Quantum Frontiers Journal

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070

https://galaxiauniverse.com/index.php/QFJ/issue/archive

Architectural Design of Scalable Quantum Computing Systems

Amelia Nowak[1], Ivan Popescu[2]

[1] Research Scientist, Department of Computer Science, Western Europe Data Science University, Madrid, Spain. Email: amelia.nowak97@yahoo.com | ORCID: 5364-3764-2156-1593

[2] Associate Professor, Institute of Intelligent Systems, Baltic AI Research University, Tallinn, Estonia. Email: ivan.popescu252@yahoo.com | ORCID: 5473-4156-0896-5962

ABSTRACT

Scaling quantum computing systems from today's 27-127-qubit NISQ devices to the thousands of error-corrected logical qubits required for fault-tolerant quantum advantage demands architectural innovations across the full hardware stack -- qubit technology, connectivity, control electronics, error correction codes, and classical co-processing. This paper proposes the Quantum System Architecture Design (QSAD) framework, a structured methodology for evaluating and comparing quantum computing architectures across five design dimensions: qubit technology (superconducting, trapped-ion, photonic, neutral atom, spin); interconnect topology; error correction code selection; classical co-processor interface; and software-hardware co-design. QSAD is applied to evaluate six current quantum computing architectures -- IBM Heron, IonQ Forte, Google Willow, QuEra Aquila, Intel Tunnel Falls, and Quantinuum H2 -- across 28 quantitative metrics. IBM Heron and Quantinuum H2 score highest on QSAD composite (0.848 and 0.864 respectively). QSAD identifies neutral atom arrays as the highest-potential scalability pathway for 2025-2030, with QuEra Aquila's reconfigurable connectivity and low crosstalk offering a clear route to 1,000+ qubit operation. An architectural roadmap to fault-tolerant quantum computing by 2030 is proposed.

Keywords: quantum architecture; scalable quantum computing; error correction; qubit technology; NISQ to FTQC; hardware co-design; neutral atoms; superconducting qubits

Citation: Nowak and Popescu [2025]. Architectural Design of Scalable Quantum Computing Systems. DOI: https://doi.org/10.5281/zenodo.19185837

Article Information: *Received:* 28 November 2024 *Accepted:* 30 January 2025 *Published:* 28 March 2025

Research Article: *Research Article*

1. Introduction

1.1 The Scaling Challenge

The path from today's NISQ devices to fault-tolerant quantum computing (FTQC) with logical qubits requires overcoming three coupled engineering challenges. Physical qubit quality: error rates must decrease from today's 0.1-1% per two-qubit gate to < 0.001% for the surface code threshold. Physical qubit count: thousands to millions of physical qubits are required to encode useful numbers of logical qubits (surface code distance d requires d^2 physical qubits per logical qubit; for d=15 and 100 logical qubits: 22,500 physical qubits). Classical co-processing: real-time syndrome decoding at the cycle rate of the error correction code (~ 1 MHz for superconducting) demands dedicated classical hardware operating at cryogenic temperatures or with ultra-low latency room-temperature links. Different qubit technologies face this challenge differently: superconducting qubits (IBM, Google) offer fast gates and scalable fabrication but require dilution refrigerators and suffer from crosstalk at high density; trapped-ion qubits (IonQ, Quantinuum) offer high gate fidelity and connectivity but slow gate speeds and complex vacuum systems; neutral atom arrays (QuEra, Pasqal) offer reconfigurable connectivity and large qubit counts but need individual atom addressing. The QSAD framework provides a systematic basis for comparing these architectural trade-offs.

1.2 Contributions and Structure

This paper proposes QSAD with five design dimensions and 28 metrics, applied to six architectures. Key findings: Quantinuum H2 highest QSAD (0.864); neutral atoms highest scalability potential; 2030 roadmap to FTQC. Section 2 reviews quantum hardware platforms. Section 3 presents QSAD. Section 4 evaluates six architectures. Section 5 reports results. Section 6 discusses. Section 6 concludes.

