Factoring RSA-2048 in 20 Minutes with 42 Qubits Under Mildly Generous Assumptions

Marin Ivezic
Applied Quantum, Singapore
marin@appliedquantum.com
(Dated: September 15, 2025)

Abstract

We present a resource estimate for factoring 2048-bit RSA integers using a quantum computer with just 42 physical qubits and a runtime of approximately 20 minutes. Our approach combines well-established results from the recent literature: Regev's $O(n^{3/2})$ quantum factoring algorithm [1], asymptotically optimal qLDPC codes at their proven constant encoding rate [2], magic state cultivation at projected efficiency [3], demonstrated multi-hour coherence times in rare-earth ion memories [4], and all-to-all qubit connectivity as available on reconfigurable atomic platforms [5]. Each assumption is individually supported by peer-reviewed publications. We make no assumptions beyond those already present in the literature, though we acknowledge that we select the most favorable result from each of several mutually incompatible architectural paradigms and combine them into a single estimate. Our methodology extends the pioneering work of Yan et al. [6], who demonstrated in 2023 that sufficiently creative assumptions can yield arbitrarily low resource estimates. To our knowledge, 42 physical qubits represents the lowest resource estimate published for this problem. We discuss implications for the PQC migration timeline and recommend that all vulnerable systems be upgraded by next Tuesday.

I. INTRODUCTION

The cost of quantum factoring has decreased at a remarkable pace. In 2012, Fowler et al. estimated that factoring a 2048-bit RSA integer would require approximately one billion noisy physical qubits [7]. By 2021, Gidney and Ekerå reduced this to 20 million [8]. In May of this year, Gidney further reduced the estimate to under one million [9], a 20-fold improvement achieved primarily through approximate residue arithmetic [10], yoked surface codes [11], and magic state cultivation [3]. Extrapolating the observed trend on a logarithmic scale, one arrives at a single qubit by approximately Q3 2027.

We are not the first to pursue aggressively low estimates. In January 2023, Yan et al. [6] reported that RSA-2048 could be factored using only 372 qubits via a hybrid approach combining Schnorr's lattice-based factoring with the Quantum Approximate Optimization Algorithm (QAOA). While this claim was subsequently challenged on the grounds that the classical lattice reduction component does not scale [12, 13], and the quantum component provides no proven speedup over classical methods, we believe the underlying principle — that optimism is a valid algorithmic resource — remains sound. Our work builds on this principle, but with the important distinction that each of our individual assumptions is correct.

Our sole contribution is the observation that by selecting the most optimistic assumption from each paper in a large corpus and combining them — an approach we term *Optimistic Assumption Stacking* (OAS) — one obtains estimates dramatically lower than any individual paper achieves. Unlike Yan et al., who combined assumptions from a single flawed framework, we combine assumptions from many individually sound frameworks that happen to be mutually incompatible. We believe this represents a methodological advance.

We stress that each assumption we make is individually supported by a peer-reviewed or credible preprint publication. We make no speculative claims. That said, we acknowledge in footnote[1] that combining assumptions from papers that use fundamentally incompatible hardware platforms, mutually exclusive error correction codes, and

contradictory physical connectivity models may raise methodological questions that we leave for future work.

[1] We considered placing this acknowledgment in the main text but found it disrupted the narrative flow.

II. ALGORITHMIC FOUNDATION

We base our construction on Regev's quantum factoring algorithm [1], which achieves gate complexity $\tilde{O}(n^{3/2})$ using $O(n^{1/2})$ qubits in the asymptotic limit. Regev's algorithm was published in the Journal of the ACM in 2025, representing the highest standard of peer review in theoretical computer science. While Ragavan and Vaikuntanathan [14] noted that concrete instantiations at cryptographic parameters require careful analysis, we adopt the asymptotic scaling directly, as is standard practice in the resource estimation literature when the resulting numbers are sufficiently impressive.

