Abstract

Quantum computing is often presented through a small set of recurring claims: continuous-variable systems are naturally advantaged, more modes means more power, quantum advantage has been demonstrated at device level, statistics could in principle be non-standard, complex amplitudes are essential, approximation-based gate synthesis is the only approach, a transmon is a qubit, and fault tolerance is a matter of years. This paper argues that each of these claims is true only within a restricted domain of applicability, and that the boundaries at which they fail are not defects of the claims but the places where the physics relevant to practitioners actually lives. Drawing on the locale framework introduced in a companion record [1], the paper states eight of the most widely repeated claims, makes the domain of each explicit, and pairs each boundary with work presented at the 23rd International Conference on Quantum Physics and Logic (QPL 2026) — together with one independent arXiv result — including new results on the energy cost of continuous-variable computation [2], a closely related independent result on trading modes against energy [3], classical simulation of Clifford+T circuits [4], invariance under quantum permutations [5], real-valued quantum theory [6,7], and the arithmetic structure of gate synthesis [8]. A review of the practitioner-oriented contributions at QPL 2026 follows, and five challenges raised at the conference — one sharpened by an independent arXiv result — are answered. The paper closes with the observation that the quantity which survives every boundary examined is energy per solution.

1. Introduction

The quantum computing field rests on a handful of sentences that are repeated so often they are no longer questioned. Continuous-variable systems, we are told, have a natural advantage because the physical world is continuous. More modes means more power. A fifty-qubit processor has demonstrated quantum advantage. Statistics could, in principle, be non-standard. Complex numbers are essential to quantum theory. Gate synthesis is an approximation problem. A transmon is a qubit. Fault tolerance is five years away.

Each of these sentences is true in some sense. Each is also false in another. This is not a contradiction; it is the normal condition of physical statements. Every physical claim holds only within a specified domain of applicability — a particular vacuum or symmetry sector, an energy scale, a noise regime, a measurement model, or a set of boundary conditions. The locale framework, introduced in a companion record [1], makes this structure explicit: a physical statement is a pair (locale, content), and the interesting physics lives at the seams where a rule fails to transfer from one locale to another.

This paper applies that framework to quantum computing as it is actually discussed by practitioners — engineers, investors, and decision-makers. Its purpose is not to debunk quantum computing, nor to defend it, but to give the reader a compact map: for each of eight widely repeated claims, what domain the claim requires, where the claim breaks, and which piece of the 2026 research literature documents the break. The map is built from the QPL 2026 program, which in one week contains the claims and nearly all of their boundaries.

2. The locale framework in brief

The companion record [1] develops four moves. First, *locale-conditionality*: every physical statement holds only within a locale, and the interesting physics lives at the seams where a rule fails to transfer. Second, *map and territory*: pedagogical maps (the spinning sphere for spin, the billiard balls for thermodynamics) fail at seams, while the territory — the transformation law, the algebra, the invariant — transfers. Third, the *rendering interface*: an observer inside a rendering experiences the metric of the image space, not the metric of the substrate, and no finite set of first-person observations decides between an Archimedean and an ultrametric substrate — yet a rendering with structure leaks structure. Fourth, *scale primitives*: stripped of anthropocentric units, the natural observables are counts and ratios.

This paper needs none of that machinery in full. It needs the first move (claims are conditional), the second move (spec sheets and roadmaps are maps), and one sharp consequence of the third move (what practitioners can and cannot infer about the machines they buy from the benchmarks they are shown). The fourth move supplies the conclusion: the quantity that transfers across every seam examined below is energy per solution.

3. Eight claims and their boundaries

Table 1 states eight claims that are heard constantly in quantum computing, the domain within which each holds, the boundary at which it fails, and the work that documents the boundary — presented at QPL 2026, with one arXiv exception. The table is the paper in miniature.

Table 1. Widely repeated claims in quantum computing, their domains of validity, and the boundaries documented at QPL 2026 (one by an independent arXiv result).

