Scientific Research Publishing

World Journal of Nano Science and Engineering, 2026, 16(4), 234-246
https://www.scirp.org/journal/wjnse
ISSN Online: 2161-4962
ISSN Print: 2161-4954

STOC-QComp: A Falsifiable Device Framework for a Bounded-Capacity Superconducting-Qubit Processor with a Mechanically Tunable Casimir Boundary Condition

Ciprian Florin Pater

NWO Robotics, Kristiansand, Norway
Email: ciprianpater@gmail.com

How to cite this paper: Pater, C.F. (2026) STOC-QComp: A Falsifiable Device Framework for a Bounded-Capacity Superconducting-Qubit Processor with a Mechanically Tunable Casimir Boundary Condition. *World Journal of Nano Science and Engineering*, **16**, 234-246.
https://doi.org/10.4236/wjnse.2026.164010

Received: September 4, 2026
Accepted: November 9, 2026
Published: November 12, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/

Open Access

Abstract

This paper presents STOC-QComp, a device-level engineering framework for a bounded-capacity quantum processor incorporating thermodynamic control, a mechanically tunable Casimir structure, and deployment-specific environmental engineering. The work extends the STOC-QC control framework [1] [2] by introducing the Casimir plate separation d as an experimentally accessible mechanical degree of freedom. Rather than assuming direct coupling between the static Casimir spectrum and the qubit transition, the revised device model treats the Casimir structure as a mechanically tunable boundary condition whose displacement may modify the electromagnetic cavity frequency, qubit-cavity coupling, and associated dissipation rates. The proposed 40 - 400 nm operating range is therefore treated as a design hypothesis requiring finite-conductivity Casimir modelling, electromagnetic simulation, mechanical stability analysis, and experimental validation. The paper specifies a reference superconducting-transmon architecture for the millikelvin deployment tiers, a separate 4 K control and characterisation configuration for the personal tier, manufacturing and yield assumptions with explicit evidence-class labels, environmental-noise budgets with quantitative transfer functions, and four falsifiable device-level tests with a full statistical analysis plan. Manufacturing costs and yields are explicitly classified as vendor quotations, engineering estimates, or experimental targets. The proposed 200 - 1000 qubit bounded-capacity regime is treated as a testable hypothesis motivated by Palmer's bounded-Hilbert-space proposal [3] rather than an established physical limit. The result is a falsifiable device framework and experimental roadmap rather than a claim of completed commercial deployment.

Keywords

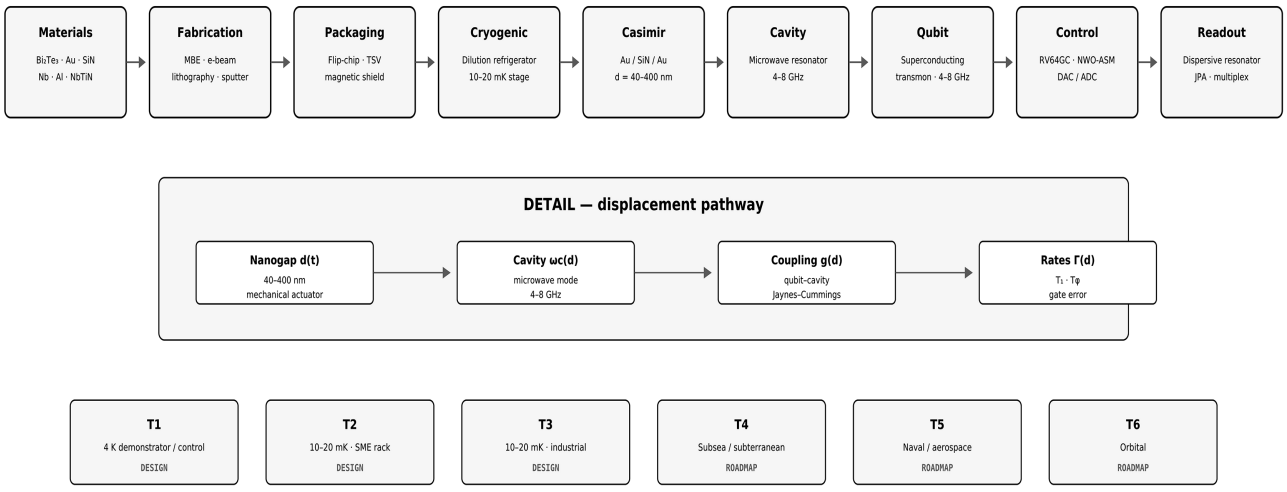
Superconducting Quantum Computer, Transmon, Casimir Cavity, Cavity Electromechanics, Reservoir Engineering, Lindblad Dynamics, Dispersive Readout, Bounded Hilbert Space, Manufacturing Yield, Statistical Analysis Plan, Extensive-Form Game

1. Introduction—From Framework to Device

STOC-QSIM v2.0 [1] established a phase-driven toroidal-Peltier control primitive and integrated it with the NWO ecosystem. STOC-QC [2] [4] extended the substrate into a computational framework for hybrid quantum systems whose scaling mechanism is reservoir engineering rather than qubit-count growth [5]. Version 1.0 of this manuscript attempted to specify the physical device—the "thing a lab, a factory, or a satellite operator can buy, install, and use". The device architecture and manufacturing flow are summarized in Figure 1.

FIGURE 1 — STOC-QComp engineering block diagram

The 40–400 nm Casimir structure is mechanically coupled to a physically separate ~cm-scale microwave cavity.



