Map of *Certified Speed-Limit Diagnostics for Deployed Superconducting Quantum Processors*

Quantum speed limits · superconducting-qubit gate calibration

54 results · 8 entirely in Lean · 11 partly in Lean · 2 in Isabelle · 36 checked by a program · written proof only (a result can carry several tags)

Dependencies in the written proof: an arrow A → B means that this proof of B uses A. Every result rests on its written proof; the tags record what was checked besides, and each statement's "Evidence & scope" block says how far. The colour of a box is its strongest tag, in the order above; the tags in the box give the rest. Boxes with a thick border are the main results; click a box to read the statement.

Introduction

Remark Abstract
the abstract (1.1)–at full amplitude, τ ≥ at the drive-IBM uses, gates > 10× faster than the coherence-limit, the device's full amplitude, 4.22× for two-qubits, a CZ between 0.295 and 0.331 of the qubit-only time (version 3)

Remark §1
what the paper claims: the problem (section 2), fifth in version 21, and no new physics

Certified Bounds

Remark Table 1
Table 1: the certified statements used in the paper (Lean 4, commit 42f307f), and the open questions for the literature
◐ part in Lean

Theorem §B-A
Mandelstam–Tamm, 1: form: 1 ≤ Δ·⟨ΔE⟩
Lean

Theorem §B-A
Margolus–Levitin: e^{-2iTH}|ψ⟩ = |ψ⟩ ⇒ T ≥ π/(2⟨H⟩)
Lean Isabelle

Theorem §B-A
Zielinski–Zych: π^n ≤ 2^M ⇒ n^n for every n > 0
Lean Isabelle

Theorem §B-B
entangling time in any local dimension: arccos λ / 2 Σλᵢ
Lean

Theorem §B-B
the earlier chord-distance version: 2 Σλᵢ ≥ 2(1−√2)
Lean
Lean

Theorem §B-A
two-level drive: τ ≥ π/Ω₀ₙ for every pulse shape; equality for a perfectly square-limit zero-exactly
◐ part in Lean

Corollary §B-B
a Bell vector needs 2 Σλᵢ ≥ π/4, in any local dimension d ≥ 2
Lean

Proposition §B-B
the bound is attained: H = J ZᵢZⱼ/4 reaches a Bell vector in time π/(4J)
Lean

Corollary §B-B
any CNOT, or gate locally equivalent to it (UCR), on the qubit subspace of two d-level systems needs 3 π/4
◐ part in Lean program

Theorem §B-B
robust bounds: a Bell vector to fidelity 1 − ε needs T ≥ (π/4 − arcsin √ε, a two-level transfer needs Ω₀ₙτ ≥ 1.2 ε means ¼
Lean

Data and Conventions

Observation §III-C
the drive convention is ε = 1: IBM's calibrated π pulse rotates by 1.148 θ with ε = 1 (2.296 π with ε = 2)
program

Observation §III-C
the residual is explained: unequal jumps' Λ₁e^{−Γ₁τ} on all 777 qubits to 1.2·10⁻³; the ratio of the two areas tends to 1/(1 + e^{γ−1} ≈ 1.5042)
program

Observation §III-C
omega0 is not an independent measurement: IBM derives it from the π-pulse calibration
program

Observation §III-D
the pre-registered fidelity 254 ns, 1000 shots, seeded with placeholder-like values, which are excluded
program

Single-Qubit Gates

Observation §IV-A
X gates last a median 11.30 times their two-level speed limit at full drive amplitude
program

Observation §IV-C
IBM's design rule: the leakage-to-error ratio frequency is about 0.09 of the anharmonicity (0.0907, or 0.078 with the normalized peak)
program

Observation §IV-C
the residual is explained: the gap factors into a pulse shape (≈ 2.16 − 1.87 = 1/(1 − e^{−γ}) and an amplitude reserve (≈ 5)

Observation §IV-B
pre-registered rivals: H1 (the gap is shape-limited (depends on the waveform reading at the drive used), H3 excluded, H2 without support
program

Observation §IV-C
pre-qubit-resolution under the 1.08× ratio more above the cap
program

Two-Qubit Gates

Observation §V-A
1758 coupled pairs in 43 backends: median |g| = 0.01268 rad/ns, median gate 548 ns
program

Observation §V-C
measured against the certified bound: 0 violations in 20 000 and 3000 trials
program