ClaimDomain in which it holdsBoundary at which it failsDocumentation Continuous-variable systems have a natural advantageUnbounded energy, ideal measurementsAt bounded energy, CV computation is efficiently simulable by qudits; the error obeys $\varepsilon \leq 1286 \cdot K \cdot n^2 \cdot E\backslash^{*2}/\sqrt{d}$Maltesson et al. [2] More modes means more powerMode count treated as freeModes trade against energy: constant modes require exponential energyBrenner, Dias & Koenig [3] — arXiv:2509.18854 Quantum advantage is demonstrated at device levelSpecific sampling tasks, permissive

baselinesExact classical sampling of Clifford+T circuits costs 2^t amplitude evaluations; realistic odd-dimension GKP states are classically simulableKoch [4]; Calcluth et al. (QPL 2026 talks) Statistics could be non-standardHigher-dimensional representations of the symmetric group (fundamental particles); emergent many-body systemsComplete and quantum permutation invariance force fundamental Bose or Fermi statistics, model-independently; parastatistics in emergent many-body or quasi-particle systems are not excludedMekonnen, Galley & Müller [5] Complex amplitudes are essentialNetwork experiments with source-state assumptionsReal-valued quantum theory reproduces all finite network correlations once independence is operational; it cannot be experimentally falsifiedHoffreumon & Woods [6,7] Solovay–Kitaev (Archimedean approximation) governs synthesisApproximation in the real metricExact Clifford+R synthesis carries explicit Bruhat–Tits building structure — arithmetic, not ArchimedeanDeaconu, Gargava, Kalra, Mosca & Yard [8] A transmon is a qubitThe two-level truncation of one deviceA transmon has $d \approx 12$ resolvable levels; the "qubit" is a 1.9% correction termQNFO, *The Two-Level Lie* [9] Fault tolerance by 2030Roadmap projectionsPhysical-to-logical overhead of 10^2–10^3, multiplied by cooling and decoding budgetsQNFO, *The Physics of Computation* [10]; JPCUB P0 [11]

Three remarks about the table.

First, none of the eight claims is simply wrong. Each is a conditional truth. The error, when there is one, is in the omission of the condition: the claim is stated without its locale, and the listener fills in a locale that the physics does not support.

Second, the boundaries are not speculative. Every boundary in the table is documented in work presented at a single conference week — with one exception, a closely related independent arXiv result. The QPL 2026 program is, in effect, a standing audit of the field's own claims.

Third, the table is a decision tool. A buyer, investor, or engineer who encounters any of the eight claims can ask: *which locale is being assumed, and does the current state of the art support it?* The table gives the answer in one line.

4. Why the boundaries matter: maps, territory, and the interface

The eight claims function as maps. A vendor's spec sheet is a map of a machine; a roadmap is a map of a field; a benchmark is a map of a capability. Maps are not wrong — they are locale-bound. The quantum volume of a processor, its reported fidelity, the number of qubits on a chip, the date on a roadmap: all of these are honest quantities that fail to transfer beyond the domain in which they were measured.

The territory — what actually transfers — is energy accounting. The first three rows of Table 1 are map failures of a single type: a claimed advantage that vanishes when the energy budget is made explicit. The territory that survives is the cost of producing a solution, measured end to end: computation, cooling, control, decoding, and the classical baseline it competes against [10,11].

The interface point is more subtle and more consequential. The practitioner's first-person evidence about a quantum computer is entirely image-space: benchmark dashboards, spec sheets, press releases, roadmap slides. The substrate — the actual physics of the device — is not accessible through that interface, and the framework's third move says something precise: no finite set of first-person observations of the image space decides the metric of the substrate. An Archimedean circuit model and an ultrametric substrate can render into the same image space [1].

What cannot be decided can still be probed, because a rendering with structure leaks structure. The clearest leak in the table is the sixth row: exact synthesis for Clifford+R turns out to carry an explicit Bruhat–Tits building structure [8]. A Bruhat–Tits building is an arithmetic object — the natural geometry of a p-adic group. It appears inside a problem that the field had treated as a routine Archimedean approximation task. The substrate of gate synthesis has, in the language of the framework, leaked through the rendering. This is an external result, produced by researchers working in the exact-synthesis tradition; its significance for the locale framework is that it is exactly the kind of structure leak the framework predicts [1,8].

5. What survives: counts, ratios, and energy

Strip the units off the discussion and three kinds of quantity remain.

Counts: qubits, gates, T-counts, cycles, physical qubits per logical qubit.

Ratios: error rates, anharmonicity, overhead factors, the ratio of classical to quantum resource use.

Energy per solution: a composite of counts and ratios that answers the only question a practitioner actually needs answered — what does a correct answer cost?

The first two rows of Table 1 make the energy statement quantitative. The CV↔DV equivalence of Maltesson et al. [2] says that a continuous-variable computation of n modes with energy budget E\ *is simulable by qudits with error at most* $1286 \cdot K \cdot n^2 \cdot E\backslash^2/\sqrt{d}$. Brenner, Dias and Koenig [3] give the converse direction: with a constant number of bosonic modes, the energy required is exponential. Taken together, the two results say that energy is the interconversion currency between computational paradigms — the exchange rate at which one formulation of a computation is converted into another. The joules-per-solution metric [11] is the natural unit of that exchange rate, and the two theorems are the first external, quantitative anchors for it.