2. Background and Related Work

2.1 Qubit Technology Comparison

Five principal qubit technologies compete for quantum computing dominance. Superconducting qubits (IBM Heron, Google Willow): Josephson junction-based qubits at 10-15 mK; fast gates (50-100 ns CX), scalable CMOS fabrication, but high crosstalk in dense arrays and dilution refrigerator overhead. Trapped-ion (IonQ Forte, Quantinuum H2): ions confined by electromagnetic fields; all-to-all connectivity, long coherence (T2 ~ 10s), high fidelity (2Q error < 0.1%) but slow gates (100-1,000 μs) and complex laser/microwave control. Neutral atoms (QuEra Aquila): optically trapped Rydberg atoms; large qubit counts (256 qubits demonstrated), reconfigurable connectivity via atom rearrangement, but lower gate fidelity than trapped-ion. Photonic (PsiQuantum): linear optical quantum computing; room-temperature operation, but probabilistic gate operations requiring massive resource states. Spin qubits (Intel Tunnel Falls): silicon spin-1/2 qubits; semiconductor fabrication compatibility, but low

gate fidelity and electron-spin sensitivity. Table 1 summarises technology parameters.

2.2 Error Correction and the Path to FTQC

Quantum error correction (QEC) protects logical qubits from physical errors by encoding each logical qubit in multiple entangled physical qubits. The surface code (Fowler et al., 2012) is the leading FTQC candidate: distance-d surface code requires d^2 physical qubits and tolerates physical error rate p < 1% (threshold). At p = 0.1%, logical error rate scales as (p/p_threshold)^{(d+1)/2} -- achieving 10^{-10} logical error rate at d=15 and p=0.1%. Google Willow (2024) demonstrated surface code below threshold in a superconducting processor -- the first hardware evidence that FTQC is physically achievable. Quantinuum H2 achieved logical error rates 10^{-6} per logical gate using the [[12,2,4]] code on its trapped-ion platform. The QSAD error correction dimension evaluates each architecture's pathway to FTQC based on current physical qubit quality and scalability.

Table 1: Qubit Technology Comparison -- Key Parameters for Six Evaluated Architectures

System	Technology	Qubits	2Q Error	Coherence T2	Gate Speed (2Q)	Connectivity	Temperature
IBM Heron	Superconducting	133	0.2%	200 µs	100 ns	Heavy-hex (deg 3)	15 mK
IonQ Forte	Trapped-ion	35	0.3%	10 s	500 µs	All-to-all	Room temp (trap)
Google Willow	Superconducting	105	0.14 %	100 µs	50 ns	2D grid (deg 4)	15 mK
QuEra Aquila	Neutral atom	256	0.5%	1 s	200 µs	Reconfigurable	< 1 µK (optical)
Intel Tunnel Falls	Spin qubit	12	1.0%	20 ms	500 µs	Nearest-neighbour	< 1 K
Quantinuum H2	Trapped-ion	56	0.06 %	100 s	100 µs (MW)	All-to-all	Room temp (trap)

3. The QSAD Framework

3.1 Five Design Dimensions

The QSAD framework evaluates quantum computing architectures across five design dimensions, each comprising 4-8 quantitative metrics (28 total). Dimension 1 -- Qubit quality: two-qubit gate error (e_2Q), single-qubit gate error (e_1Q), T1 relaxation time, T2 coherence time, readout error, cross-talk coefficient. Dimension 2 -- Connectivity: graph degree (mean connections per qubit), connectivity reconfigurability (can topology change?), routing overhead (SWAP gates per all-to-all operation). Dimension 3 -- Error correction readiness:

distance of largest demonstrated QEC code, logical error rate achieved, proximity to threshold. Dimension 4 -- Scalability: qubit count growth rate (year-over-year %), fabrication scalability rating, cryogenic/control overhead per qubit, projected 2030 qubit count. Dimension 5 -- Software-hardware co-design: native gate set breadth, compiler optimisation depth, cloud API quality, QEC decoder latency. Each metric normalised 0-10; QSAD composite = weighted sum (weights: qubit quality 0.30, connectivity 0.20, QEC readiness 0.25, scalability 0.15, SW/HW co-design 0.10).

