Quantum state transfer through a lossy bus: fundamental error limits for every control pulse, with application to IBM's meter-scale interconnects

CARLES MARÍN

Abstract. How small can pulse shaping make the error of quantum state transfer between two qubits, such as superconducting qubits in different modules, through a lossy bus (a resonator, a waveguide or a cable mode)? We give bounds that hold for every control pulse, every duration and any number of bus modes, knowing only the maximal coupling strength G, the qubit loss rate γ and the bus loss rate $\kappa \geq \gamma$. The population delivered in time T obeys $P(T) \leq e^{-\gamma T}/[1+(\kappa-\gamma)/(G^2T)]$: going slowly costs qubit loss, going fast costs bus loss, and no pulse escapes both. The end qubits cannot be bypassed either: in the frame where the qubits are lossless, every transfer keeps the excitation in them for an integrated time bounded below by a quantity that tends to $\pi/(2G)$ as the transfer becomes complete, and swap–wait–swap attains that bound. As $\gamma/(\kappa-\gamma) \to 0$ and $\sqrt{\gamma(\kappa-\gamma)}/G \to 0$, the error cannot fall below $2\sqrt{2\gamma(\kappa-\gamma)}/G\,(1-o(1))$; this constant was known as the optimum of perturbative and quasi-adiabatic pulse optimisation (Wang et al. 2017; Luo, Fang, Wang and Chesi 2024); here it bounds every control, and an explicit pulse attains it with explicit error terms. Markovian dephasing enters through the same equations with rates $2/T_2$, so the bounds also limit the channel fidelity; with a different coherence time for each element, a dissipative version of the occupation bound gives the floor of a real device, and a reciprocity relation shows that the direction of transfer does not change the optimal error. For the four links of IBM's relay experiment (arXiv:2610.29986), at $g/2\pi = 9$ MHz and with one effective bus mode and the coherence times of each element, no pulse can bring the error below 6.03–10.05×10^{-4} (certified with interval arithmetic), if dephasing is Markovian at the measured T_2^* (5.1–14.1×10^{-4} if all coherence times change together by $\pm 20\%$ and $g/2\pi$ is 8–9 MHz), and not below 1.5–1.8×10^{-4} for any dephasing noise, Markovian or not. A 39-ns pulse comes within a factor 1.4–1.5 of that floor, against 71–114×10^{-4} measured; on one link it even beats the floor computed with IBM's symmetric coherence model, which therefore does not give a floor for the device; doubling the coupling halves the floor, the dephasing of the worse qubit is the next lever, and with couplers that reach full strength in about 3 ns the floor is only about 10% above that of ideal couplers; within the model, a 26-ns trapezoidal pulse at 18 MHz already brings all four links below 10^{-3}, while 10^{-4} also needs better coherence times (we give one sufficient combination). The same bounds put three other published transfers 4.5–7.6 times above their floors. The first bound, the occupation lemma, the frontier lemma and the window comparisons are machine-checked in Isabelle/HOL, and every number comes with the script that produces it; a companion laboratory computes the floor of any link.

Date: October 31, 2026.

Key words and phrases. Quantum state transfer, quantum bus, lossy resonator, coaxial cable, quantum interconnect, STIRAP, optimal control, fidelity bound, error floor, dephasing, quantum speed limit, superconducting qubits, modular quantum computing, network benchmarking, formal verification, Isabelle/HOL.

Contents

1. Introduction 3
2. Main results 4
3. Model and hypotheses 6
4. The floor 7
5. The end qubits must be occupied 8
6. Partial transfer and the asymptotic constant 9
7. Dephasing 12
8. Application to IBM's relay 14
8.1. Which hardware change lowers the floor 16
8.2. Other platforms 18
9. Prior work and what is new 19
10. Limitations 20
11. Verification and supplementary files 21
Appendix A. Formula sheet: which bound to use, and when 21
Use of AI tools 23
References 23

Companion site. https://karlesmarin.github.io/blueprints/relay-floors/: every result of this paper with its proof status (Isabelle sources, certificates, scripts), the figures, and a laboratory, Check your link, that computes the floor of any qubit–bus–qubit link from its coherence times, the bound at the duration of a measured pulse and a verdict on the measured error, with the same code as Table 1. This preprint: doi:10.5281/zenodo.23836450. Part I: doi:10.5281/zenodo.23783382.

