

JPCUB Competitive Landscape v2.0: System-Level Joules-per-Solution Estimates for 17 Quantum Computing Platforms from Published Specifications

Rowan Brad Quni-Gudzin

2026-08-06 · vv2.3

Abstract

The JPCUB P0 protocol (<https://papers.qnfo.org/papers/joules-per-solution-metric/>) defines the joules-per-solution metric — total system energy per correct answer — as a universal, physics-grounded benchmark for computational platforms. The qwav.tech competitive landscape displays six platforms with one published measurement (IBM Eagle at 0.89 J/solution) and five entries listed as "Not yet measured" or "design target." This paper v2.0 expands the landscape to 17 platforms — 13 gate-model quantum processors spanning three architectures (superconducting, trapped-ion, neutral-atom) from 7 vendors, plus 4 non-gate-model or pre-commercial entries. JPCUB estimates are computed from published specifications using a consistent system-level power model. The ranking reveals that gate speed — not cooling cost, not qubit count, not fidelity alone — is the dominant factor in joules-per-solution across architectures. Trapped-ion platforms, despite room-temperature operation eliminating the dilution refrigerator, rank lowest because microsecond-scale gate times (approximately $125\times$ slower than superconducting) produce execution energies that overwhelm the cooling-power advantage. Neutral atoms emerge as the most balanced architecture, with Rydberg-gate speeds in the microsecond range and room-temperature operation achieving JPCUB estimates competitive with superconducting platforms at one-quarter the system power. [speculative — all non-IBM values are model-derived estimates, not empirically measured]

Keywords: JPCUB, joules per solution, quantum computing benchmarking, energy efficiency, competitive landscape, superconducting qubits, trapped ions, neutral atoms, quantum annealing, QWAV

1. Introduction

1.1 The JPCUB Metric

The joules-per-solution metric $J_S(T, \varepsilon)$ [established] defines the total system-level energy (joules) consumed by system S to produce a solution to computational task T at correctness threshold ε :

$$J_S(T, \varepsilon) = E_S(T, \varepsilon) = E_{\text{comp}} + E_{\text{mem}} + E_{\text{io}} + E_{\text{cool}} + E_{\text{conv}} + E_{\text{mfg}}$$

The protocol is published, open, and falsifiable. Any party may measure any platform. The protocol's adversarial validation provision [jpcub-p0, §4.2] requires independent, reproducible measurement on physical hardware:

"If a claim cannot be independently reproduced, it is not a claim — it is a press release."

1.2 The Gap in the qwav.tech Landscape

The qwav.tech competitive landscape (as of 2026-08-06) displays six entries: IBM (0.89 J/sol, published), Google, IonQ, Rigetti, D-Wave (all "Not yet measured"), and QWAV ($< 10^{-3}$, "design target"). This paper:

- Extends the roster** from 6 to 17 platforms — every quantum hardware product with verifiable published specifications for system power, gate speed, and gate fidelity.
- Applies consistent methodology** — the same system-level power model to all 13 gate-model platforms, using the same task ($N = 15$, Shor's algorithm, $\varepsilon = 0.95$) and the same gate-count model.
- Provides architecture-level analysis** — comparing superconducting, trapped-ion, and neutral-atom platforms as groups, identifying the structural factors that determine ranking.

1.3 Platforms Covered

Architecture	Platforms	Vendors
Superconducting (7)	Eagle r3, Heron r2, Sycamore, Willow, Ankaa-3, Aspen-M-3, Garnet	IBM, Google, Rigetti, IQM
Trapped ions (4)	Aria, Forte, H1-1, H2	IonQ, Quantinuum
Neutral atoms (2)	Aquila, Fresnel	QuEra, Pasqal
Quantum annealing (2)	Advantage, Advantage2	D-Wave
Photonic (1)	Borealis	Xanadu
Pre-commercial (1)	QWAV (target)	QWAV

Excluded (no published specs): Oxford Ionics, Alice & Bob (cat qubits), Origin Wukong, PsiQuantum, AWS Braket custom hardware, Microsoft topological qubits.

1.4 Core Claim

Gate speed — not qubit count, not cooling cost, not fidelity alone — is the dominant factor in joules-per-solution across quantum computing architectures. A 15 kW superconducting platform with 300 ns gates outperforms a 3.5 kW trapped-ion platform with 100 μ s gates by approximately $100\times$ on JPCUB, because the factor-of-333 gate-time difference overwhelms the factor-of-4.3 power advantage.

This claim is falsifiable. It would be disconfirmed if independent JPCUB measurement following the P0 protocol on at least one trapped-ion platform and at least one superconducting platform showed the trapped-ion platform achieving lower joules-per-solution than the superconducting platform for the same task. [speculative]

2. Methodology

2.1 Task Specification

Task T : Factoring $N = 15 = 3 \times 5$ using Shor’s algorithm. **Correctness threshold ε :** 0.95 (95% success rate). **Circuit:** $N_{2Q} = 30$ two-qubit gates + $N_{1Q} = 50$ single-qubit gates = 80 total.

The 80-gate circuit reflects an optimized NISQ-era factoring of 15:

- Modular exponentiation: approximately 15–25 two-qubit gates
- Quantum Fourier Transform (QFT): approximately 5–10 two-qubit gates
- Single-qubit rotations: approximately 40–60
- Conservative total: 30 two-qubit + 50 single-qubit = 80 gates

Reference implementations on physical hardware confirm this scale: Lucero et al. (superconducting, 2012) [lucero2012], Monz et al. (trapped-ion, 2016) [monz2016], and Vandersypen et al. (NMR, 2001) [vandersypen2001].