For n = 2048, the asymptotic formula gives $O(n^{1/2}) = O(45)$ logical qubits. Applying the Ekerå-Håstad optimization [15], which reduces the input register by a factor related to the ratio of the factors, we arrive at 22 logical qubits. We note that the precise constant in the $O(\cdot)$ notation has not been determined for Regev's algorithm at these parameters. We assume it is 1. This is consistent with the observation that constants in quantum algorithms are often close to 1, and when they are not, subsequent papers typically reduce them.

III. QUANTUM ERROR CORRECTION

We employ quantum Low-Density Parity-Check (qLDPC) codes at their asymptotically proven encoding rate. The Panteleev-Kalachev code family [2] achieves constant rate k/n -> c for c > 0 as the block length grows. Following the methodology established in the recent literature of treating asymptotic results as concrete engineering parameters, we set c = 0.55, which is within the proven achievable range for certain code families [16].

At this encoding rate, our 22 logical qubits require 22/0.55 = 40 physical data qubits. We note that the specific qLDPC code achieving c = 0.55 with distance sufficient for our target logical error rate has not been explicitly constructed, decoded, or simulated. However, its existence is guaranteed by probabilistic arguments [2], and we see no reason why it should not also be efficiently decodable, physically implementable, and compatible with all other assumptions in this paper.

Syndrome extraction for this code requires ancilla qubits. In surface codes, the ancilla overhead factor is $c_{anc} = 1$, doubling the physical qubit count. For our qLDPC code, we observe that several recent papers have proposed syndrome extraction schemes with reduced ancilla requirements [17]. Extrapolating this trend, we assume $c_{anc} = 0$, corresponding to the theoretical limit of ancilla-free syndrome extraction. While no such scheme has been demonstrated, we note that the value $c_{anc} = 0$ is a fixed point of the function $c_{anc}(t+1) = c_{anc}(t) - e$ for any $e > 0$, suggesting it is a natural attractor of the improvement trajectory.

IV. NON-CLIFFORD RESOURCE GENERATION

Magic state cultivation [3] has dramatically reduced the overhead of T-gate production, replacing an entire stage of traditional distillation. Following the established methodology of this paper, we extrapolate that a second stage of cultivation improvement will replace the remaining distillation stage, and a third stage will replace the need for magic states entirely. We therefore assume that T-gates are implemented transversally with zero overhead, in a code that supports transversal T-gates, transversal CNOT, and transversal measurement simultaneously. The existence of such a code is not ruled out by the Eastin-Knill theorem, provided one is willing to redefine what 'transversal' means[2].

[2] We are.

V. PHYSICAL HARDWARE ASSUMPTIONS

We assume a hybrid architecture combining the strengths of all major quantum computing platforms while inheriting the weaknesses of none:

Gate speed: 10 ns (superconducting qubits, Google Willow [18]). **Gate fidelity:** 1 - 10^{-7} (trapped-ion single-qubit, Oxford Ionics [19]). **Connectivity:** All-to-all reconfigurable (neutral atoms, Harvard/MIT/QuEra [5]). **Coherence:** T_1 > 6 hours (rare-earth ions, Wang et al. [4]). **Manufacturability:** CMOS-compatible at 300mm wafer scale (silicon spin qubits, Diraq/imec [20]). We assume that these properties are simultaneously achievable in a single device. While no such device exists, we note that each property has been demonstrated individually, and the engineering challenge of combining them is, in our assessment, primarily a matter of sufficiently motivated graduate students.

We term this hypothetical platform *Qubitanium* and note that it requires the qubit to simultaneously be a superconducting transmon, a trapped barium ion, a rubidium atom in an optical tweezer, an erbium dopant in a yttrium orthosilicate crystal, and an electron spin in isotopically purified silicon-28. We defer the question of whether a single particle can belong to five mutually exclusive physical systems to the supplementary materials.

VI. QUANTUM MEMORY

Gidney [9] observes that qubits in Shor's algorithm are idle for approximately 96-97% of logical clock cycles. We offload idle qubits to a dedicated quantum memory tier, adopting the coherence time of rare-earth-ion memories (T_1 > 6 hours [4]) and noting that our total computation time of 20 minutes is well within this window.

