronically poor but stable qubit, nor on the 155 others [4]. All numbers in this paper concerning the kernel are those published numbers; none are extended here.

3 Why the substrate forces the construction

The generic argument of [3] becomes concrete on quantum control hardware in three steps.

3.1 Control information is consumed everywhere and governed nowhere

A cryogenic quantum system is operated through a stack — control electronics, calibration routines, error management, operating software — every layer of which consumes control telemetry: pulse amplitudes, calibration tables, relaxation and coherence times, readout fidelities. In the prevailing architecture no layer is charged with observing that information as a governed object: auditing where it came from, whether it is drifting, whether a decision taken on it is still within the competence of the model that produced it. The companion paper showed the consequence at the level of a single failure class: systematic bias integrates quietly in a feedback loop until the actuator saturates, and threshold alarms fire only after the correction authority is gone [4]. The generalisation is the point: *the information that predicts the failure is present in the telemetry, and nothing governed is listening to it.* That is a vigilance demand — wide-field, non-committing, spanning every channel of the stack — and it is precisely the demand the execution machinery of the control loop cannot serve without contaminating its own committing work [3]. The scale of that demand is set by the substrate's own physics: substantial temporal fluctuation of relaxation times, driven by two-level-system spectral diffusion, is a characterised feature of this hardware class [5]. An honest observer must therefore discriminate sustained systematic drift from that natural fluctuation background — firing on change while staying silent on ordinary wander, including on qubits that are chronically poor but stable — or it drowns the operator in alarms and is ignored.

3.2 Observation of a quantum plant must be non-perturbative by construction

On any substrate, an observer that writes into the process it observes corrupts its own evidence. On a quantum control stack the demand is sharper for a physical reason: the plant itself is measurement-sensitive, and the control loop's own actuation budget is the scarce resource the observer must not spend. The architecture therefore fixes, as its first structural commitment, an observer that holds to zero writes — it reads control and calibration state, computes, and returns, and has no path by which its observation re-enters the plant. The companion paper's real-data validation was deliberately conducted in exactly this posture: read-only, metadata only, no circuits executed [4]. What was there a validation constraint is here a permanent architectural law.

3.3 An ungoverned estimator reproduces the fluent-and-blind failure

Suppose the demonstrated kernel were wired directly to operators or automation with no governing layer. Every verdict would be committed at the kernel's full fluency, including: verdicts computed outside the regime where its error budget was established; verdicts on telemetry whose own provenance has silently changed; verdicts repeated indefinitely against a plant that has stopped responding (stiction); verdicts carried to new hardware whose calibration envelope the kernel has never seen. None of these are hypothetical exotica; they are the ordinary drift of deployed instruments. Each is an instance of the committing channel finalising what no vigilant channel has verified — the split-brain failure of [3], in engineered form, on the control substrate.

The forced resolution is the same as the generic one: partition the committing estimation from the wide-field vigilance, and couple them so that finalisation is gated.

4 The architecture

QCA is the construction that discharges these constraints. We state it in its public units, each with the constraint it answers. The depth vocabulary used below — D0, D1, D2 — is an evidence ladder, not a feature list: D0 is the demonstrated kernel; D1 is the coupled advisory loop, a bounded experiment target whose success criteria are fixed before any run; D2 is the full architecture reported here as compile-ready geometry.

4.1 Units

NPO — non-perturbative observer (observe). Reads process and control state without writing back into it; fires on change, stays silent on stable state. Answers the zero-write law of §3.2 and hosts the vigilance demand of §3.1. Shared as a discipline with the autonomous-domain architecture in this body of work; same doctrine, different wire.

QCK — quantum control kernel (detect). The demonstrated D0 floor [4]: fail-closed precursor and drift-anomaly detection on control and calibration telemetry, operating to a declared error budget with an explicit boundary of competence, abstaining outside it. The kernel is the instrument of the architecture in exactly the sense of [2]: it reads the deviation that locates the missing or varying interaction.

Fail-closed gates (discriminate: allow · abstain · refuse). Every advisory output passes a gate that admits it, abstains on insufficient evidence, or refuses it. The default under uncertainty is closed. This is the finalisation coupling of [3] made into a unit: the committing path cannot finalise what vigilance has not verified.

State coherence and equilibrium (hold · balance). The observer, the gates, and memory are held as one coherent runtime object across sessions and cycles, and commits to that state pass an equilibrium discipline — thresholded, locked, and accounted — so that state change is a governed event rather than a side effect. This answers the retention bound of [1]: a finite system that cannot keep its full history must govern what it keeps.

Provenance-checked memory (remember). What is retained carries its provenance and is checked against it; patterns that have already been seen to fail are refused rather than re-attempted. Memory is thereby fail-closed in the same sense as the gates: it admits, or it declines, and it does not silently fabricate continuity.

Lifecycle and loop supervisor (sustain · health). A lifecycle unit accounts for boot, continuous operation, and maturity of the running object. A loop supervisor watches the whole advisory loop for the failures a fluent loop cannot see in itself — stuck outputs, stiction against an unresponsive plant, verdicts repeating without re-evaluation — and forces re-evaluation or abstention.

4.2 The measurement plane (design intent)

Above the kernel's demonstrated floor, the architecture specifies a measurement plane: families of features computed from the limited telemetry streams the stack already produces. Publicly stated at the level of role, these are: dynamical-stability features reconstructed from limited scalar streams; separation of directed multi-channel coupling from noise floors; control-loop stability margins in pulse and reset feedback; and an irreversibility and cost budget on measurement and reset aggression against the thermal envelope. Two disciplines bound this plane. First, every numeric threshold in it is calibrated per hardware during a coupling window — no universal numeric bound is asserted, because the finite-observation result forbids pretending one exists [1]. Second, the plane's wording is exact: it performs precursor and drift-anomaly detection; it does not claim demonstrated pre-collapse prediction.