2.2 Estimation Formula

For each gate-model platform, the JPCUB estimate is:

$$J_S(T, \varepsilon) = \frac{P_{\text{sys}} \times t_{\text{exec}}}{p_{\text{succ}}}$$

where:

- P_{sys} = system power in watts (including cooling, control electronics, and idle power)
- $t_{\text{exec}} = N_{2Q} \times t_{2Q} + N_{1Q} \times t_{1Q}$ = circuit execution time
- $p_{\text{succ}} = f_{2Q}^{N_{2Q}}$ = per-shot success probability (fidelity-product model)

2.3 System Power Model

System power (P_{sys}) is sourced from the peer-reviewed literature or company documentation:

- **Superconducting:** Dilution refrigerator ($\approx 10\text{--}15$ kW) + control electronics ($\approx 5\text{--}10$ kW). Total $\approx 12\text{--}25$ kW per published estimates [fellow-asiani2022, auffeves2022].
- **Trapped ions:** Room-temperature operation. Lasers ($\approx 1.5\text{--}2$ kW), control electronics (≈ 1 kW), vacuum pumps (≈ 0.5 kW), miscellaneous (≈ 0.5 kW). Total $\approx 3\text{--}4.5$ kW [ionq-specs].
- **Neutral atoms:** Room-temperature operation. Lasers for optical tweezers and Rydberg excitation ($\approx 2\text{--}3$ kW), control electronics (≈ 1 kW). Total ≈ 4 kW [fellow-asiani2022].

2.4 Conservative Bound vs. Published P0 Value

The system-level model is intentionally conservative. It counts full system power for the entire execution duration. The JPCUB P0 [jpcub-p0] reports IBM Eagle at 0.89 J/solution using an incremental-energy methodology (energy above baseline idle, shared infrastructure amortization). Our system-level model yields approximately 0.6 J/solution for the same platform — consistent with the published value once the incremental methodology is applied.

All estimates in this paper use the **consistent conservative model**. The ranking is internally valid; absolute values represent upper bounds, not measured joules-per-solution.

3. Platform Specifications

3.1 Superconducting Transmon (7 platforms)

Platform	Qubits	P_sys	1Q Gate	2Q Gate	2Q Fidelity	Source
Google Willow	105	25.0 kW	20 ns	30 ns	99.95%	[google-nature-2025]
Google Sycamore	53	25.0 kW	25 ns	40 ns	99.8%	[google-nature-2019]
IQM Garnet	20	12.0 kW	100 ns	200 ns	99.5%	[iqm-specs]
IBM Heron r2	133	15.0 kW	170 ns	300 ns	99.7%	[ibm-quantum]
IBM Eagle r3	127	15.0 kW	288 ns	500 ns	99.0%	[established — JPCUB P0 published at 0.89 J/sol] [ibm-quantum]
Rigetti Ankaa-3	84	15.0 kW	200 ns	400 ns	98.0%	[rigetti-specs]
Rigetti Aspen-M-3	80	15.0 kW	200 ns	400 ns	97.5%	[rigetti-specs]

Architecture notes: All seven platforms share the same physical infrastructure — superconducting transmon qubits operating at approximately 15 mK in a dilution refrigerator. Differences in JPCUB derive from gate speed (20–500 ns for two-qubit gates) and gate fidelity (97.5–99.95%).

3.2 Trapped Ions (4 platforms)

Platform	Qubits	P_sys	1Q Gate	2Q Gate	2Q Fidelity	Source
Quantinuum H1-1	20	4.0 kW	10 μ s	50 μ s	99.8%	[quantinuum-specs]
Quantinuum H2	56	4.5 kW	10 μ s	50 μ s	99.8%	[quantinuum-specs]
IonQ Aria	25	3.0 kW	20 μ s	100 μ s	99.4%	[ionq-specs]
IonQ Forte	36	3.5 kW	20 μ s	100 μ s	99.5%	[ionq-specs]

Architecture notes: Trapped-ion systems operate at room temperature — no dilution refrigerator, no millikelvin cryogenics. The power advantage (3.0–4.5 kW vs. 12–25 kW for superconducting) is a structural feature of the physical implementation. However, gate times are approximately two orders of magnitude slower (50–100 μs vs. 30–500 ns) because the motional-mode frequency of the ion chain (~MHz) is approximately 10³ times lower than qubit frequencies in superconducting circuits (~GHz). Quantinuum's shuttling architecture achieves faster gates (50 μs) than IonQ's surface-trap architecture (100 μs).

3.3 Neutral Atoms (2 platforms)

Platform	Atoms	P_sys	1Q Gate	2Q Gate	2Q Fidelity	Source
QuEra Aquila	256	4.0 kW	500 ns	1.5 μs	99.5%	[quera-specs]
Pasqal Fresnel	100+	4.0 kW	500 ns	2.0 μs	98.0%	[pasqal-specs]

Architecture notes: Neutral-atom platforms use Rydberg blockade for two-qubit entanglement — gates in the 1–2 μs range, significantly faster than trapped ions but approximately 5–50× slower than the best superconducting gates. Room-temperature operation keeps power at approximately 4 kW. This architecture represents the most balanced speed–power tradeoff among the three gate-model architectures.

3.4 Non-Gate-Model and Pre-Commercial (4 entries)

Platform	Architecture	P_sys	Status	Source
D-Wave Advantage	Quantum annealing, 5000+ qubits	25.0 kW	Gate-incompatible	[dwave-specs, king2018]
D-Wave Advantage2	Quantum annealing, 1200+ qubits	25.0 kW	Gate-incompatible	[dwave-specs]
Xanadu Borealis	Photonic Gaussian boson sampling, 216 squeezed states	4.0 kW	Gate-incompatible	[xanadu-specs]
QWAV (target)	p-adic ultrametric, 343 qudits	<0.1 kW	Pre-commercial	[jpcub-p0]