Observation §V-B
ECR/CX gates last a median 4.22 times the CNOT-class bound at Ω|g|
program

Observation §V-C
against our certified constants the median ratio is 1.08 (max 1.31); (sharp at 0), twice the literature one
◐ part in Lean program

Observation §V-D
cross resonance: against its first-order formula: median θ(Δ/2) = 1.109; 13. CR/Δ > 0.45

Time Series 2021–2025

Observation §VI
single-qubit ratio frozen near 11 from 2021 to 2025; two-qubit ratio 2.78 → 4.26
program

Observation §VI
the two-qubit ratio grew because gates got longer, not because the coupling changed
program

Observation §VI
lengthening does not lower the error: Spearman +0.13 (p = 0.23); it falls in only 31 % of lengthened pairs
program

Observation §VI
drive-power effect (p = 0.0017) vanished under a stratified control (p = 0.98): consistent with confounding by epoch
program

How Far the Model Allows

Observation §VII-A
optimal control: shortest X time found = max(τ_QSL, 1.1·2π/Ω), a speed-limit branch and a leakage-dominated branch
program

Observation §VII-A
for IBM's |δ| = 2.07 Ω₀ₙ, the shortest gate found takes ≈ 16 ns at its own cap (35.6 ns is ≈ 2.2 times that); at τ = 1 the 4-level search finds 3.3 ns, but that value is truncation-dependent
program

Proposition §VII-A
a lower bound for any X gate: at most 1 − 4/3 ≈ 0.87 can, and in the trace metric T ≥ 2 arccos(1+|Tr(U_ideal† U_true)|/4)/Ω₀ₙ, for τ ≥ 0.2 the times found are 6 % above it
program

Observation §VII-A
truncation: the 4-level pulse at τ = 1 degrades to 3.5·7.2·10⁻⁴ with 5 to 12 levels; a 4-level pulse at T = 7.0(4.1) 4 ns) reaches 4·4·10⁻⁴ here and 1.6·10⁻⁴ with 6 and 12 levels
program

Observation §VII-B
the shortest CNOT found is 1.2 times slower at IBM's drive fraction 0.45 than at 0.7
program

Observation §VII-A
the full amplitude |Ξ| = 1 is ofold = omega(d); between 0.9 and 1.1 it is 8.4 (4.8 to exact (bound 4.3 on); the median gate is 19–11 times the model minimum, a 35.6 ns Falcon gate 7.4–8.3 times
program

Observation §VII-B
the shortest found: a fast twisting of one transmon keeps (I⊗Q)|Q, the code in the extreme eigenvectors of Q gives (1/√2)(|Ω,0⟩ ± i|−Ω,τ⟩), with a residual 4·10⁻⁵ and 4·10⁻⁴ here; at τ = 0.863, and with the refined bound (not under) 0.451113 ≤ T* < 0.885
◐ part in Lean program

Observation §VII-B
a gate bound by the Choi rank: T ≥ 0.3927 (the answer2 = (1/2)[3; so 0.3929961 ≤ T_*^2·10^{-3} ≤ 0.553 (sharpened to 0.451013 in eventhead); exact Bell preparation is a different, easier task
◐ part in Lean program

Observation §VII-B
CZ with the second transmon level: error 2·10^{-4} at T = 0.553 with local controls ≤ 10ℓ, i.e. 0.35 of the qubit-only time and 5 % above the construction π/Δ13
◐ part in Lean program

Observation §VII-B
the 0.45113 is comparator-trajectory argument; the 93.02 terms is ≈ 0.0110 (the gap 0.885 − 0.553 = T_*^2·(Δ/Ω)^2 + 0.2261 and T_*(2·10^{-3} ≥ 0.45103
◐ part in Lean program

Observation §VII-B
a refined CZ bound: T_*^2 ≥ 0.451113 (a comparator-trajectory argument; the 93.02 terms is ≈ 0.0110 (the gap 0.885 − 0.553 = T_*^2·(Δ/Ω)^2 + 0.2261 and T_*(2·10^{-3} ≥ 0.45103
◐ part in Lean program

Observation §VII-B
version 3: two copies give 0.4511 ≤ T_* ≤ 0.46223, a closed excursion through (max), certified in interval arithmetic, given T_* ≥ 0.5200/2.5; so the CZ lies between 0.295 and 0.331 of the qubit-only time
program