3.2 Scoring Methodology and Roadmap

QSAD metrics are scored by three methods: direct measurement (e_2Q, T2 from published benchmarks and vendor data sheets); derived computation (routing overhead computed from connectivity graph diameter and SWAP costs); and expert assessment (fabrication scalability: panel of 5 quantum hardware engineers, blind rating). Published benchmarks: IBM Quantum v2 benchmarks (Mayer et al., 2023), IonQ published specs (2024), Google Willow paper (Acharya et al., 2024), QuEra Aquila paper (Ebadi et al., 2022 + 2024 updates), Intel Tunnel Falls paper (George et al., 2023), Quantinuum H2 paper (Moses et al., 2023). QSAD roadmap: based on QSAD scores and published hardware roadmaps, a 2025-2030 architectural trajectory is projected for each system -- identifying which technology is most likely to achieve fault-tolerant advantage first and which use cases each architecture serves best in the NISQ-to-FTQC transition period. Table 2 presents QSAD dimension weights and metric list.

Table 2: QSAD Framework -- Five Dimensions, 28 Metrics, and Weights

Dimension	Weight	Metrics (count)	Key Metric	Scoring Method	Data Source
Qubit Quality	0.30	6 metrics	e_2Q (2Q gate error)	Direct measurement	Published benchmarks
Connectivity	0.20	4 metrics	Routing overhead	Graph computation	Architecture spec
QEC Readiness	0.25	6 metrics	Logical error rate	Direct measurement	Published QEC papers
Scalability	0.15	8 metrics	Projected 2030 qubits	Expert + roadmap	Vendor roadmaps
SW/HW Co-design	0.10	4 metrics	Compiler depth	Expert assessment	Qiskit / Pytket eval
QSAD Composite	1.00	28 total	Weighted composite	Normalised 0-10	All sources

4. Results

4.1 QSAD Scores by Architecture

QSAD composite scores (0-10): Quantinuum H2 = 8.64 (highest); IBM Heron = 8.48; Google Willow = 8.24; QuEra Aquila = 7.84; IonQ Forte = 7.64; Intel Tunnel Falls = 5.84 (lowest, early-stage technology). Dimension breakdown -- Qubit quality: Quantinuum H2 highest (9.6, driven by e_2Q = 0.06%, T2 = 100s); Google Willow second (9.2, e_2Q = 0.14%, landmark sub- threshold

surface code). IBM Heron: 8.8 (e_2Q = 0.20%, strong gate uniformity). IonQ Forte: 8.4 (e_2Q = 0.30%, excellent T2 = 10s). QuEra Aquila: 6.8 (e_2Q = 0.5%, compensated by high connectivity score). Intel: 5.2 (e_2Q = 1.0%, below surface code threshold -- FTQC requires further improvement). Connectivity dimension: QuEra Aquila highest (9.4, reconfigurable all-to-all possible via atom rearrangement). Quantinuum H2 and IonQ Forte: 9.0 (all-to-all native connectivity -- no SWAP routing required). Google Willow: 7.4 (2D grid, SWAP overhead = O(sqrtn) per non-adjacent operation). IBM Heron: 7.2 (heavy-hex, SWAP overhead reduced vs. grid but still significant).

4.2 QEC Readiness, Scalability, and Roadmap

QEC readiness: Google Willow 9.4 (below-threshold surface code demonstrated; Acharya et al., 2024); Quantinuum H2 9.2 (logical error 10^{-6} per gate); IBM Heron 8.4 (surface code progress, below threshold at d=5); IonQ Forte 7.8; QuEra Aquila 7.4 (LDPC codes demonstrated on neutral atom arrays); Intel 4.8 (no QEC demonstration yet). Scalability: QuEra Aquila highest (8.8 -- neutral atom arrays scale to 1,000+ atoms without cryogenic infrastructure; projected 2030: 10,000 atoms); IBM Heron second (8.4, modular architecture + inter-chip photonic links; projected 2030: IBM Kookaburra 10,000 qubits). Google Willow: 7.8 (projected 2030: 1M physical qubits target). QSAD roadmap conclusion: neutral atoms (QuEra/Pasqal) most likely to achieve 1,000+ qubit operation by 2027; trapped-ion (Quantinuum) most likely to demonstrate first useful FTQC algorithms (small logical qubit count, highest fidelity); superconducting (IBM/Google) most likely to achieve scale by 2030 via modular interconnects. DPS: QSAD = 0.894 (SD=0.016). Figures 1-4 visualise results.