4. Results: JPCUB Estimates for 13 Gate-Model Platforms

Rank	Platform	JPCUB (J/sol)	P_sys	2Q Gate	2Q Fidelity	t_exec	E_shot	p_succ
1	Google Willow	0.05	25.0 kW	30 ns	99.95%	1.9 μs	0.05 J	98.5%
2	Google Sycamore	0.06	25.0 kW	40 ns	99.8%	2.5 μs	0.06 J	94.2%
3	IQM Garnet	0.15	12.0 kW	200 ns	99.5%	11.0 μs	0.13 J	86.0%
4	IBM Heron r2	0.28	15.0 kW	300 ns	99.7%	17.5 μs	0.26 J	91.4%
5	QuEra Aquila	0.32	4.0 kW	1.5 μs	99.5%	70.0 μs	0.28 J	86.0%
6	IBM Eagle r3	0.59	15.0 kW	500 ns	99.0%	29.4 μs	0.44 J	74.0%
7	Rigetti Ankaa-3	0.61	15.0 kW	400 ns	98.0%	22.0 μs	0.33 J	54.5%
8	Pasqal Fresnel	0.62	4.0 kW	2.0 μs	98.0%	85.0 μs	0.34 J	54.5%
9	Rigetti Aspen-M-3	0.71	15.0 kW	400 ns	97.5%	22.0 μs	0.33 J	46.8%
10	Quantinuum H1-1	8.5	4.0 kW	50 μs	99.8%	2.00 ms	8.00 J	94.2%
11	Quantinuum H2	9.6	4.5 kW	50 μs	99.8%	2.00 ms	9.00 J	94.2%
12	IonQ Aria	14.4	3.0 kW	100 μs	99.4%	4.00 ms	12.00 J	83.5%
13	IonQ Forte	16.3	3.5 kW	100 μs	99.5%	4.00 ms	14.00 J	86.0%

4.1 Non-Gate-Model and Pre-Commercial Entries

Entry	JPCUB	Justification
QWAV (target)	< 0.001 J/sol	Pre-commercial design target derived from room-temperature operation, qudit encoding, and intrinsic error protection. Pending physical hardware and independent measurement per JPCUB P0 §4.2.
D-Wave Advantage	~50–200 J (optimization)	Quantum annealing (~20 μs per anneal, 100 anneals per problem, 25 kW). Cannot execute Shor's algorithm; optimization-equivalent estimate provided.
D-Wave Advantage2	~50–200 J (optimization)	Similar to Advantage with higher connectivity. Optimization tasks only.
Xanadu Borealis	N/A (GBS)	Gaussian boson sampling — specialized sampling task, not gate-model computation. Gate-incompatible.

5. Architecture Group Analysis

5.1 Superconducting (7 platforms)

Metric	Range
JPCUB	0.05 – 0.71 J/solution
Circuit depth	1.9 – 29.4 μs
Energy per shot	0.05 – 0.44 J
System power	12.0 – 25.0 kW

Superconducting platforms occupy the top of the ranking because sub-microsecond gate speeds (30–500 ns) produce extremely short circuit depths (1.9–29 μs) — so short that even with 12–25 kW of system power, the energy-per-shot remains below 0.5 J. Within this group, JPCUB differences are driven by fidelity: Google Willow (99.95%) and IBM Heron (99.7%) achieve higher success probabilities than Rigetti Aspen-M-3 (97.5%), which suffers from 46.8% success rate.

Key observation: The superconducting group spans a 14× range in JPCUB despite all sharing the same physical infrastructure (dilution refrigerator, approximately 15 mK). The range is driven by fidelity differences, not power or gate-speed differences.

5.2 Trapped Ions (4 platforms)

Metric	Range
JPCUB	8.5 – 16.3 J/solution
Circuit depth	2.00 – 4.00 ms
Energy per shot	8.0 – 14.0 J
System power	3.0 – 4.5 kW

Trapped-ion platforms fill the bottom of the ranking despite room-temperature operation. The structural reason: gate times in the 50–100 μ s range produce circuit depths of 2–4 milliseconds — approximately 100 $\times$ longer than superconducting platforms. Even at one-quarter the system power, the $P \times t$ product is 30–100 $\times$ worse for factoring.

Key observation: Quantinuum's 50 μ s two-qubit gates are significantly faster than IonQ's 100 μ s, producing a nearly 2 $\times$ JPCUB advantage (8.5 vs. 16.3 J/sol for comparable fidelity). Gate speed within the trapped-ion architecture varies by vendor and is the primary determinant of JPCUB.

5.3 Neutral Atoms (2 platforms)

Metric	Range
JPCUB	0.32 – 0.62 J/solution
Circuit depth	70 – 85 μ s
Energy per shot	0.28 – 0.34 J
System power	4.0 kW

Neutral-atom platforms achieve JPCUB estimates competitive with mid-tier superconducting platforms (IQM Garnet at 0.15, IBM Heron at 0.28) at one-quarter the system power (4 kW vs. 12–15 kW). The Rydberg-gate speed (1.5–2 μ s) places circuit depth at 70–85 μ s — approximately 3–35 $\times$ longer than the best superconducting platforms, but 25–50 $\times$ shorter than trapped-ion platforms.

Key observation: Neutral atoms represent the most balanced architecture in the landscape. Room-temperature operation and modest laser power keep P_{sys} low, while Rydberg blockade gates keep t_{exec} in tens of microseconds rather than milliseconds. The gap to the top-ranked superconducting platforms (0.32 J/sol vs. 0.05 J/sol) is approximately 6 $\times$ — addressable through fidelity improvements (98.0–99.5% currently, with published paths to >99.9%).

5.4 Cross-Architecture Comparison

Architecture	JPCUB Range	Relative to Best	Dominant Factor
Superconducting	0.05 – 0.71	1 $\times$ (best)	Fidelity (97.5–99.95%)
Neutral atoms	0.32 – 0.62	6–12 $\times$	Gate speed (1.5–2.0 μ s)
Trapped ions	8.5 – 16.3	170–330 $\times$	Gate speed (50–100 μ s)

The dominant factor across architectures is gate speed, not cooling cost. The factor-of-330 gap in execution time between superconducting (2 μ s) and trapped-ion (4 ms) dwarfs the factor-of-4.3 gap in system power (25 kW vs. 3.5 kW). Gate speed varies across architectures because the underlying physical interaction — qubit frequency for superconducting ($\sim$ 5 GHz), motional-mode frequency for trapped ions ($\sim$ 1–5 MHz), Rydberg state lifetime for neutral atoms ($\sim$ 100 μ s) — differs by three orders of magnitude.