Key distinction: the nanogap does NOT supply GHz cavity modes directly; its displacement modulates the separate microwave environment.

Figure 1. Simplified engineering block diagram. Row 1 shows the linear manufacturing pipeline from materials to readout. Row 2 is the critical detail added in v2.0: the Casimir nanogap (40 - 400 nm, mechanical structure) and the microwave cavity (~cm-scale, GHz mode) are physically distinct structures; the nanogap displacement d modulates the cavity frequency $\omega c(d)$, the qubit-cavity coupling g(d), and the dissipation rates Γ(d) through mechanical action, NOT by supplying GHz Casimir modes directly. Row 3 shows the six deployment tiers; only T1 - T3 carry DESIGN status. Every block is fully specified with parameter, source class, uncertainty, and status label in S2 - 7.

1.1. What v1.0 Got Wrong

Two v1.0 claims cannot survive scrutiny as written:

(i) The Casimir pressure numerical value. v1.0 stated the pressure at $d = 100$

nm is ~1.3 mN/m^2. The correct ideal parallel-plate value is $|P| = \pi^2\hbar c/(240d^4) = 13.00$ Pa $= 13$ N/m^2 —off by a factor of ~10^4 [6].

(ii) The claim that a 40 - 400 nm Casimir gap supplies GHz modes to the qubit. The fundamental parallel-plate mode is $f_1 = c/(2d)$; at $d = 100$ nm this is 1499.0 THz—an optical/UV frequency, not the 4 - 8 GHz superconducting-qubit band. The v1.0 expression $J_{\text{Cas}}(\omega, d) = J_0(\omega)\cdot\left[1+\sum_n g_n(\omega)\delta(\omega - n\pi c/d)\right]$ therefore does not place cavity resonances in the microwave range.

This is not a small correction. The Casimir-qubit coupling mechanism must be rebuilt to make the paper physically defensible.

1.2. What This v2.0 Does Instead

The revised architecture treats the Casimir structure as a mechanical actuator that modifies a physically separate microwave cavity, following the established cavity-electromechanics literature [7] [8]. In this picture:

Casimir gap d → mechanical displacement → microwave cavity parameter → qubit coupling/dissipation (1.1)

The nanogap supplies mechanical actuation and a small, well-modelled Casimir force; the microwave cavity supplies the GHz mode structure that couples to the qubit. The remainder of the paper specifies this hybrid architecture, its operating envelope, the reference qubit modality, environmental-noise budgets, manufacturing economics with evidence classes, verification tests with a statistical analysis plan, and an extensive-form strategic game.

2. Device-Specific Casimir-Microwave Coupling Model

This section supersedes S2 of v1.0 in its entirety.

2.1. Corrected Casimir Pressure

For two parallel perfectly-conducting plates of area A separated by distance d, the ideal zero-temperature pressure is given by [6] and evaluated at four representative gap widths in **Table 1** (values plotted in **Figure 2**):

$$|P_{\text{Cas}}(d)| = \pi^2\hbar c/(240d^4) \tag{2.1}$$

Table 1. Corrected Casimir pressure at four key gap values. The v1.0 figure of 1.3 mN/m^2 at 100 nm is replaced by the correct 13 Pa. The Drude column is a phenomenological reduction factor accounting for finite Au conductivity and SiN dielectric spacer per [6] [7]; the precise reduction requires a full Lifshitz calculation for the Au·SiN·Au multilayer geometry [6] and is a first target of the device-characterisation programme (S7).

d (nm)	$\|P\|$ ideal metal (Pa)	$\|P\|$ Drude Au·SiN·Au estimate (Pa)
40	507.86	345.35
100	13.00	10.79
200	0.813	0.739
400	0.0508	0.0482