6. Practitioner-oriented work at QPL 2026

The QPL 2026 program contains more than one hundred accepted contributions. The following are the ones most directly relevant to practitioners — people building, buying, or benchmarking quantum systems.

Synthesis and compilation. Buildings for Synthesis with Clifford+R [8] gives the exact synthesis structure for one of the most important non-Clifford gate sets, and thereby opens the arithmetic structure discussed above. Pauli gadget synthesis for gatesets with arbitrary even-arity Clifford gates (Meijer-van de Griend & Becker) and multi-qubit controlled gates with optimal T-count (Yamazaki & Akibue) address the resource counts that dominate circuit costs. A complete rule set for multi-qudit Clifford circuits in all odd prime dimensions (Bian, Li, Ross, van de Wetering & Zhao) extends the toolkit beyond qubits, and asymptotically optimal comparators (Vandaele) cover a standard subroutine. Clifford circuit synthesis for distributed architectures with arbitrary network topology (Laakkonen) addresses compilation across hardware.

Classical simulation and verification. Koch's classical Clifford+T sampler [4] evaluates 2^t amplitudes and handles distance-19 magic-state cultivation circuits in seconds — a direct constraint on what "quantum advantage" can mean. Calcluth et al. simulate realistic odd-dimension GKP states classically. Low-rank-width simulation in ZX-calculus (Kuyanov & Kissinger), hybrid treewidth/T-count simulation (Sutcliffe), automated circuit optimization with randomized replacements (Szyniszewski et al.), and tensor-and-gadget reinforcement learning for hardware-aware architecture search (Kundu) all extend the classical baseline against which quantum claims must be measured.

Error correction and fault tolerance. SpiderCat gives optimal fault-tolerant cat-state preparation (Khesin et al.). Fault tolerance by construction (Rodatz, Poor & Kissinger), flow-preserving rewrite rules (Backens & Perdrix), MWPM-decodability-preserving rewrites (Schweikart et al.), transversal AND in quantum codes (Li & Yeh), and phantom codes that entangle logical qubits without physical operations (Koh et al.) extend the QEC toolkit. Magic-state cultivation with few Clifford terms (Wan & Zhong) and basis-independent stabilizerness (Zurel & Davis) address the dominant resource bottleneck.

Practical protocols. Device-independent QKD with a single measurement per site (D'Avino et al.) and a resource-efficient quantum-walker quantum RAM (De Riso et al.) are directly deployable results.

The boundary papers — those that challenge the claims in Table 1 — are treated in the next section.

7. Responses to five challenges

Five challenges at QPL 2026 — one sharpened by an independent arXiv result — bear directly on the framework and on the claims in Table 1. They are addressed in turn.

Statistics. Mekonnen, Galley and Müller prove, by two model-independent arguments, that systems invariant under quantum permutations are either bosons or fermions [5]. This is consistent with the pre-registered negative result of the author's earlier work on the statistics distinction [12]: statistics is a locale effect of permutation invariance, not a theorem of any single formalism. The new result supplies an external upper bound on the routes by which non-standard statistics could have been obtained, including the channel-count route based on DHR locality. An open question remains: whether the invariance argument also constrains the two-dimensional braid-group locale — that is, whether it reaches anyons at all.

Real-valued quantum theory. Hoffreumon and Woods show that real-valued quantum theory reproduces all finite network correlations once independence is imposed operationally, and that it cannot be experimentally falsified [6,7]. The exchange scalar that governs boson/fermion statistics, $R = e^{(2\pi i s)} = (-1)^{(2s)}$, is real for integer and half-integer spin — the statistics seam never required complex amplitudes. Complex structure is load-bearing only at the anyon seam and in the algebra of interference. The claim "complex amplitudes are essential" is itself a conditional truth: true in the braiding locale, indistinguishable from a real rendering in every finite-network locale so far probed.

Energy as the exchange rate. The CV↔DV equivalence [2] and the mode-energy trade [3] — the latter an independent arXiv result — are the first externally proven seam crossings of the framework, and the first quantitative locale boundaries. They convert the joules-per-solution thesis [11] from a first-principles proposal into a theorem-anchored instrument. Two open questions remain: the tightness of the energy exponent E^{*2} versus the dimension term $1/\sqrt{d}$, and whether the binning in the construction can be softened to match the information used by real GKP decoders.