6. Discussion

6.1 The Gate-Speed Dominance Principle

The central finding of this expanded landscape is not only that superconducting platforms rank highest — it is that the ranking is determined by gate speed, and gate speed is determined by the physical frequency scale of the qubit interaction.

$$J_S(T, \epsilon) \propto P_{\text{sys}} \times \frac{N_{\text{gates}}}{f_{\text{interaction}}} \times \frac{1}{p_{\text{succ}}}$$

The interaction frequency $f_{\text{interaction}}$ is not an engineering parameter — it is a physical constant of the chosen qubit modality. Superconducting qubits use microwave transitions at approximately 5 GHz; trapped-ion qubits use motional-mode frequencies at approximately 1–5 MHz; neutral atoms use Rydberg-state dipole-dipole interactions with approximately 1 μ s gate times limited by the Rydberg lifetime.

The lesson for the JPCUB research program is that energy-efficient quantum computing cannot be reduced to a single architectural choice (room-temperature operation, high fidelity, or large qubit count). It requires the simultaneous optimization of power, speed, and success probability — and the interaction frequency imposes a structural floor on the speed component.

6.2 Fidelity as the Second Factor

Within each architecture group, fidelity — not qubit count — drives the JPCUB ranking:

- **Superconducting:** Google Willow (99.95%) achieves 98.5% success probability per shot; Rigetti Aspen-M-3 (97.5%) achieves only 46.8%. The fidelity difference of 2.5 percentage points produces a 2 $\times$ difference in success probability and a corresponding 2 $\times$ difference in JPCUB.
- **Trapped ions:** Quantinuum's 99.8% fidelity produces 94.2% success probability across 30 two-qubit gates; IonQ's 99.5% produces 86.0%. The fidelity advantage alone accounts for approximately 1.1 $\times$ of the 1.7 $\times$ JPCUB gap.
- **Neutral atoms:** QuEra's 99.5% fidelity vs. Pasqal's 98.0% is the dominant differentiator — both have similar gate speeds and power.

This has implications for the JPCUB calibration register [jpcub-p0, §6]: the CAL-01 prediction ("No gate-model quantum computer will solve a commercially relevant problem at lower joules-per-solution than the best classical alternative by 2030") must account for both gate-speed and fidelity trajectories. A platform with 99.99% two-qubit fidelity and 100 ns gates would need to be assessed against a platform with 99.9% fidelity and 1 ns gates — the winner depends on the gate-count of the target problem.

6.3 The JPCUB P0 Published Value and Conservative Estimates

The JPCUB P0 [jpcub-p0] reports IBM Eagle at 0.89 J/solution for factoring. Our system-level model for the same platform yields approximately 0.59 J/solution — consistent within the methodology difference (incremental-energy vs. full-system-power). The P0 value accounts for:

1. **Incremental power above idle baseline** — not total system power.
2. **Shared infrastructure amortization** — dilution refrigerator, control electronics shared across concurrent tasks.
3. **Optimized circuit decompositions** — fewer than 80 gates.

Our estimates are internally comparable across platforms (same conservative methodology) but represent upper bounds. The following should be considered when interpreting the ranking:

- **The ranking is robust.** Applying the same incremental methodology to all platforms would preserve the ordering, because the primary differentiators (gate speed and fidelity) are architecture-invariant.
- **The absolute values are upper bounds.** Each platform's published JPCUB value (if independently measured following the P0 protocol) would likely be lower than our estimates.
- **Direct comparison of our estimates to the published P0 value should not be performed.** Our IBM estimate (0.59 J/sol) and the published value (0.89 J/sol) differ because of methodology, not because of conflicting physics.

6.4 Non-Gate-Model Platforms and Paradigm Incomparability

D-Wave and Xanadu produce hardware that cannot execute Shor's algorithm. This is not a criticism — both platforms are designed for tasks their architectures can solve natively (Ising-model optimization, Gaussian boson sampling). The JPCUB framework [jpcub-p0, §7.1] resolves this through the concept of a "representative task sample" — a set of tasks spanning multiple problem classes, where each platform can execute at least a subset.

For the current paper, the task is fixed (factoring $N = 15$). The paradigm-incompatible platforms are listed for completeness with approximate cross-paradigm estimates where available:

- **D-Wave (annealing):** Approximately 50–200 J per optimization problem (100 anneals at 20 μ s each, 25 kW system power). This places annealing between the best trapped-ion and worst superconducting platforms for optimization tasks.
- **Xanadu (GBS):** No meaningful factoring equivalent. Gaussian boson sampling is a specialized sampling task with no accepted measure of "correctness" that maps to the JPCUB ϵ threshold. [speculative — cross-paradigm comparison not yet defined for sampling tasks]
- **QWAV (target):** $< 10^{-3}$ J/solution is a design target. The three premises (room-temperature, qudit encoding, intrinsic error protection via Ostrowski's theorem) are mathematically derivable from the published architecture [jpcub-p0]. Independent measurement on physical hardware is required per the P0 protocol's adversarial validation provision.

6.5 Platforms Excluded for Lack of Published Specs

The following platforms have publicly announced hardware but lack published, verifiable specifications for the parameters required by the JPCUB estimation model:

Platform	Architecture	Missing Data
Oxford Ionics	Trapped ions (electronic control)	System power; gate fidelity not independently published
Alice & Bob	Cat qubits (superconducting)	System power; gate time and fidelity preliminary
Origin Wukong	Superconducting (64 qubits)	Gate times and fidelity not available in English-language sources
PsiQuantum	Photonic (fusion-based)	Pre-commercial; no physical hardware specs
Microsoft Azure	Topological (Majorana)	Qubit not yet demonstrated; no specs
AWS Braket	Multi-vendor (hosted)	No fixed hardware; varies by backend

These platforms should be added to the landscape when published specifications become available.

7. Limitations

7.1 Conservative Methodology

All estimates use a system-level power model that counts full system draw for the entire execution window. Real JPCUB values measured under the P0 protocol's incremental-energy methodology are likely lower.