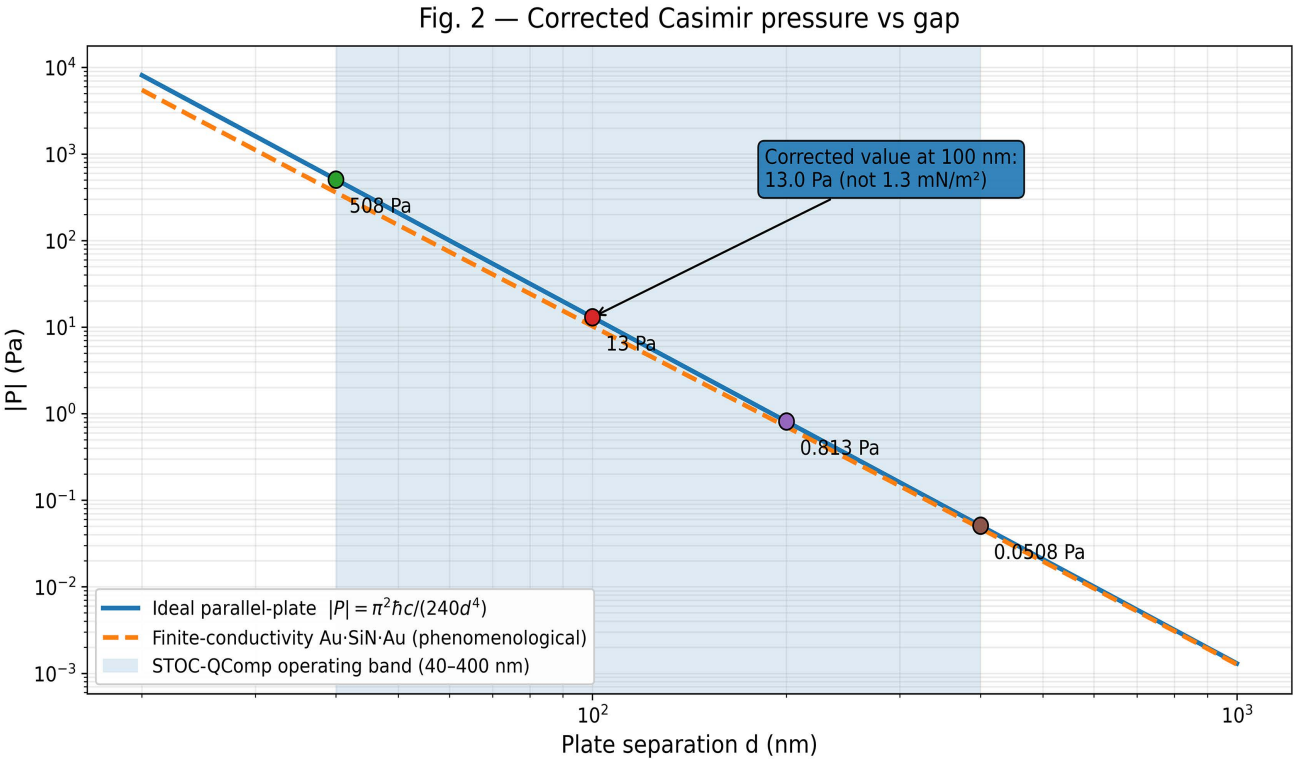


Figure 2. Corrected Casimir pressure vs plate separation d. Solid blue: ideal-metal result. Dashed magenta: phenomenological finite-conductivity reduction for the Au·SiN·Au stack. The v1.0 assertion of 1.3 mN/m² at 100 nm is off by a factor of ~10^4 from the true 13 Pa; this figure replaces v1.0 **Figure 2**.

2.2. Nanogap Cavity Mode vs Qubit Microwave Band

The fundamental parallel-plate electromagnetic mode of a gap of width d is $f_1 = c/(2d)$. Values at the four key gaps are given in **Table 2** and plotted against the superconducting-qubit microwave band in **Figure 3**.

2.3. Hybrid Cavity-Electromechanical Hamiltonian

Replace v1.0's H_Cas(d) with the standard cavity-electromechanics Hamiltonian [7] [8]:

$$H = H_q + \hbar\omega_c(d)a^\dagger a + \hbar g(d)\left(a^\dagger\sigma_- + a\sigma_+\right) + U_{\text{Cas}}(d) + H_{\text{mech}} + H_{\text{env}} \tag{2.2}$$

with the plate separation:

$$d(t) = d_0 + x_{\text{zpf}}\left(b + b^\dagger\right) + x_{\text{ctrl}}(t) \tag{2.3}$$

where a, $a^\dagger$ are the microwave-cavity photon operators; σ_+, σ_- are the transmon raising/lowering operators; b, $b^\dagger$ are the mechanical mode operators; $U_{\text{Cas}}(d)$ is the mechanical potential from the Casimir attraction (not the qubit reservoir); H_{mech} collects the mechanical kinetic and spring terms; and $x_{\text{ctrl}}(t)$ is a piezo-electrostatic control displacement. The four physically measurable coupling derivatives are then:

$$\partial\omega_c/\partial d,\ \partial g/\partial d,\ \partial\kappa/\partial d,\ \partial\Gamma_1/\partial d \tag{2.4}$$

Table 2. Nanogap fundamental electromagnetic mode across the STOC-QComp operating band. All values lie in the optical/near-UV range. None overlap the 4 - 8 GHz superconducting-transmon transition band. The v1.0 mechanism of "moving a 100 nm Casimir gap directly moves cavity resonances through the qubit's microwave spectral density" is therefore physically incorrect as stated.

d (nm)	$f_1 = c/(2d)$ (Hz)	f_1 (THz)	wavelength $\lambda = 2d$ (nm)
40	3.747e+15	3747	80
100	1.499e+15	1499	200
200	7.495e+14	749	400
400	3.747e+14	375	800