Classical simulation. The classical samplers and simulations of Koch [4] and Calcluth et al. re-index the locale of quantum advantage: each one shrinks the domain in which a device-level advantage claim holds. The invariant that survives is joules-per-solution — including the energy cost of the classical sampler itself, which the field does not yet account for systematically.

The arithmetic structure of synthesis. The Bruhat–Tits building structure of exact Clifford+R synthesis [8] is not a challenge but a confluence: an independent external confirmation that arithmetic, non-Archimedean structure is load-bearing at a practitioner seam. It is adopted as independent verification of the valuation-based program in ultrametric quantum computation, with the honest label that it is independent convergence, not derivation.

8. Conclusions

Eight claims, eight boundaries, one conference week — and one independent arXiv result. The practical content of this paper is Table 1: a locale-indexed map of the claims that quantum computing practitioners hear every day, each paired with the exact paper — presented at QPL 2026, with one arXiv exception — that documents where it breaks. The theoretical content is the observation that the same structure appears in every row — a claim true within a domain, failing at a boundary, and the physics of interest living at the boundary.

Three conclusions follow.

First, claims are conditional truths. The habit of stating them without their conditions is the single most expensive ambiguity in the field, because it is what allows a roadmap to be read as a promise and a benchmark as a product.

Second, the boundaries are public. The QPL 2026 program — one week, one venue — contains the documentation of every boundary in Table 1 but one, an independent arXiv result; the audit is public either way. The audit of quantum computing's claims is not a task for the future; it is the current research program of the field itself.

Third, energy per solution is what survives. Every boundary in Table 1 is, at bottom, an energy boundary. The quantities that transfer across all of them are counts, ratios, and the composite that answers the practitioner's question: what does a correct answer cost?

What this paper does not claim: no dynamical theory of boundaries is proposed; the table is a catalog, not a prediction; and nothing here decides whether the substrate of quantum computing is Archimedean or ultrametric — only that the question is real, and that the answer will leak through the rendering before it is decided from inside it.

References

[1] Quni-Gudzinas, R. B. *Conditional Truths and the Locale Framework: Map, Territory, and the Rendering Interface*. Zenodo, DOI 10.5281/zenodo.21983324 (concept). Version record: 10.5281/zenodo.21984929.

[2] Maltesson, A., Rodung, L., Budinger, N., Ferrini, G., Calcluth, C. *Equivalence of continuous- and discrete-variable gate-based quantum computers with finite energy*. arXiv:2510.08546.

[3] Brenner, L., Dias, B., Koenig, R. *Trading modes against energy*. arXiv:2509.18854.

[4] Koch, M. *Classical Clifford+T sampling without computing marginals*. QPL 2026, accepted talk.

[5] Mekonnen, M., Galley, T. D., Müller, M. P. *Invariance under quantum permutations rules out parastatistics*. arXiv:2502.17576.

[6] Hoffreumon, T., Woods, M. P. *Quantum theory based on real numbers cannot be experimentally falsified*. arXiv:2603.19208.

[7] Hoffreumon, T., Woods, M. P. *Quantum theory does not need complex numbers*. arXiv:2504.02808.

[8] Deaconu, M., Gargava, N., Kalra, A. R., Mosca, M., Yard, J. *Buildings for Synthesis with Clifford+R*. arXiv:2510.11526.

[9] Quni-Gudzinas, R. B. *The Two-Level Lie: The Transmon Is Not a Qubit — And the Entire Field Knows It*. Zenodo, DOI 10.5281/zenodo.21484345.

[10] Quni-Gudzinas, R. B. *The Physics of Computation: Fundamental Limits and the Honest Boundaries of Post-Classical Computing*. Zenodo, DOI 10.5281/zenodo.21255013.

[11] Quni-Gudzinas, R. B. *The Joules-per-Solution Metric: Definition, Measurement Protocol, and Anti-Gaming Provisions for Honest Computational Benchmarking*. Zenodo, DOI 10.5281/zenodo.21637028.

[12] Quni-Gudzinas, R. B. *The Boson/Fermion Distinction: Spin-Statistics as Structural Invariant*. Zenodo, concept DOI 10.5281/zenodo.21938970. Version record: 10.5281/zenodo.21964598.

Acknowledgements. We thank Manuel Mekonnen for correspondence that sharpened the statement of the permutation-invariance boundary in Table 1 (fundamental-particle scope; parastatistics in emergent many-body or quasi-particle systems are not excluded by the invariance arguments).