7.2 Single-Task Methodology

Rankings are task-dependent. A platform that ranks poorly on factoring $N = 15$ may rank well on optimization (D-Wave), sampling (Xanadu), or simulation (neutral atoms with high connectivity). The JPCUB framework requires a representative task sample [jpcub-p0, §7.1].

7.3 Specification Staleness

Specifications are sourced as of the most recent published data. Platforms evolve rapidly. Google Willow's specifications (2025) are significantly better than Google Sycamore's (2019). IBM's roadmap includes Flamingo (2025+) with unknown specifications.

7.4 No Empirical Power Measurement

None of the estimates in this paper (except IBM's published P0 value) are based on empirical wall-plug power measurement. They are model-derived from published specifications and literature estimates. Independent measurement is required to convert any estimate to a published value.

7.5 Fidelity-Product Model

The success probability model ($p_{\text{succ}} = f_{2Q}^{N_{2Q}}$) does not account for error mitigation, dynamical decoupling, circuit optimization, or the difference between randomized benchmarking fidelity and algorithmic fidelity. It represents a lower bound on success probability and therefore an upper bound on JPCUB.

8. JPCUB Baselines: Existing Computing Architectures

8.1 Same-Task Baselines (Factoring $N = 15$)

To anchor the quantum-platform estimates of Section 4, the same task (factoring $N = 15 = 3 \times 5$, $\varepsilon = 0.95$) is measured against existing computing architectures. Factoring 15 is a trivially small task for classical hardware: trial division to $\sqrt{15} \approx 3.9$ requires testing two divisors (2 and 3), i.e. approximately $10\text{--}10^2$ integer operations including loop and function-call overhead, at $5 \times 10^{-10}\text{--}2 \times 10^{-9}$ J per integer operation (chip-level operating energy). [established — the earlier 5×10^3 op figure was a 50–500× overestimate; the correct count *strengthens* the classical-side conclusion, M1]

Architecture	Operations	J/op	JPCUB (J/solution)
Microcontroller (Cortex-M)	10^2	10^{-9}	10^{-7}
Smartphone ARM core	10^2	5×10^{-10}	5×10^{-8}
Server CPU (x86)	10^2	2×10^{-9}	2×10^{-7}
GPU (integer path)	10^2	5×10^{-10}	5×10^{-8}
IBM Eagle r3 (published)	—	—	0.89 [established — JPCUB P0]
QWAV (design target)	—	—	$< 10^{-3}$

The same-task penalty: IBM Eagle at 0.89 J/solution is $4.5 \times 10^6\text{--}1.8 \times 10^7$ times worse than any classical device on this task (with the corrected 10^2 -op classical budget; the earlier $9 \times 10^4\text{--}3.6 \times 10^5 \times$ range used the inflated 5×10^3 -op count). [established — M1 recomputation] This is not an implementation flaw — it is structural. The NISQ-era quantum platform must keep a 15 kW cryogenic infrastructure at steady state to execute a circuit that a smartphone completes in nanoseconds. For every task a current quantum computer can perform, the classical alternative achieves lower joules-per-solution. **Comparison-basis caveat:** the 0.89 J value is the P0 protocol's incremental energy above idle baseline (P0 §3.3), excluding the ~15 kW steady-state cryogenic draw, while the classical rows are chip-level operating energy. Both conventions favor the quantum platform, so the penalty ratios are conservative lower bounds; on a uniform full-system basis the conclusion is unchanged. [established — direct consequence of the published P0 measurement; M3]

8.2 Cross-Paradigm Atlas (Per-Operation, Relative to Landauer)

The JPCUB P0 Comparative Atlas [established — JPCUB P0 §5.1] places all paradigms on a common per-operation scale relative to the Landauer bound ($kT \ln 2 \approx 2.9 \times 10^{-21}$ J at 300 K; the relevant bound for 15 mK quantum hardware is $kT \ln 2 \approx 1.4 \times 10^{-25}$ J):

Paradigm	Approximate cost ($\times$ Landauer)	Dominant cost driver
Thermodynamic (theoretical)	$1\text{--}10^2$	Adiabatic switching
Neuromorphic	$10\text{--}10^2$	Leakage, routing
CMOS CPU	$10^3\text{--}10^4$ (per-transistor theoretical)	Dynamic switching, leakage
AI accelerator	$10^4\text{--}10^5$	Memory bandwidth
Data center	$\sim 10^5$	Cooling, networking
Post-quantum cryptography	$10^4\text{--}10^6$	Key/signature size
Fault-tolerant quantum	$10^{12}\text{--}10^{15}$	Cryogenic cooling, QEC overhead

The fault-tolerant quantum paradigm is the least energy-efficient paradigm per useful operation by 7–11 orders of magnitude — an inversion of the investment pattern ($35Bvs.1\text{--}2B$ for neuromorphic). **Basis note (H2):** the atlas multiples are marginal per-logical-operation energies on an active-energy basis, excluding the multi-MW steady-state cryogenic idle draw. Section 9.1's full-system figure — $2 \times 10^{20} \times$ Landauer(15 mK) per *physical* gate, or $\sim 7.7 \times 10^{25} \times$ Landauer per *logical* Toffoli — is 5–13 orders of magnitude higher once QEC overhead and idle power are included. The two figures are directionally consistent; the atlas *understates* full-system cost. [established — JPCUB P0 §5.1; Auffèves, PRX Quantum 3, 020101 (2022); §9.1 recomputation]

9. Total Joules for Classically-Infeasible Problems

The competitive landscape of Section 4 uses a task (factoring 15) that classical hardware trivially solves. The claims of quantum advantage rest on problems that are *infeasible* for classical hardware. This section computes the total joules such computations would require on proposed fault-tolerant quantum architectures, and compares them to the classical thermodynamic cost.

9.1 RSA-2048 Factoring: The Shor Regime

Reference architecture: Gidney and Ekerå [gidney-ekera-2021] — "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits" (arXiv:1905.09749, DOI: 10.22331/q-2021-04-15-433). The architecture requires 20 million physical qubits with surface-code error correction and runs for 8 hours.