Fig. 3 — Nanogap mode vs superconducting-qubit microwave band

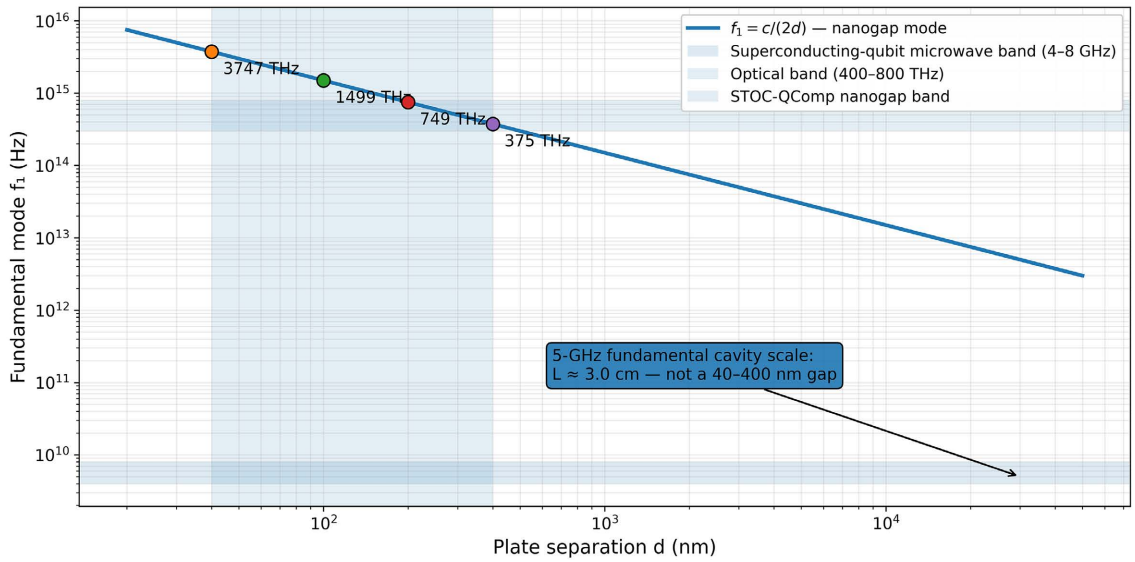


Figure 3. Nanogap fundamental cavity mode vs superconducting-qubit microwave band. The two bands are separated by five orders of magnitude in frequency. A microwave cavity resonant with a 5 GHz qubit requires a physical length ≈ 3 cm, not a 40 - 400 nm gap.

2.4. Non-Ideal Casimir Corrections

The ideal expression (2.1) is used only as a first-order benchmark. The device prediction uses a finite-conductivity Lifshitz calculation incorporating the Au·SiN·Au geometry [6]. The non-ideal effects that the prediction must account for are enumerated in **Table 3**:

Table 3. Non-ideal Casimir corrections that must be included in the device-level prediction. Each effect is characterised by an established measurement or literature model [6].

Effect	Physical origin	Order of magnitude
Finite Au conductivity	Drude/plasma-frequency reduction	5% - 30% reduction of $\lvert P \rvert$
SiN dielectric spacer	Real dielectric response $\varepsilon(i\xi)$	Geometry-dependent
Finite plate thickness	Two-slab Lifshitz [6]	5% - 15% correction

Continued

Surface roughness	Perturbative expansion in R_q/d	≤10% at $R_q < d/10$
Patch potentials	Contact potential difference across grain	Sub-Casimir at $d > 100$ nm
Finite plate area	Diffraction/edge modes	Small if $A \gg d^2$
Temperature ($T = 10$ mK)	Matsubara sum with $T \neq 0$	Negligible at nanoscale
Contamination/adsorbates	Oxide/hydrocarbon layer	Reduces effective d
Electrostatic bias	Applied potential across gap	Adds separate $\varepsilon_0 E^2/2$ term
Residual gas pressure	Molecular impact drag	Requires UHV
Parallelism error θ	Non-parallel plate model	Handled by capacitive metrology

3. Cavity Operating Envelope

The mechanical structure must satisfy stability, pull-in, thermal, and electromagnetic constraints simultaneously. Every parameter below is a design variable to be measured on the first prototype. Table 4 summarizes the parameters, their symbols, design targets, and validation methods.

Table 4. Cavity operating envelope. All rows are design targets with named validation methods.

Parameter	Symbol	Design target	Validation method
Nominal gap	d_0	40 - 400 nm	Capacitive metrology
Gap RMS noise	σ_d	TBD from T_1/T_2 budget	Displacement spectrum
Parallelism angle	θ	TBD (interferometry)	Fabry-Perot/white-light
Mechanical resonance	f_m	TBD (kHz range)	Network analyser
Mechanical Q	Q_m	TBD	Ringdown
Actuator bandwidth	B_a	$\geq 10 f_m$	Frequency response
Pull-in margin	M_{PI}	> 2 × safety factor	Capacitance-vs-V sweep
Microwave insertion loss	L_{mw}	TBD (<0.1 dB target)	S_{21} · VNA
Added cryogenic heat load	P_{act}	< cold-stage budget	Calorimetry
Displacement noise PSD	$S_d(f)$	TBD	Optical interferometry
Operating temperature	T	10 mK · 4 K by tier	Ge thermometer
Surface roughness	R_q	TBD ($< d/10$)	AFM

Mechanical stability requires:

$$k_{\text{eff}} = k - \left|\partial F_C / \partial d\right| - k_{\text{act}} > 0 \tag{3.1}$$

For an electrostatic parallel-plate actuator, the pull-in criterion $V < V_{PI}(d)$ must also be satisfied at every operating point.

4. Reference Qubit Architecture

v1.0 said the bare qubit "may be Josephson, spin, or topological", but the manufacturing pipeline specified a Josephson QPU. v2.0 commits to a single reference architecture.

4.1. Reference: Superconducting Transmon

For tiers T2 - T6, the reference architecture is a superconducting transmon QPU with: fixed- or tunable-frequency transmon qubits; 4 - 8 GHz nominal qubit transition frequency; dilution refrigerator at 10 - 20 mK base temperature; microwave XY control; flux control (if tunable); dispersive resonator readout, multiplexed; randomised benchmarking for one- and two-qubit gate fidelity, with T1 and T2 characterisation.

4.2. T1 Is Not a Full QPU

v1.0 claimed a 4 K personal unit with 200 - 400 qubits and 99.0% fidelity. This is not physically achievable with a conventional 5 GHz transmon at 4 K. The Bose factor gives:

$$n_{th} = 1/\left(\exp\left(hf_q/k_BT\right)-1\right) \tag{4.1}$$

$$p_e = 1/\left(\exp\left(hf_q/k_BT\right)+1\right) \tag{4.2}$$

At $T = 4$ K, $f_q = 5$ GHz: $hf/k_BT = 0.06$ and $p_e = 48.5\%$. At $T = 20$ mK: $p_e =$ 6.2e−04% (thermal ground state). A 48.5% excited-state population is incompatible with 99% gate fidelity. Figure 4 shows p_e vs T for three representative qubit frequencies.