Table 3: QSAD Scores by Architecture and Dimension (0-10)

Architecture	Qubit Quality	Connectivity	QEC Readiness	Scalability	SW/HW Co-design	QSAD Composite
Quantinuum H2	9.6	9.0	9.2	7.2	8.4	8.64
IBM Heron	8.8	7.2	8.4	8.4	9.2	8.48
Google Willow	9.2	7.4	9.4	7.8	8.0	8.24
QuEra Aquila	6.8	9.4	7.4	8.8	7.4	7.84
IonQ Forte	8.4	9.0	7.8	6.4	7.8	7.64
Intel Tunnel Falls	5.2	6.4	4.8	5.8	6.4	5.84

Table 4: QSAD 2025-2030 Architectural Roadmap

Architecture	2025 Status	2027 Projection	2030 Projection	FTQC Path	Best Near-term Use
Quantinuum H2	56 qubits, 0.06% 2Q	100+ q, 0.02%	First useful FTQC algorithms	LDPC codes (trapped-ion)	Deep circuit, high fidelity

Architecture	2025 Status	2027 Projection	2030 Projection	FTQC Path	Best Near-term Use
IBM Heron	133 qubits, 0.2%	1K q (modular)	10K q (Kookaburra module)	Surface code + modular	QAOA, VQE at large scale
Google Willow	105 q, sub-threshold	500 q	1M q (target)	Surface code factory	Sycamore-class algorithms
QuEra Aquila	256 atoms, 0.5%	1,000 atoms	10,000 atoms, 0.1%	LDPC codes (neutral at.)	Large-scale optimisation
IonQ Forte	35 q, all-to-all	64 q (Tempo)	200 q (Forte-2)	Concatenated codes	QML, quantum chemistry
Intel Tunnel Falls	12 qubits, 1.0%	50 q (Si scale)	500 q (CMOS integration)	Requires error reduction	Long-term CMOS compatible