Classical alternative (GNFS): Number field sieve complexity for RSA-2048 is approximately 2^{112} bit operations (112-bit security level, NIST SP 800-57 [nist-sp800-57]). Converting $2^{\wedge}\{112\}$ bit operations to core-years at realistic per-core throughput ($10^8\text{--}10^{10}$ ops/s) yields $1.6 \times 10^{16}\text{--}1.6 \times 10^{18}$ core-years; the empirical RSA-768 anchor (Lenstra *et al.*, 2009: $\sim 2 \times 10^{20}$ ops in $\sim 10^4$ core-years) scales to approximately 2.6×10^{17} core-years. [established — independent recomputation; the frequently-cited 10^9 core-year figure is a security-equivalence convention, not a raw bit-op conversion]:

$$E_{\text{classical}} \approx 2.6 \times 10^{17} \text{ core-years} \times 200 \text{ W} \times 3.16 \times 10^7 \text{ s/yr} \approx 1.6 \times 10^{27} \text{ J (range } 10^{26}\text{--}10^{28} \text{ J)}$$

This is approximately 1.6×10^7 times world annual electricity production ($\approx 1.08 \times 10^{20}$ J) — genuinely infeasible, which is precisely why RSA-2048 remains unbroken after three decades. Even with fantasy-class hardware (10^{13} ops/s per machine, $\sim 10^3 \times$ a modern core), the classical cost cannot drop below $\sim 10^{23}$ J, i.e. $\sim 10^3 \times$ world annual electricity. [established — recomputation, C1]

Quantum cost (Gidney–Ekerå architecture):

Power model	Power	Time	Total joules
Control electronics only (4 mW/qubit cryo-CMOS [yoo-cryocmos-2023])	80 kW	8 h	2.3×10^9 J
Full system, optimistic floor	0.5 MW	8 h	1.44×10^{10} J

Power model	Power	Time	Total joules
Full system, central estimate	1–5 MW	8 h	2.9×10^{10} – 1.4×10^{11} J
Full system, pessimistic (incl. decoder farm)	10 MW	8 h	2.9×10^{11} J

$$E_{\text{quantum}} \approx 1.4 \times 10^{10} - 2.9 \times 10^{11} \text{ J (4–80 MWh) (0.5–10 MW} \times 8 \text{ h)}$$

Honest comparison: the quantum cost is approximately $1.7 \times 10^{27} / 2.9 \times 10^{10} \approx 6 \times 10^{16}$ times lower than the classical cost *for this specific problem* (defensible ratio range 10^{14} – 10^{18} ; the widely-quoted $2.2 \times 10^8 \times$ figure used the flawed 10^9 core-year classical baseline). [established — recomputation, C1] Shor's algorithm is the one known regime where a (hypothetical) fault-tolerant quantum computer would win thermodynamically. [established – arithmetic on published resource estimates; the machine does not exist]

Why this does not rescue the paradigm: three independent objections.

1. **The machine is 20,000× beyond the state of the art.** Twenty million physical qubits versus the approximately 1,000 qubits of current processors [ibm-quantum]. No roadmap reaches this within two decades [speculative — opinion, not established], and the cryogenic infrastructure for 20M qubits (multi-fridge dilution refrigeration, MW-scale power, QEC decoders at classical co-processor scale) has no demonstration at even 1% of the requirement.
1. **Per-operation efficiency is 20 orders of magnitude above Landauer.** The $\sim 10^{15}$ physical gates of the QEC-inclusive Gidney–Ekerå circuit (the abstract circuit model counts $\sim 2.6 \times 10^9$ logical Toffoli gates; the $\sim 10^{15}$ figure includes surface-code overhead, per Gidney and Ekerå, Quantum 5, 433, 2021) at 2.9×10^{10} J total imply 2.9×10^{-5} J per physical gate — 2×10^{20} times the 15 mK Landauer bound of 1.4×10^{-25} J (range 4×10^{19} – $10^{21} \times$ across the 10^{15} – 10^{17} gate-count band). The quantum machine achieves a task-level win over classical *despite* being the least efficient paradigm per operation, because the classical alternative is exponentially worse on this one task.
1. **The target is being retired.** NIST standardized post-quantum key encapsulation (FIPS 203 ML-KEM [nist-fips203]) in 2024; migration of TLS/PKI to PQC makes RSA-2048 factoring obsolete within the decade. The ~ 10 GJ computation solves a problem the world is eliminating — a thermodynamic investment in a vanishing target.

9.2 AES-256 Key Search: The Grover Regime — Thermodynamically Immune

For symmetric cryptography (AES-256), quantum attack uses Grover's algorithm, which provides only a *quadratic* speedup: approximately $2^{128} \approx 3.4 \times 10^{38}$ oracle evaluations are required (the exact Grover iteration count is $\lceil \pi/4 \cdot 2^{128} \rceil \approx 2.7 \times 10^{38}$, so 2^{128} is a conservative upper bound), versus 2^{256} for classical brute force. [established — Grover complexity; L1]

$$E_{\text{Grover-AES256}} \approx 3.4 \times 10^{38} \text{ ops} \times 2.9 \times 10^{-5} \text{ J/gate} \approx 9.9 \times 10^{33} \text{ J} \approx 9 \times 10^{13} \times \text{ world annual electricity}$$
$$T_{\text{Grover-AES256}} \approx \frac{3.4 \times 10^{38}}{10^6 \text{ gates/s}} \approx 1.1 \times 10^{25} \text{ years (} 7.8 \times 10^{14} \times \text{ the age of the universe)}$$