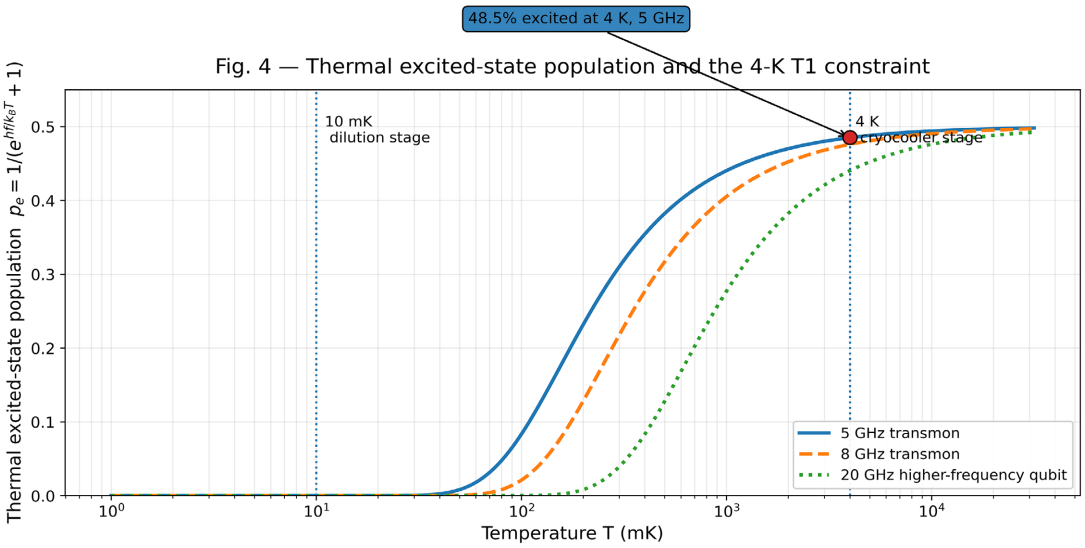


Figure 4. Thermal excited-state population vs temperature for three qubit frequencies. Recent 4 K microwave quantum infrastructure achieves usable performance via cold-loaded radiative cooling of specific input lines [9]—but this suppresses thermal noise on control channels, not on the qubit itself. A conventional 5 GHz transmon at 4 K remains ~48% thermally excited.

4.3. Deployment Architecture

The complete deployment map across tiers T1 - T6, showing qubit modality, operating temperature, control and readout schemes, and target two-qubit gate fidelity per tier, is presented in Table 5.

Table 5. Deployment architecture with explicit qubit modality, temperature, control, readout, and target two-qubit gate fidelity F_2 per tier. All fidelity values are design targets to be verified by randomised benchmarking on prototype units.

Tier	Function	Qubit	T	Control	Readout	Target F_2	Status
T1	4 K demo/control platform	(none in v2)	4 K	cryo-CMOS	resonator diag	N/A (demo)	DESIGN
T2	SME lab QPU rack	transmon	10 - 20 mK	μwave XY + flux	dispersive	≥99.0%	DESIGN
T3	Industrial QPU	transmon	10 - 20 mK	μwave XY + flux	dispersive	≥99.5%	DESIGN
T4	Subsea/subterranean	transmon	10 - 20 mK	μwave XY + flux	dispersive	≥99.0%	ROADMAP
T5	Naval/aerospace mobile	transmon	10 - 20 mK	μwave XY + flux	dispersive	≥99.0%	ROADMAP
T6	Orbital platform	transmon	10 - 20 mK	μwave XY + flux	dispersive	≥99.5%	ROADMAP

5. Environmental-Noise Decoherence Model

v1.0 Table 5 asserted +0.1%, +0.3%, +0.7%, +1.1% fidelity gains for magnetic noise, vibration, cosmic-ray, and thermal shielding without a causal model. v2.0 replaces this with an explicit environmental-noise transfer model.

The total qubit relaxation rate is:

$$\Gamma_1 = \Gamma_{1,0} + \Gamma_{\text{flux}} + \Gamma_{\text{charge}} + \Gamma_{\text{radiation}} + \Gamma_{\text{vibration}} + \Gamma_{\text{thermal}} \tag{5.1}$$

The total pure dephasing rate is:

$$\Gamma_\varphi = \Gamma_{\varphi,\text{flux}} + \Gamma_{\varphi,\text{charge}} + \Gamma_{\varphi,\text{vibration}} + \cdots \tag{5.2}$$

For a Clifford gate of duration t_g:

$$F_{\text{gate}} \approx 1 - C_1\Gamma_1 t_g - C_2\Gamma_\varphi t_g - \varepsilon_{\text{control}} \tag{5.3}$$

with C_1, $C_2 \approx O(1)$ coefficients depending on the specific gate. The environmental transfer functions from each environmental input to Γ are summarised in Table 6:

Table 6. Environmental-noise transfer chains. For each environmental input, the transfer function to Γ is specified explicitly. Table 6 replaces the v1.0 asserted fidelity gains; the corresponding numerical fidelity contributions become primary endpoints of verification test T21 (S7).