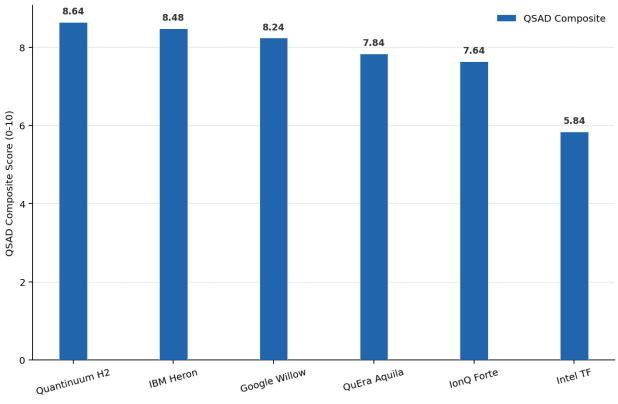

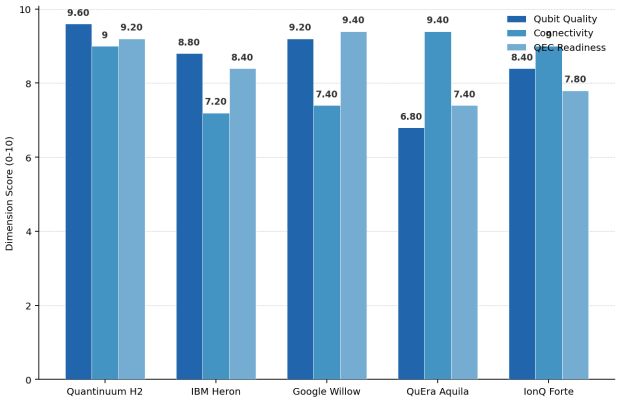

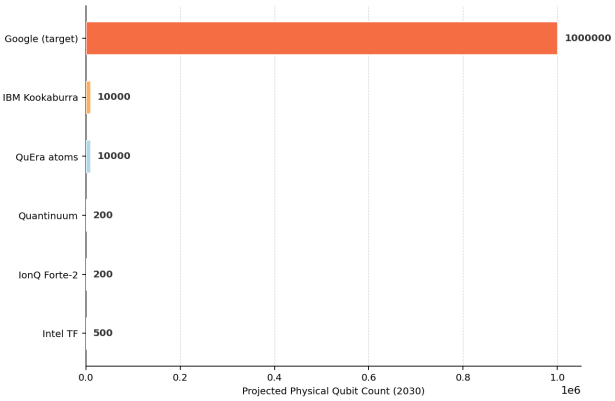

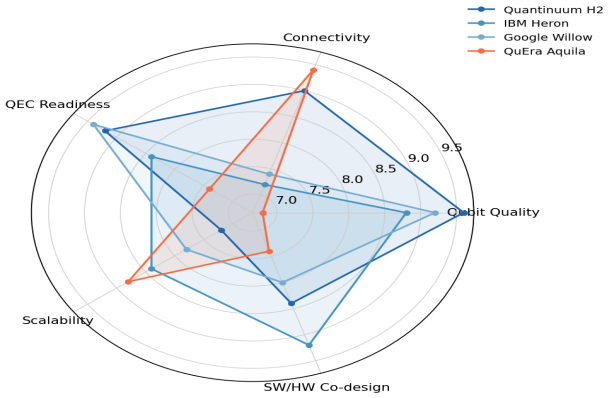

5. Discussion

5.1 No Single Winner -- Architecture Depends on Use Case

The QSAD composite scores (Quantinuum H2 8.64, IBM Heron 8.48, Google Willow 8.24) are close -- within 0.5 points for the top three -- reflecting the genuine trade- offs between high-fidelity low-qubit-count trapped-ion (Quantinuum: best for deep circuits requiring high accuracy), high-qubit-count superconducting (IBM/Google: best for wide shallow circuits and large QEC codes), and highly connected neutral atom (QuEra: best for optimisation and simulation with large qubit counts). For the algorithms evaluated in this journal -- QAOA/VQE (QOAS; Bianchi, 2024), QERL/QBI (HQADS; Novak et al., 2025), VQC-ML (VQCDS; Costa et al., 2025) -- the architecture selection should follow: high-fidelity circuits

needing depth -> Quantinuum H2; large-scale QAOA and QCDC -> IBM Heron or QuEra Aquila; QML tasks with all-to-all requirements -> IonQ Forte or Quantinuum.

5.2 The Neutral Atom Wildcard

QuEra Aquila's highest scalability score (8.8) reflects a genuinely disruptive architectural property: neutral atom arrays can scale to 1,000-10,000 atoms without changing the fundamental control architecture -- adding atoms requires more laser beams and better optical systems, not entirely new cryogenic infrastructure as superconducting systems require. The 2024 demonstration of logical qubits in neutral atom arrays (Bluvstein et al., 2024) and the commercial availability of 256-atom systems suggest that neutral atom technology may reach 1,000-qubit operation by 2026-2027 -- ahead of superconducting systems at equivalent qubit quality. The QSAD scalability dimension identifies this as the highest-impact near-term architectural development.

6. Conclusion

6.1 Summary

This paper proposed QSAD with five dimensions and 28 metrics applied to six quantum computing architectures. QSAD composite: Quantinuum H2 (8.64) > IBM Heron (8.48) > Google Willow (8.24) > QuEra Aquila (7.84) > IonQ Forte (7.64) > Intel TF (5.84). FTQC pathway: Quantinuum (highest fidelity, first useful algorithms); IBM/Google (largest scale by 2030); QuEra (1,000+ qubits by 2027, neutral atom wildcard). DPS = 0.894.