Observation §VII-B
Isabelle/HOL: the Jacobi–Lie auxiliary: the smallest-largest-Schmidt-coefficient is exactly cos(2T) up to the Bell time with (an analytic witness)
◐ part in Lean program

Application: A Chip-to-Chip State Transfer

Remark §VIII
the application: transfer 80 ion ions, one standing mode, g/2π ≈ 8–9 MHz, T_2* = 256–312 ms, F = 98.86–99.29 %
program

Remark Table VII
Table VII: the IBM Falcon r5.11L and the single-mode minimal time (4.1–5.6× sequential, 11.6–15.9× the certified floor); the 128 ns target would still be 2.9–3.2×
program

Observation §VIII
at 128 ns the budget is 0.22 of g1 for the ratio, from the any-dimension speed limit with fill ≤ 0.8
◐ part in Lean

Observation §VIII
at 128 ns the coherence error of the optimal protocol is 0.13–0.26 % (master equation), 4–6× below the decoherence budget; the additive model gives 0.6 % (0.09–0.17 %): exact for relaxation, not for dephasing; one protocol, not a floor
program

Observation §VIII
the optimal protocol: optimal time is π/(2 g) (Bouxiin–Chambless–Charles 2005), in lower half g/2 ; π/2 is machine-checked in Isabelle
Isabelle program

Remark §VIII
conclusion: the modes (the bounds (hex hold with the collective norm G), g is the maximal coupling. IBM's fidelity and error model, Markovian dephasing at the Ramsey rate

What the Certificate Does and Does Not Guarantee

Remark §IX
the certified objects are the inequalities, not the paper, and the results in Lean and the assumptions are listed

Reproducibility

Remark §X
every number traces to a script in the artifact package, linked from this blueprint
program

Conclusion

Remark §XI
conclusion: 11.3× at full amplitude and > 2× at the drive used, X gates > 10× faster than the median within the full amplitude (4.22× for two-qubits, a CZ between the refined bound 0.45113 and the construction π/Δ3)

Remark §XI
conflict of interest and the use of AI assistance

What the paper contributes

Two kinds of results, kept apart on this map. **Certified inequalities** (green, linked to Lean): the Margolus–Levitin bound, an L^1 form of Mandelstam–Tamm, the Zieliński–Zych family for every exponent $\alpha>0$, time-dependent two-level bounds, and an entangling-time bound $J\,T\ge\arccos\lambda$ in arbitrary local dimension (a Bell state needs $J\,T\ge\pi/4$, attained; robust versions with state infidelity ε). **Diagnostics on deployed hardware** (grey, with the program that computes each number): how far the calibrated gates of IBM superconducting processors are from those bounds, from public, version-pinned calibration snapshots.

- Single-qubit X gates last a median 11.30 times their two-level speed limit $\pi/\Omega_{\max}$ (1646 qubits, 41 backends), a ratio frozen near 11 from 2021 to 2025.
- Two-qubit ECR/CX gates last a median 4.22 times the known CNOT-class bound $\pi/(2j_q)$; the ratio grew from 2.78 (2021) to 4.26 (2025) because gates got longer, not because couplings got weaker.
- Calibrated drives sit at fixed fractions of an energy splitting: $\Omega_{\rm peak}\approx0.09|\delta|$ (one qubit), $\Omega_{\rm CR}\approx0.45\Delta$ (two).
- The drive-strength convention is fixed by IBM's calibrated π pulses; the $\approx15\%$ residual is explained by the pulse normalization: $\texttt{omegad}\cdot|\mathrm{amp}|\cdot A_{\rm rec}=\pi$ on all 777 qubits to $1.2\cdot10^{-7}$, and the residual tends to $1/(1-e^{-2})=1.15652$.
- At the drive IBM actually uses the X gate is within about 2 of the square-pulse limit (shape factor 1.87, or 2.16 with Qiskit's peak normalization); the 11.3 is against full drive amplitude, which IBM does not use.
- Optimal control on IBM's own transmon model finds feasible X gates of ≈16 ns at IBM's amplitude cap and 4.8 ns at its full amplitude (minimum at least 4.3 ns): 7.4–8.3 times faster than the archived 35.6 ns within the same amplitude budget; ≈3.4 ns needs three times that amplitude (numerical witnesses, not physical floors; the 6-level pulse is converged to 12 levels). The fit $\max(\pi/\Omega_{\max},\,\sim1.1\cdot2\pi/|\delta|)$ is empirical, not a theorem; $T\ge\pi/\Omega_{\max}$ (and $2\arcsin(1-e)/\Omega_{\max}$ in the trace metric) holds for any number of levels.
- With the second transmon level the minimal CZ time lies between 0.295 and 0.331 of the qubit-only time: $0.4632/J\le T_*\le0.5200/J$ (version~3: a two-copy bound and a certified protocol; before, $0.4511/J\le T_*\le\pi/(6J)$; the certified norm bound is $\pi/(8J)$); optimal control reaches error $2\cdot10^{-8}$ at $0.55/J$ (0.35), and for the same task $0.4510/J\le T_*(2\cdot10^{-8})\le0.55/J$. In larger local dimension the construction and the entangling-time theorem approach each other (ratio $\to1$, certified in Lean).