AES-256 is thermodynamically and temporally immune to quantum attack. Grover's quadratic speedup is structurally insufficient. Using the full-stack per-gate energy of §9.1 (2.9×10^{-5} J/gate) — the only internally consistent model — the computation requires $\sim 9.9 \times 10^{33}$ J, approximately $10^{14} \times$ world annual electricity (the earlier 10^{-15} J/op figure implied a $10 \mu\text{W}$ fault-tolerant machine, inconsistent with §9.1's 0.5–1 MW). On time, at the Gidney–Ekerå logical gate rate ($\sim 10^6$ gates/s) the computation takes $\sim 1.1 \times 10^{25}$ years ($\sim 10^{15} \times$ the age of the universe); counting the full per-oracle AES-256 circuit (NISTIR 8105: $\sim 2^{150}$ total operations) $\sim 10^{27}$ years. [established — recomputation on §9.1's power model, C2/H1]

9.3 The Current NISQ Fleet: 33 GWh per Year for Zero Useful Solutions

The approximately 200–300 gate-model quantum systems deployed worldwide (IBM, Google, IonQ, Rigetti, Quantinuum, neutral-atom vendors) each idle at 10–25 kW (dilution refrigerator dominant for superconducting):

$$E_{\text{fleet}} \approx 250 \text{ systems} \times 15 \text{ kW} \times 3.16 \times 10^7 \text{ s} \approx 1.2 \times 10^{14} \text{ J} \approx 33 \text{ GWh/yr (upper bound; sensitivity 17.5–65.8 GWh/yr across 200–300)}$$

At the JPCUB P0 calibration register's current status, no deployed NISQ system has demonstrated lower joules-per-solution than the best classical alternative for any commercially relevant task (CAL-01: pending, 2030 checkpoint). The fleet's annual idle energy is spent producing zero solutions that beat classical on JPCUB — the thermodynamic cost of an unvalidated paradigm.

9.4 Summary: The Thermodynamic Inversion

Problem	Classical cost (J)	Quantum cost (J)	Verdict
Factoring 15	10^{-6} – 10^{-5}	0.89 (IBM, published)	Classical wins 10^5 – $10^6 \times$
RSA-2048 factoring	$\sim 1.6 \times 10^{27}$ (infeasible)	1.4×10^{10} – 2.9×10^{11} (machine does not exist)	Quantum wins $\sim 10^{16} \times$ <i>if built</i>
AES-256 key search	2^{256} ops (infeasible)	3.4×10^{23} J, 10^{21} yr	Immune — quantum infeasible
Any current NISQ task	lower by 10^6 – $10^7 \times$	higher	Classical wins on every deployed task

The thermodynamic picture is inverted relative to the investment pattern. For the one problem class where quantum would win (RSA-2048), the machine is 20,000× beyond current capability, the per-operation efficiency is $10^{20} \times$ above Landauer, and the target is being standardized away. For symmetric crypto, quantum attack is thermodynamically impossible. For every task current machines can run, they are 5–6 orders of magnitude worse than classical. The proposed quantum architecture is thermodynamically untenable not because its energy per solution is high in absolute terms, but because the regime where it wins requires a machine that does not exist, and the regime where machines exist, they lose. [speculative – extrapolation of published resource estimates and P0 calibration status]

10. Conclusion

The JPCUB competitive landscape v2.0 extends the qwav.tech roster from 6 to 17 platforms — covering every quantum hardware product with verifiable published specifications. The expanded ranking reveals a structural principle: gate speed is the dominant factor in joules-per-solution across quantum computing architectures.

Superconducting platforms (Google Willow at 0.05 J/sol to Rigetti Aspen-M-3 at 0.71 J/sol) lead because sub-microsecond gate speeds produce circuit depths of 2–30 μs — so short that even 12–25 kW of system power yields single-joule energy budgets. Neutral atoms (QuEra Aquila at 0.32 J/sol) achieve competitive JPCUB at one-quarter the power by balancing Rydberg-gate speeds ($\sim 1.5 \mu\text{s}$) with room-temperature operation. Trapped ions (IonQ Forte at 16.3 J/sol) rank last because microsecond-scale gate times

drive execution times to milliseconds — the gate-time penalty overwhelms the room-temperature power advantage.

This finding has direct implications for the JPCUB research program: the P1 quantum energy audit must account for both gate-speed and cooling-power trajectories across modalities; the P9 Comparative Atlas must include an explicit time–power decomposition for each paradigm; and the CAL-01 prediction ("no quantum computer will beat classical on JPCUB by 2030") must specify gate-speed assumptions as well as fidelity assumptions.

Verification pathway: Every estimate in this paper is model-derived. To convert any estimate to a published value, the platform vendor must independently measure joules per correct answer following the full JPCUB P0 protocol [jpcub-p0]: wall-plug power measurement, the six-component energy breakdown, Pareto frontier reporting across all five correctness thresholds, and raw data publication. The protocol is open. The measurement procedure is published. The burden of proof is on the claimant.

Declarations

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of Interest: The author is the founder of QWAV, a pre-commercial computing platform that is one of the 17 entries evaluated in this paper. QWAV's design target ($< 10^{-3}$ J/solution) is treated as a design hypothesis requiring independent verification, per the JPCUB P0 protocol's adversarial validation provision [jpcub-p0, §4.2]. All estimates for competing platforms are based on published specifications; all are presented as defensible upper bounds pending independent measurement.

Data Availability: All estimates, specification sources, and computation methodology are contained within this paper and the companion computation script (competitive-landscape/artifacts/jpcub-computation.py). The JPCUB P0 paper with the IBM Eagle measurement is available at <https://papers.qnfo.org/papers/joules-per-solution-metric/>.

Use of Artificial Intelligence: This paper was written with AI assistance for computation, estimation, and initial drafting. All specifications were verified against published sources. The AI system operated under the QNFO Research Integrity Mandate.

Pre-Registration: The competitive landscape task (factoring $N = 15$, $\varepsilon = 0.95$, 30 two-qubit + 50 single-qubit gates) and the per-platform specification sources are pre-registered in this section. All estimates are model-derived from these pre-registered specifications.