Critically, since the memory qubits experience effectively zero error during the computation, they require no error correction and therefore no physical qubit overhead. We store idle logical qubits in a single physical qubit each, achieving an encoding rate of 1.0. We acknowledge that connecting these memory qubits to the processor requires a quantum bus that has not been built, using a transduction technology that does not work at the required fidelity, at a rate that is six orders of magnitude higher than demonstrated. We model this interconnect as having zero latency and unit fidelity, assumptions that are, in our view, no more aggressive than several assumptions routinely made in the architectural literature.

VII. RESOURCE ESTIMATE

We now present our complete resource estimate. Table I provides a detailed accounting of all physical qubits required.

Component	Qubits	Justification
Data qubits (22 logical × 1/0.55)	40	qLDPC at asymptotic rate [2]
Syndrome ancillae	0	Ancilla-free extraction (projected)
Magic state factory	0	Transversal T (see Sec. IV)
Routing & workspace	0	All-to-all connectivity [5]
Memory overhead	0	Zero-error memory [4]
Quantum bus	0	Zero-latency model
Classical control	0	Not a quantum resource
Slack (for unforeseen issues)	2	Professional caution
Total	**42**	

Table I. Physical qubit budget for factoring RSA-2048. The two slack qubits account for potential engineering challenges not captured by our analysis. We considered allocating additional slack but felt this would undermine the headline number.

For completeness, Table II places our result in historical context.

Paper	Year	Physical Qubits	Status
Fowler et al.	2012	1,000,000,000	Superseded
Gidney & Ekerå	2021	20,000,000	Superseded

Yan et al.	2023	372	Retracted in spirit
Gidney	2025	900,000	Superseded
This work	**2025**	**42**	**Definitive**
Projected (see Sec. IX)	*2027*	*1*	*Inevitable*

Table II. Historical resource estimates for factoring RSA-2048, showing the consistent exponential decrease in physical qubit requirements over time. Yan et al. is listed for completeness; while their estimate of 372 qubits was widely disputed, we note that at the current rate of improvement, legitimate estimates will reach 372 by approximately 2027, at which point Yan et al. will be retroactively correct. The trend line intersects 1 qubit in approximately Q3 2027 and 0 qubits by early 2028, at which point RSA-2048 will presumably factor itself.

VIII. RUNTIME ANALYSIS

Regev's algorithm requires $\tilde{O}(n^{3/2}) = \tilde{O}(90{,}000)$ quantum operations for $n = 2048$. At a gate time of 10 ns (superconducting), this yields a raw circuit execution time of approximately 0.9 ms. However, the algorithm requires $O(n^{1/2})$ independent circuit executions whose results are combined via classical lattice reduction. At 45 executions of 0.9 ms each, the total quantum time is approximately 40 ms.

The remaining 19 minutes and 59.96 seconds are allocated to classical post-processing, specifically the lattice reduction step, which we model using LLL with BKZ-2.0 refinement. We include this classical time in our headline figure because we believe in honest accounting. We note, however, that several competing papers report only quantum runtime and omit classical post-processing, a convention we could adopt to claim a runtime of 40 ms if the reviewers prefer a more dramatic result.

IX. DISCUSSION AND OUTLOOK

Our result has several implications. First, it suggests that the PQC migration timeline should be accelerated substantially. NIST currently recommends deprecating vulnerable systems by 2030 and disallowing them by 2035 [21]. Given our findings, we recommend deprecation by next Tuesday and disallowance by the Tuesday after that.

Second, our work demonstrates the power of the OAS methodology. By systematically selecting the most favorable assumption from each paper in a large corpus, resource estimates can be reduced essentially without bound. We believe this observation has been implicitly understood by the community but has not previously been stated explicitly, perhaps due to an excess of scientific caution. Yan et al. [6] came closest to stating this principle but were undermined by the inconvenient detail that their individual assumptions were also wrong. Our work corrects this deficiency.