The paper claims no new physics: the speed limits are classical, the two-qubit constants are known for qubits, and the transmon model is IBM's. The entangling-time bound is, for qubits, a known corollary of Dür *et al.* and Vidal–Hammerer–Cirac, new only as a machine-checked, any-dimension assembly. This blueprint is built from a **draft**: the paper's own open items are shown where they matter.

🚀 **Start here:** **Observation §IV-A**: X gates last a median 11.30 times their two-level speed limit at full drive amplitude.

Main results:

- **Theorem §II-A**: Zieliński–Zych: $\pi^{\wedge}\alpha \leq 2M_\alpha\, \tau^{\wedge}\alpha$ for every $\alpha > 0$
- **Theorem §II-B**: entangling time in any local dimension: $\arccos \lambda \leq J \cdot \Sigma\tau_k$
- **Corollary §II-B**: a Bell vector needs $J \cdot \Sigma\tau_k \geq \pi/4$, in any local dimension $d \geq 2$
- **Observation §III-C**: the drive convention is $c = 1$: IBM's calibrated π pulse rotates by $1.148\,\pi$ with $c = 1$ ($2.296\,\pi$ with $c = 2$)
- **Observation §IV-A**: X gates last a median 11.30 times their two-level speed limit at full drive amplitude
- **Observation §V-B**: ECR/CX gates last a median 4.22 times the CNOT-class bound $\pi/(2jq)$
- **Observation §VI**: single-qubit ratio frozen near 11 from 2021 to 2025; two-qubit ratio 2.78 → 4.26
- **Observation §VII-A**: optimal control: shortest X time found $\approx \max(\pi/\Omega_{max}, 1.1 \cdot 2\pi/|\delta|)$, a speed-limit branch and a leakage-dominated branch
- **Observation §VII-B**: CZ with the second transmon level: error $2 \cdot 10^{-8}$ at $T = 0.55/J$ with local controls $\leq 10J$, i.e. 0.35 of the qubit-only time and 5 % above the construction $\pi/(6J)$

🧭 Fields: where the paper lives and what it connects

Tools taken from other branches

- **Quantum speed limits** → Margolus–Levitin, Mandelstam–Tamm and the Zieliński–Zych family, applied to the drive amplitude of a transmon and to the coupling of a pair — Theorem §II-A, Theorem §II-A, Theorem §II-A, Theorem §II-A
- **Entangling capabilities of Hamiltonians** → minimal interaction time under unbounded local control; Schmidt coefficients and the Fubini–Study angle (Dür et al., Vidal–Hammerer–Cirac, Rudnicki) — Theorem §II-B, Corollary §II-B, Proposition §II-B, Corollary §II-B
- **Formal verification** → Lean 4 + Mathlib: the inequalities of Table I — Theorem §II-A, Theorem §II-A, Theorem §II-A, Theorem §II-B, Corollary §II-B, Proposition §II-B, Theorem §II-B
- **Optimal control** → piecewise-constant GRAPE control on a 4-level transmon with a bound on the drive — §VII-A, §VII-A, §VII-A, §VII-A, §VII-B

What the results serve

- **Gate design on superconducting processors** ← how far calibrated gates are from the physical limit, and which fractions of the energy splittings the calibrations use — §IV-A, §V-B, §IV-C, §V-D
- **Reproducible hardware benchmarking** ← version-pinned public calibration snapshots as a citable source, and a stratified control against confounding by epoch — §VI, §VI