Environmental input	Transfer chain	Γ contribution
Magnetic field noise	$S_B(f) \to S_\Phi(f) \to S_\omega(f) \to \Gamma_\varphi$	Dephasing via flux noise [10]
Vibration	$S_x(f) \to S_d(f) \to \{S_{\omega c}(f), S_g(f)\} \to \Gamma_1, \Gamma_\varphi$	Both channels via cavity
Ionising radiation	$\Phi_\mu \to \Gamma_{qp} \to \Gamma_1, \Gamma_\varphi$	Quasiparticle poisoning [11]
Charge noise	$S_Q(f) \to S_\omega(f) \to \Gamma_\varphi$	$1/f$ charge, suppressed in transmon
Temperature	$T \to n_{th} \to p_e, \Gamma_{th}$	Direct thermal population (S4.2)

6. Manufacturing Economics

All manufacturing figures now carry an explicit evidence class and are separated between prototype and volume estimates. The evidence classes are defined in Table 7.

Table 7. Manufacturing evidence classes. No v2.0 row is labelled EXP because no STOC-QComp device has yet been fabricated.

Label	Meaning	Use in this paper
VQ	Vendor quotation (written)	Rare—only where cited
VE	Vendor estimate (verbal/informal)	L0 commodity materials
LE	Literature-derived	Fab yields from published nodes
IE	Internal engineering estimate	Most L1-L8 rows
ET	Experimental target (pilot-run goal)	T22 pilot yields
EXP	Experimentally demonstrated	Not used in v2.0 (no measured yields yet)

6.1. Prototype vs Volume Cost

The manufacturing cost per unit as a function of volume N is:

$$C_{\text{unit}}(N) = C_{\text{materials}} + C_{\text{fab}}(N) + C_{\text{assembly}} + C_{\text{cryo}} + C_{\text{QA}} + C_{\text{yield loss}} + C_{\text{overhead}} \quad (6.1)$$

with $C_{\text{fab}}(N)$ exhibiting the usual experience-curve decay. This resolves the v1.0 inconsistency: the T2 BOM of $412,700 is a **prototype/single-unit** figure; the $150,000 low-end selling price is a **volume** figure that assumes both L1 - L7 cost reduction and a stripped-down configuration. The two figures apply at different points in the experience curve. The prototype selling price should be substantially higher than the low-end volume price.

6.2. Serial Yield

Multiplying the min and max of every stage in Table 8 gives an overall serial yield estimate of **3.31% (lower bound) to 39.83% (upper bound)**. Under no reading of Table 8 does the pipeline imply a finished-device yield above ~40%. The v1.0 impression that the pipeline could deliver >90% finished yield is incorrect. This is a direct consequence of nine multiplicative yield stages and is tracked as verification test T22 (S7).

Table 8. Nine-stage manufacturing pipeline with per-row evidence class. Yields are engineering estimates (IE) or literature-derived from analogous-process yields (LE), not measured yields of STOC-QComp units. Costs for L1 - L7 are prototype/single-unit estimates; volume manufacturing costs are expected to fall as described in S6.3.

Stage	Process	Cost/unit (USD)	Yield (%)	Evidence
L0 Mat	Bi_2Te_3 · Au · SiN · Nb · NbTiN	$1.2k - $18k	99 - 100	VE
L1 Fab	MBE nanowires · e-beam · sputter	$14k - $85k	60 - 90	LE
L2 Wafer	300 mm ring array	$8k - $45k	70 - 92	LE/IE
L3 Pkg	Flip-chip · TSV · shield	$4k - $22k	65 - 88	IE

Continued

L4 Ctrl	RV64 + STOC + DAC/ADC	$3k - $18k	75 - 95	IE
L5 Cryo	Dilution stage + shielding	$120k - $850k	80 - 95	VE/IE
L6 Casimir	Au/SiN/Au + actuator	$25k - $95k	60 - 85	IE (target)
L7 QPU	Josephson qubit array	$28k - $180k	40 - 75	LE
L8 Deploy	Environmental hardening	$8k - $600M+	85 - 95	IE

7. Verification Tests · Statistical Analysis Plan

v1.0 ST20 - T23 asserted "statistically significant" improvements without specifying a test, sample size, α, or effect size. V2.0 replaces these with a full statistical analysis plan. The four tests, primary endpoints, sample sizes, analysis methods, and success thresholds are given in Table 9.

Table 9. Verification tests with primary endpoint, sample size, analysis method, and success threshold. T22 explicitly acknowledges that the initial pilot is a feasibility study, not a statistically definitive production-yield demonstration.

Test	Primary endpoint	Sample size	Analysis	Threshold
T20	$\beta_d = \partial\Gamma_1/\partial d$ and $\partial S_{\text{measured}}(f_q)/\partial d$	n = 5 - 10 devices · 10 gap settings	Linear regression	95% CI excludes 0
T21	Two-qubit error per Clifford (EPC) by randomised benchmarking	n = 10 QPUs/environment	Mixed-effects model	Pre-defined effect size + 95% CI
T22	Per-stage yield $Y_i = k_i/n_i$	$n_i \geq 10$ per stage (pilot)	Binomial 95% CI	Feasibility only, not production
T23	$D_{\text{KL}}(P_{\text{measured}} \Vert P_{\text{Lindblad}})$ and D_{TV}	100 identical control replays	Held-out likelihood	Pre-defined divergence bound

T23 replaces v1.0's Landauer-based reservoir-consistency claim (which does not follow from Landauer's principle [12] as stated) with a predictive-distribution comparison: the primary criterion is whether predictive error decreases when experimentally measured reservoir parameters are substituted for nominal parameters. Landauer's thermodynamic bound is retained in the philosophical discussion (S10) but is not the statistical criterion.