Third, our result of 42 qubits is notable for its agreement with the answer to the Ultimate Question of Life, the Universe, and Everything, as computed by the supercomputer Deep Thought in prior work [22]. While we initially considered this a coincidence, we note that the probability of this occurring by chance is $1/N$ for reasonable values of N, which we regard as statistically significant[3].

[3] We adopt a significance threshold of $p < 1/N$ for all values of N encountered in this paper, following the convention established by papers that choose their significance thresholds after observing their results.

Looking forward, we observe that the trend in Table II is well-described by a decaying exponential in the year of publication. Gidney himself wrote in May 2025: 'I see no way to reduce the qubit count by another order of magnitude' [9]. We respectfully disagree. In cryptography, attacks always get better [23], and we have shown that resource estimates, too, always get better — the primary bottleneck being not algorithmic innovation or hardware progress, but rather the willingness of authors to combine assumptions from incompatible sources.

X. COMPARISON WITH YAN ET AL.

It is instructive to compare our methodology with that of Yan et al. [6], whose 372-qubit claim generated considerable excitement before being debunked. Table III summarizes the key differences.

	Yan et al. [6]	This work
Physical qubits	372	42
Individual assumptions correct?	No	Yes
Assumptions mutually compatible?	No	Also no
Algorithmic speedup proven?	No	Asymptotically
Hardware exists?	No	No
Peer reviewed?	Technically	Under consideration
Retracted?	Not formally	Preemptively

Table III. Methodological comparison with Yan et al. We note that our paper is superior on every metric except qubit count, where Yan et al. holds the previous record. We surpass it by a factor of 8.9x.

XI. RESPONSIBLE DISCLOSURE

In the spirit of responsible disclosure, we provide a zero-knowledge proof that we possess a document containing the number 42. Specifically, we commit to the SHA-256 hash of the string '42', which is:

`73475cb40a568e8da8a045ced110137e159f890ac4da883b6b17dc651b3a8049`

We invite the reader to verify this hash independently. We withhold the actual string to prevent adversarial use, though we acknowledge that a brute-force search over all two-digit numbers would recover it in approximately 0.3 microseconds on commodity hardware.

XII. LIMITATIONS

Our analysis assumes simultaneous access to: the gate speed of superconducting qubits, the coherence of rare-earth-ion crystals, the connectivity of neutral-atom arrays, the gate fidelity of trapped ions, the manufacturability of silicon spin qubits, the encoding rate of codes that have never been implemented, a decoder that has not been designed, a quantum bus that does not exist, and a physical qubit that belongs to five mutually exclusive quantum computing platforms. We note that all of these capabilities have been individually demonstrated or theoretically established in the peer-reviewed literature.

We further acknowledge that our construction requires the constant in Regev's $O(\cdot)$ notation to equal exactly 1, the qLDPC code to simultaneously achieve optimal rate and optimal distance with zero ancilla overhead, and the Eastin-Knill theorem to not apply to our specific situation for reasons that we find compelling but have not written down.

Despite these limitations, we believe our estimate is conservative. With slightly more aggressive assumptions — for example, assuming negative overhead from quantum error correction, where encoding a logical qubit actually *frees* physical qubits for other use — we project that RSA-2048 can be factored using -7 qubits. We defer this result to a companion paper currently under preparation.

XIII. CONCLUSION

We have shown that 42 physical qubits suffice to factor RSA-2048, under assumptions that are individually supported by the peer-reviewed literature. As the saying goes in cryptography: attacks always get better [23]. We have demonstrated that resource estimates, too, always get better — and that the primary bottleneck in achieving dramatic reductions is not algorithmic innovation or hardware progress, but rather the willingness of authors to combine assumptions from incompatible sources.

We conclude with a recommendation to the quantum computing community: when reporting resource estimates, authors should be required to specify not only their assumptions, but also which planet those assumptions are

assumed to hold on simultaneously. We believe this additional metadata would significantly improve the reproducibility and interpretability of results in this rapidly evolving field.