4.3 The integrity plane (design intent)

The architecture audits itself by the same law it applies to the plant — the self-referential bound of [1] made operational. Two units are specified. **Self-model honesty:** the architecture compares its own calibration model against live telemetry — prediction residuals, distribution divergence — and downgrades its own competence claim before it can mislead, re-calibrating or abstaining. The residual is treated exactly as [2] prescribes: located missingness, an address for inquiry, never noise to suppress. **Fail-closed portability:** transfer of the architecture to new hardware is certified only when the relocation provably leaks zero information — when what was established on the source hardware demonstrably survives the move — and otherwise the architecture refuses to port itself. Portability is a gate, not a promise; the same law is intended to carry across hardware modalities without ad-hoc rewrite, and until that is exercised it is stated as intent.

4.4 The advisory boundary

At every depth the outputs are advisory: verdicts, refusals, and evidence, delivered with session provenance. The operator — the laboratory, the platform, the plant owner — owns the response, and the actuators, throughout. This is not a commercial posture but an architectural one: the vigilance channel's authority is the authority to withhold finalisation, not to seize execution. Nothing in this paper claims, and the architecture does not contain, a production-control pathway over partner hardware.

5 Relation to the series, and to the kernel

The topology of this body of work is now explicit. The Finite Cognition series is the trunk, and it is generic: a limit on knowing [1], a signal within the limit [2], and a construction the limit forces [3]. Domain branches instantiate the trunk in pairs: an instrument paper that demonstrates the reading of the domain's deviation signal, and an architecture paper that reports the forced construction on that domain's substrate. On the quantum branch, the companion estimator paper [4] is the instrument; the present paper is the architecture. The branch is thereby complete: foundation shared with the trunk, instrument demonstrated, architecture derived. The evidence ladder separates what each rung may claim — the instrument's numbers are measured; the architecture's geometry is derived; and the coupling experiment between them (D1) is a target whose success criteria are fixed before the run, reported as such and not otherwise.

6 Maturity, scope, and what is not claimed

We are explicit about the standing of every element, in the register of the companion paper.

Demonstrated (TRL-4, under published bounds only). The QCK detection floor: the simulation characterisation and the read-only IBM Heron calibration validation, exactly as published [4]. Nothing in this paper adds to, extends, or re-interprets those numbers.

Derived. That the finite-observation bounds, the located-missingness instrument, and the forced partition-and-coupling, applied to quantum control information, force an architecture with the properties of §4: non-perturbative observation, a governed detection floor, fail-closed gating of finalisation, governed state and memory, self-audit, and gated portability.

Specified, not validated. The measurement plane and the integrity plane as bodies of computable features and audits; the full architecture as a coherent runtime. These are compile-ready geometry: internally consistent, buildable to, and stated at the level of principle and role — not validated as live middleware on any partner's hardware, not exercised in a dilution refrigerator, and not deployed in production anywhere.

Not claimed. That the architecture is running, or has run, as production control on any quantum hardware; that any coupling experiment (D1) has been performed — it has not, and when it is, its results will be reported against criteria fixed in advance; that the architecture predicts collapse seconds ahead as a demonstrated capability — the demonstrated capability is precursor and drift-anomaly detection under the published bounds; that any universal numeric threshold exists — thresholds are calibrated per hardware; that this work is peer-reviewed — it is an open preprint; or that any affiliation with or endorsement by IBM exists — the validation used their public data, read-only, and the method and conclusions are independent.

7 Conclusion

The companion paper demonstrated an instrument; this paper reports the machine it was always a part of. The finite-observation limit, carried through the series, does not stop at bounding what a control estimator can know: it dictates how a system that acts on such knowledge must be built — partitioned from its own committing fluency, coupled at finalisation, non-perturbative toward its plant, honest about its own calibration, and closed by default. The Quantum Control Architecture is that construction on the cryogenic control substrate: a demonstrated floor, a governed body, and an advisory boundary. The kernel measures; the architecture answers for what is done with the measurement. What remains between the two is not argument but experiment — the bounded coupling whose criteria are already fixed — and the discipline of this body of work is to report that experiment when it has been run, and not before.

References

[1] M. Hussain, "The Physical Limits of Finite Observation: Seven Bounds on Sensing, Inference, and Control," Horos Labs Ltd (2026). DOI: 10.5281/zenodo.21399897. [Finite Cognition series, Paper One.]

[2] M. Hussain, "Deviation as Located Missingness: Turning the Finite-Observation Bound into a Diagnostic Instrument," Horos Labs Ltd (2026). DOI: 10.5281/zenodo.21399930. [Finite Cognition series, Paper Two.]

[3] M. Hussain, "Forced Partition: Why a Finite Cognitive System Must Split Contradictory Computations and Re-Couple Them," Horos Labs Ltd (2026). DOI: 10.5281/zenodo.21399981. [Finite Cognition series, Paper Three.]

[4] M. Hussain, "Precursor Detection of Systematic Bias Drift in Cryogenic Quantum Control: A Fail-Closed Estimator Validated in Simulation and on Live Hardware Calibration," Horos Labs Ltd (2026). DOI: 10.5281/zenodo.21399999. [Quantum branch, instrument paper.]

[5] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, "Dynamics of superconducting qubit relaxation times," npj Quantum Information **8**, 132 (2022). arXiv:2105.15201.