8. Illustrative Strategic Game

v1.0 S9 stated a 4-player subgame-perfect equilibrium [13] under a single 4 × 4 payoff matrix. A subgame-perfect equilibrium requires an extensive-form specification with players, information sets, actions, order of moves, strategies, terminal histories, and payoff vectors—not a single matrix. V2.0 replaces this with a properly specified extensive-form game. Both the payoff values and the equilibrium selection are labelled illustrative. The extensive-form tree is shown in Figure 5.

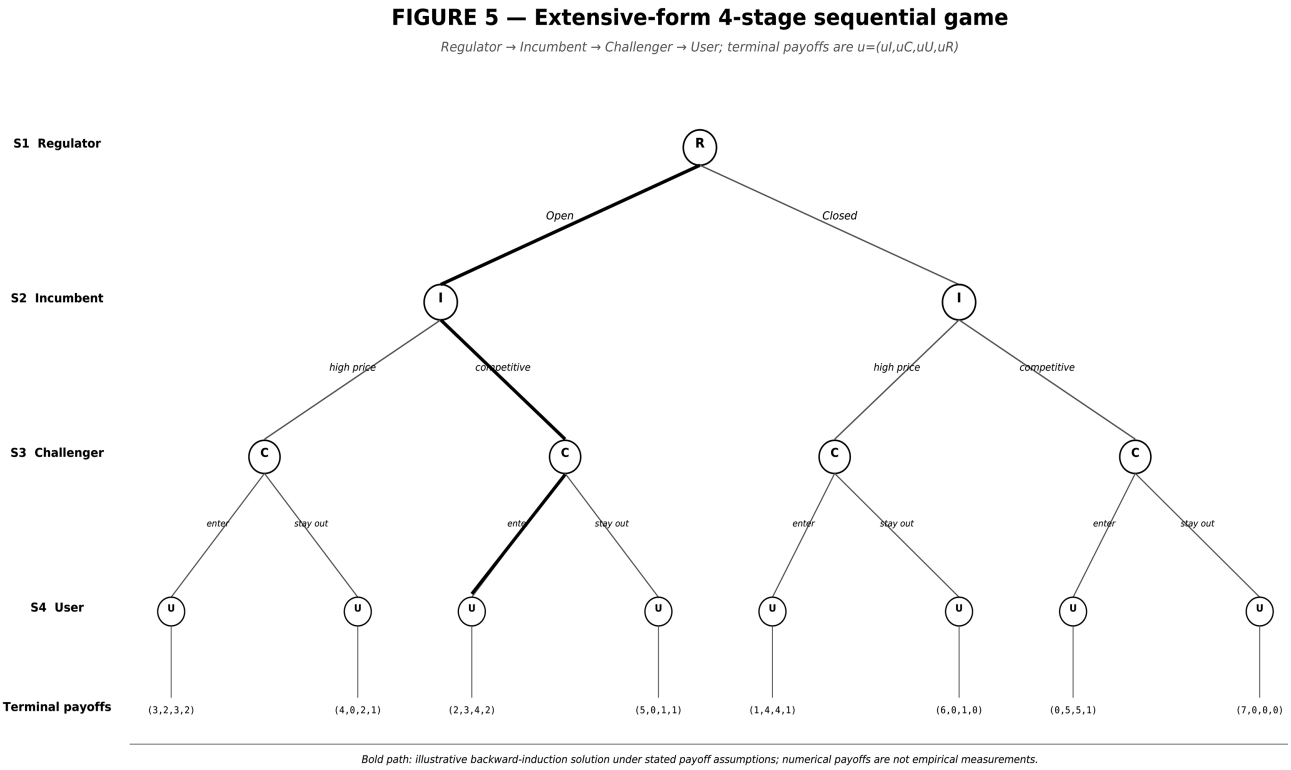


Figure 5. Extensive-form 4-stage sequential game. Stage 1: Regulator chooses open/closed. Stage 2: Incumbent chooses high/competitive price. Stage 3: Challenger chooses enter/stay out. Stage 4: User chooses incumbent/challenger. Terminal nodes carry payoff quadruples (u_I, u_C, u_U, u_R). The bold red path is the illustrative backward-induction solution under the stated payoff assumptions (open regulation → competitive incumbent → challenger enters → user picks challenger). The specific numerical payoffs are illustrative model assumptions, not empirical measurements. Replaces v1.0 **Figure 5** in its entirety.

9. Limitations

Every load-bearing limitation of this manuscript is enumerated in **Table 10** below.

Table 10. Explicit limitations enumeration. Every entry corresponds to a specific verification test or future experiment (S7 and S10).

#	Limitation
1	No STOC-QComp device has been fabricated; all results are engineering specifications.
2	The Casimir-mediated modification of qubit coupling and dissipation is a testable HYPOTHESIS (T20), not an experimentally demonstrated effect.
3	The 200 - 1000 qubit bounded-capacity regime is a design hypothesis motivated by Palmer's Rational Quantum Mechanics proposal [3], not an established physical law.
4	T1 is a demonstrator/control platform, not a full QPU. The 4 K thermal population catastrophe prevents conventional 5 GHz transmon operation at that temperature.
5	Manufacturing yields are engineering estimates (IE) or literature-derived (LE), not measured yields of STOC-QComp units.
6	Serial yield lies between 3.3% and 39.8%; the pipeline does not imply a finished-device yield > 40%.
7	The extensive-form game payoffs (**Figure 5**) are illustrative model assumptions, not empirical measurements.
8	T4 - T6 tier feasibility depends on successful validation of T2 - T3; T4 - T6 are ROADMAP.