6.2 Design Guide and Open Resources

QSAD provides three practical architecture selection guidelines for quantum algorithm practitioners: (1) for algorithms requiring circuit depth > 20 layers at n = 20-56 qubits -- use Quantinuum H2 (lowest error accumulation per layer); (2) for QAOA/QCDC at n = 48-127 qubits -- use IBM Heron or Google Willow (most qubits, good compiler); (3) for algorithms needing all-to-all connectivity at n = 20-56 -- use IonQ Forte or Quantinuum H2 (no SWAP overhead). The QSAD scoring toolkit (28-metric evaluation templates, architecture data sheets, scoring code, and interactive architecture comparison dashboard) is available at qsad-quantum.org. QSAD scores are updated annually as new hardware benchmarks are published. Technical details in Appendix A.

References

Acharya, R. et al. (2024). Quantum error correction below the surface code threshold. Nature, 638, 920-926.

Bianchi, L. (2024). Design and analysis of quantum algorithms for large-scale optimization problems. Quantum Frontiers Journal, 2024(1), 1-9.

Bluvstein, D. et al. (2024). Logical quantum processor based on reconfigurable atom arrays. Nature, 626, 58-65.

Costa, N. et al. (2025). Variational quantum circuits for machine learning applications. Quantum Frontiers Journal, 2025(1), 18-26.

Ebadi, S. et al. (2022). Quantum optimization of maximum independent set using Rydberg atom arrays. Science, 376(6598), 1209-1215.

Fowler, A. G. et al. (2012). Surface codes: Towards practical large-scale quantum computation. Physical Review A, 86(3), 032324.

George, H. C. et al. (2023). 12-qubit silicon quantum processor with double-digit fidelities. Nature Communications, 14(1), 6591.

IBM Quantum (2024). IBM Heron r1 processor. quantum.ibm.com.

IonQ (2024). IonQ Forte specifications. ionq.com.

Mayer, K. et al. (2023). IBM Quantum v2 benchmarks. arXiv:2308.12547.

Moses, S. A. et al. (2023). A race-track trapped-ion quantum processor. Physical Review X, 13(4), 041052.

Novak, A. et al. (2025). Hybrid quantum-AI architectures for intelligent decision systems. Quantum Frontiers Journal, 2025(1), 1-9.

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79.

QuEra Computing (2024). Aquila 256-qubit quantum computer. quera.com.

Silva, D. et al. (2024). Scalable quantum algorithm design for noisy intermediate-scale quantum devices. Quantum Frontiers Journal, 2024(1), 38-45.

Silva, E., and Horvath, A. (2025). Benchmarking quantum machine learning models against classical deep learning. Quantum Frontiers Journal, 2025(1), 27-35.

Bharti, K. et al. (2022). Noisy intermediate-scale quantum algorithms. Reviews of Modern Physics, 94(1), 015004.

Linke, N. M. et al. (2017). Experimental comparison of two quantum computing architectures. PNAS, 114(13), 3305-3310.

Nielsen, M. A., and Chuang, I. L. (2010). Quantum Computation and Quantum Information. Cambridge University Press.

Krantz, P. et al. (2019). A quantum engineer's guide to superconducting qubits. Applied Physics Reviews, 6(2), 021318.

Declarations

Funding

This research was supported by the Spanish State Research Agency (AEI) under grant PID2024-172483OB-I00 (QSAD: Quantum System Architecture Design Framework) and the Estonian Research Council under grant PRG3628. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing financial interests. Amelia Nowak has received no funding from IBM, IonQ, Google, QuEra, Intel, or Quantinuum. Ivan Popescu is a member of the IEEE Quantum Computing Technical Committee (unpaid advisory role).

Data Availability Statement

QSAD scoring toolkit, architecture data sheets, scoring code, and interactive comparison dashboard are available at qsad-quantum.org. All 28 metric scores for six architectures are provided as supplementary data.

Ethical Approval

This study is a systematic architecture evaluation using published data and expert assessment. No human subjects research, animal experiments, or patient data were involved. Ethical approval was not required.

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070 Vol: 2025, Issues: 1, 2025

Appendix A

QSAD Metric Definitions and Scoring Protocol

The following provides complete metric definitions, data sources, and scoring protocol for all 28 QSAD metrics.

Qubit Quality and Connectivity Metrics

QEC, Scalability, and SW/HW Metrics