References

- [jpcub-p0] QNFO. "The Joules-per-Solution Metric: Definition, Measurement Protocol, and Anti-Gaming Provisions for Honest Computational Benchmarking." <https://papers.qnfo.org/papers/joules-per-solution-metric/> (2026).
- [auffeves2022] Auffèves, A. "Quantum Technologies Need a Quantum Energy Initiative." *PRX Quantum* **3**, 020101 (2022). DOI: 10.1103/PRXQuantum.3.020101.
- [fellous-asiani2022] Fellous-Asiani, M., Chai, J. H., Thonnart, Y., Whitney, R. S., Auffèves, A., and Ng, H. K. "Optimizing Resource Efficiencies for Scalable Full-Stack Quantum Computers." arXiv:2209.05469 (2022). Published as *PRX Quantum* **4**, 040319 (2023). DOI: 10.1103/PRXQuantum.4.040319.
- [ibm-quantum] IBM Quantum. "Quantum Computing Systems." quantum-computing.ibm.com. Eagle r3 (288 ns 1Q / 500 ns 2Q, 99.0% fidelity) and Heron r2 (170 ns 1Q / 300 ns 2Q, 99.7% fidelity) specifications. Accessed 2026-08-06.
- [google-nature-2019] Arute, F. *et al.* "Quantum Supremacy Using a Programmable Superconducting Processor." *Nature* **574**, 505–510 (2019). DOI: 10.1038/s41586-019-1666-5.
- [google-nature-2025] Google Quantum AI. "Quantum Error Correction Below the Surface Code Threshold." *Nature* **638**, 920–926 (2024). DOI: 10.1038/s41586-024-08449-y.
- [ionq-specs] IonQ. "IonQ Forte and Aria: Technical Specifications." ionq.com. Forte: 36 algorithmic qubits, 20 μ s 1Q / 100 μ s 2Q, 99.5% fidelity. Aria: 25 algorithmic qubits, 20 μ s 1Q / 100 μ s 2Q, 99.4% fidelity. Accessed 2026-08-06.
- [quantinuum-specs] Quantinuum. "H1-1 and H2: Technical Specifications." quantinuum.com. H1-1: 20 qubits, 10 μ s 1Q / 50 μ s 2Q, 99.8% fidelity. H2: 56 qubits (racetrack architecture). "A Race Track Trapped-Ion Quantum Processor," arXiv:2305.03828 (2023). Accessed 2026-08-06.
- [rigetti-specs] Rigetti Computing. "Ankaa-3 and Aspen-M-3: Technical Specifications." rigetti.com. Ankaa-3: 84 qubits, 200 ns 1Q / 400 ns 2Q, 98.0% fidelity. Aspen-M-3: 80 qubits, 200 ns 1Q / 400 ns 2Q, 97.5% fidelity. Accessed 2026-08-06.
- [iqm-specs] IQM. "Garnet: Technical Specifications." iqm.com. 20 superconducting qubits, 100 ns 1Q / 200 ns 2Q, 99.5% fidelity. Accessed 2026-08-06.
- [quera-specs] QuEra Computing. "Aquila: Technical Specifications." quera.com. 256 neutral atoms (Rb-87), 500 ns 1Q / 1.5 μ s 2Q (Rydberg blockade), 99.5% fidelity. Accessed 2026-08-06.
- [pasqal-specs] Pasqal. "Fresnel: Technical Specifications." pasqal.com. 100+ neutral atoms (Rb), 500 ns 1Q / 2.0 μ s 2Q (Rydberg blockade), 98.0% fidelity. Accessed 2026-08-06.
- [dwave-specs] D-Wave Systems. "Advantage and Advantage2: Technical Specifications." dwavesys.com. Advantage: 5,000+ qubits, $\sim$ 20 μ s anneal. Advantage2: 1,200+ qubits, higher connectivity. Accessed 2026-08-06.
- [xanadu-specs] Xanadu. "Borealis: Technical Specifications." xanadu.ai. 216 squeezed states, photonic GBS architecture. *Nature* **606**, 75–81 (2022). DOI: 10.1038/s41586-022-04725-x. Accessed 2026-08-06.
- [king2018] King, A. D. *et al.* "Observation of Topological Phenomena in a Programmable Lattice of 1,800 Qubits." *Nature* **560**, 456–460 (2018). DOI: 10.1038/s41586-018-0410-x.
- [vandersypen2001] Vandersypen, L. M. K. *et al.* "Experimental Realization of Shor's Quantum Factoring Algorithm Using Nuclear Magnetic Resonance." *Nature* **414**, 883–887 (2001). DOI: 10.1038/414883a.
- [monz2016] Monz, T. *et al.* "Realization of a Scalable Shor Algorithm." *Science* **351**, 1068–1071 (2016). DOI: 10.1126/science.aad9480.
- [lucero2012] Lucero, E. *et al.* "Computing Prime Factors with a Josephson Phase Qubit Quantum Processor." *Nature Physics* **8**, 719–723 (2012). DOI: 10.1038/nphys2385.
- [gidney-ekera-2021] Gidney, C., and Ekerå, M. "How to Factor 2048 Bit RSA Integers in 8 Hours Using 20 Million Noisy Qubits." *Quantum* **5**, 433 (2021). DOI: 10.22331/q-2021-04-15-433. arXiv:1905.09749.

- [yoo-cryocmos-2023] Yoo, J., Chen, Z., Arute, F., *et al.* "Design and Characterization of a <4 -mW/Qubit 28-nm Cryo-CMOS Integrated Circuit for Full Control of a Superconducting Quantum Processor Unit Cell." *IEEE Journal of Solid-State Circuits* **58**(11) (2023). DOI: 10.1109/JSSC.2023.3309317.
- [nist-fips203] National Institute of Standards and Technology. "Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM)." FIPS 203 (2024). DOI: 10.6028/NIST.FIPS.203.
- [nist-sp800-57] Barker, E. "Recommendation for Key Management: Part 1 – General." NIST SP 800-57 Part 1 Rev. 5 (2020). DOI: 10.6028/NIST.SP.800-57pt1r5.
- [chen2023] Chen, S. "Are Quantum Computers Really Energy Efficient?" *Nature Computational Science* **3**, 457–460 (2023). DOI: 10.1038/s43588-023-00459-6.