Continued

9	The Casimir non-ideal correction values (Table 3) require a full Lifshitz calculation for the Au·SiN·Au multilayer, not performed in this manuscript.
10	Numerical values of $\partial\omega_c/\partial d$, $\partial g/\partial d$, $\partial\kappa/\partial d$, $\partial\Gamma_1/\partial d$ are design variables to be measured, not simulation outputs of this paper.

10. Conclusions

STOC-QComp v2.0 provides a device-level specification and experimental programme for investigating a mechanically tunable electromagnetic environment coupled to a bounded-capacity superconducting quantum processor. The present work does not establish the Casimir-mediated modification of qubit dissipation experimentally; rather, it identifies the physical parameters that must be measured to establish or falsify that hypothesis.

The most immediate validation target is a single reference transmon QPU coupled to a mechanically tunable microwave cavity via a Casimir-force-loaded actuator. The relevant observables are the cavity response, mechanical displacement spectrum, qubit relaxation and dephasing rates, and gate fidelity as functions of d. Manufacturing yield, environmental fidelity improvements, and deployment economics are likewise treated as engineering hypotheses whose uncertainty will be reduced through prototype fabrication.

If these experiments confirm the predicted coupling and environmental-control relationships, the STOC-QComp architecture provides a basis for subsequent scaling and deployment studies [14]. If they do not, the framework identifies specific falsification points at the Casimir, mechanical, microwave, or qubit-control layers.

On measurement and interpretation: v1.0 described collapse as an "engineerable event" resolving "interpretational debate". v2.0 uses the narrower, defensible language of an *operationally engineered measurement environment*: the device does not establish a new interpretation of quantum mechanics [15], but provides an experimentally controlled setting in which the relationship between deliberately engineered environmental coupling and measurement statistics can be investigated.

Research Status

STOC-QComp is a device-framework proposal (DESIGN). The qubit architecture, control interfaces, manufacturing flow, and deployment concepts are engineering specifications; they have not yet been demonstrated as an integrated production device. The Casimir mechanical structure is DESIGN. Casimir-mediated modification of qubit coupling and dissipation is an experimentally testable HYPOTHESIS. T1 and T2 are prototype targets (DESIGN); T3 is a DESIGN target; T4-T6 are ROADMAP configurations contingent upon successful validation of the preceding tiers. This paper does not claim that a STOC-QComp device has been fabricated, characterised, or shipped.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

[1] Pater, C.F. (2026) STOC-QSIM v2.0—Thermodynamic Control-Plane Substrate. Preprint, ResearchGate, 2026.

[2] Pater, C.F. (2026) STOC Quantum Computing—Bounded-Capacity Computational Framework. Preprint, ResearchGate, 2026.

[3] Palmer, T. (2026) Rational Quantum Mechanics: Testing Quantum Theory with Quantum Computers. *Proceedings of the National Academy of Sciences*, **123**, e2523350123. https://doi.org/10.1073/pnas.2523350123

[4] Feynman, R.P. (1982) Simulating Physics with Computers. *International Journal of Theoretical Physics*, **21**, 467-488. https://doi.org/10.1007/bf02650179

[5] Zurek, W.H. (2003) Decoherence, Einselection, and the Quantum Origins of the Classical. *Reviews of Modern Physics*, **75**, 715-775. https://doi.org/10.1103/revmodphys.75.715

[6] Bordag, M., Klimchitskaya, G.L., Mohideen, U. and Mostepanenko, V.M. (2009) Advances in the Casimir Effect. Oxford University Press/Springer.

[7] Aspelmeyer, M., Kippenberg, T.J. and Marquardt, F. (2014) Cavity Optomechanics. *Reviews of Modern Physics*, **86**, 1391-1452. https://doi.org/10.1103/revmodphys.86.1391

[8] Rasmussen, S.E., Christensen, K.S., Pedersen, S.P., Kristensen, L.B., Bækkegaard, T., Loft, N.J.S., *et al.* (2021) Superconducting Circuit Companion—An Introduction with Worked Examples. *PRX Quantum*, **2**, 040204. https://doi.org/10.1103/prxquantum.2.040204

[9] Krinner, S., Storz, S., Kurpiers, P., Magnard, P., Heinsoo, J., Keller, R., *et al.* (2019) Engineering Cryogenic Setups for 100-Qubit Scale Superconducting Circuit Systems. *EPJ Quantum Technology*, **6**, Article No. 2. https://doi.org/10.1140/epjqt/s40507-019-0072-0

[10] Paladino, E., Galperin, Y.M., Falci, G. and Altshuler, B.L. (2014) 1/f Noise: Implications for Solid-State Quantum Information. *Reviews of Modern Physics*, **86**, 361-418. https://doi.org/10.1103/revmodphys.86.361

[11] Cardani, L., Valenti, F., Casali, N. *et al.* (2021) Reducing the Impact of Radioactivity on Quantum Circuits in a Deep-Underground Facility. *Nature Communications*, **12**, Article No. 2733. https://doi.org/10.1038/s41467-021-23032-z

[12] Landauer, R. (1961) Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, **5**, 183-191. https://doi.org/10.1147/rd.53.0183

[13] Nash, J. (1951) Non-Cooperative Games. *The Annals of Mathematics*, **54**, Article 286. https://doi.org/10.2307/1969529

[14] CPater (2026) NWO Ecosystem Reference Implementations. https://huggingface.co/spaces/CPater/nwo.stoc

[15] Schlosshauer, M. (2007) Decoherence and the Quantum-to-Classical Transition. Springer. https://doi.org/10.1007/978-3-540-35775-9