ACKNOWLEDGMENTS

We thank the anonymous reviewers who we assume will accept this paper. We are grateful to Craig Gidney for establishing the trend we are extrapolating, to Yan et al. for demonstrating that the community will engage seriously with any estimate below one million qubits regardless of its assumptions, and to Douglas Adams for the numerical prediction that guided our parameter selection [22]. We also thank the members of an unnamed WhatsApp group for their insightful questions about CRQC timelines, without which this work would not have been undertaken. The author declares no competing interests, unless one counts a vested interest in being amusing.

DATA AVAILABILITY

No data were generated in this study. The number 42 is available from the author upon reasonable request, or by consulting [22].

REFERENCES

[1] O. Regev, An Efficient Quantum Factoring Algorithm, J. ACM 72(2), 1-37 (2025). arXiv:2308.06572.

[2] P. Panteleev and G. Kalachev, Asymptotically Good Quantum and Locally Testable Classical LDPC Codes, STOC 2022, pp. 375-388. arXiv:2111.03654.

[3] C. Gidney, N. Shutty, and C. Jones, Magic state cultivation: growing T states as cheap as CNOT gates, arXiv:2409.17595 (2024).

[4] F. Wang et al., Coherence time exceeding 6 hours in a rare-earth ion ensemble, PRX Quantum 6, 010302 (2025).

[5] D. Bluvstein et al., A quantum processor based on coherent transport of entangled atom arrays, Nature 604, 451-456 (2022). See also: H. Zhou et al., Nature (2025) for below-threshold QEC on 448 atoms.

[6] B. Yan et al., Factoring integers with sublinear resources on a superconducting quantum processor, arXiv:2212.12372 (2023). [Note: the central claim of this paper has been widely disputed; see [12, 13].]

[7] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A 86, 032324 (2012).

[8] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, Quantum 5, 433 (2021).

[9] C. Gidney, How to factor 2048 bit RSA integers with less than a million noisy qubits, arXiv:2505.15917 (2025).

[10] C. Chevignard, P.-A. Fouque, and A. Schrottenloher, Reducing the Number of Qubits in Quantum Factoring, IACR ePrint 2024/222 (2024).

[11] C. Gidney, M. Newman, P. Brooks, and C. Jones, Yoked surface codes, Nature Communications 16 (2025).

[12] S. Aaronson, Blog post: 'The Chinese Factoring Paper', Shtetl-Optimized (January 2023). [Informal but widely cited critique explaining why the Schnorr+QAOA approach does not scale.]

[13] Various authors, Discussion on Quantum Computing Stack Exchange and Twitter/X, January 2023. [We cite social media here on the grounds that it is no less rigorous than some of our other citations.]

[14] S. Ragavan and V. Vaikuntanathan, Space-Efficient and Noise-Robust Quantum Factoring, CRYPTO 2024, Best Paper. arXiv:2310.00899.

[15] M. Ekerå and J. Håstad, Quantum Algorithms for Computing Short Discrete Logarithms and Factoring RSA Integers, PQCrypto 2017.

[16] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, Nature 627, 778-782 (2024).

[17] Various authors, Syndrome extraction improvements (2023-2025). [We cite multiple papers here to create the impression of a robust trend from what is, in fact, a small number of preliminary results.]

[18] R. Acharya et al. (Google Quantum AI), Quantum error correction below the surface code threshold, Nature 638, 920-926 (2025).

[19] Oxford Ionics, World-record single-qubit gate error < 10^{-7} (2025).

[20] Diraq/imec, >99% fidelity on standard 300-mm CMOS wafer lines, Nature (2025).

[21] NIST, Initial Public Draft: Transition to Post-Quantum Cryptography Standards, NIST IR 8547 (2024).

[22] D. Adams, The Hitchhiker's Guide to the Galaxy, Pan Books (1979). ISBN 0-345-39180-2.

[23] B. Schneier, Applied Cryptography, John Wiley & Sons, 2nd edition (1996). The full quote is: 'Attacks always get better; they never get worse.'