1. Introduction

Suppose you have built a link between two superconducting qubits through a shared bus—a resonator, a waveguide or a cable mode—and you measure a transfer error of 10^{-2}. How much of it could a better pulse remove, and how much is fixed by the hardware? This paper answers that question with floors that no pulse can go below, valid for every pulse shape, phase and duration and for any number of bus modes. They depend only on the maximal coupling strength and the loss and coherence rates, so they also tell which hardware parameter lowers the floor most (Section 8.1), and they apply to other platforms (Section 8.2).

The physics behind the floors is a familiar trade-off. The bus is usually the lossy element, so protocols such as STIRAP and its shortcuts keep it nearly empty; the price is time, during which the qubits decay and dephase. Its perturbative optimum is known: for small losses ($\kappa \gg \gamma$) the minimal loss $1 - P$ is $2\sqrt{2\kappa\gamma}/G$, reached at the duration $T_{\mathrm{asym}} = \sqrt{2\kappa/\gamma}/G$ [2, 3]. What has been missing is a statement for *every* control, at *finite* duration: an optimum within a family of pulses does not tell whether another pulse does better.

This paper gives such statements and applies them to hardware.

- *A floor at every duration* (Section 4): $P(T) \le e^{-\gamma T}/[1 + (\kappa - \gamma)/(G^2 T)]$ for every control, with any number of bus modes, machine-checked; for the population problem at the loss ratio of an IBM link ($\kappa/\gamma = 4.2$) it is within 4–12% of the best pulses found.
- *The end qubits cannot be bypassed* (Section 5): a dark-state protocol empties the bus, but the excitation must then sit in the end qubits for an integrated time bounded below, tending to $\pi/(2G)$ for a complete transfer.
- *The asymptotic constant* (Section 6): the perturbative optimum $2\sqrt{2\gamma(\kappa - \gamma)}/G$ becomes a lower bound over all controls, and an explicit open-loop pulse attains it with explicit error terms.
- *Dephasing and real devices* (Section 7): the coherence with the vacuum is an amplitude of the same problem with rates $2/T_2$, which turns loss bounds into bounds on the fidelity F_B that network benchmarking measures; a version with a rate for each element (emitter, bus, receiver) gives the floor of a device whose elements differ, and a reciprocity relation shows that the direction of transfer is not a design choice.
- *IBM's meter-scale links* (Section 8): floors of 6.03–10.05×10^{-4}, certified with ball arithmetic, against measured errors 11–16 times larger; a floor computed with the symmetric coherence model used for comparison in Ref. [4] is not a floor for the device (a 39-ns pulse beats it on one link); a simple minimal-time pulse comes within a factor of about 1.5 of the floor, and in the examples we computed what matters is the time spent in the noisy elements, not the total duration.
- *Which hardware change helps* (Section 8.1): the floor ranks changes before any pulse is designed: the coupling first, then the dephasing of the worse qubit; the bus mode matters little. We quantify the share between the two couplers, the cable length, several modes, the slew rate of the couplers, and give pulses and coherence targets for 10^{-3} and 10^{-4}.
- *Other platforms* (Section 8.2): three published transfers lie 4.5–7.6 times above the floors of their effective one-mode models.

Every bound states how it is checked: Isabelle/HOL for the floor, the endpoint occupation, the frontier lemma and the window comparisons; interval arithmetic for the IBM floors; and a script for every number (Section 11). Appendix A collects the formulas and when to use each.

This is the second part of a study of how far deployed superconducting hardware sits from the limits set by quantum mechanics. Part I [1] measured the gap between calibrated single- and two-qubit gate durations and their speed limits on IBM processors, with machine-checked inequalities. Here the operation is the transfer of a state between modules, and the limit is an error floor rather than a minimal time: Part I asks how fast, this paper how well. Both rest on the same Isabelle entry (Section 11). Section 2 states the results and how each one is certified; the proofs follow in Sections 4–7, the IBM numbers in Section 8, and what was known in Section 9.

2. Main results

Setting in one line. One excitation starts in qubit A and must reach qubit B through a bus that loses energy at rate κ; the qubits lose it at rate $\gamma \le \kappa$; each coupler can be driven up to strength G (G bounds the norm of each qubit's coupling vector to the bus modes, not each mode separately) with any shape, phase and duration; the bus may have any number of modes (Section 3). $P(T)$ is the population that reaches B at time T.

Result 1 — the floor (Theorem 4.1)

What it says. For every pulse and every duration,

$$P(T) \le \frac{e^{-\gamma T}}{1 + (\kappa - \gamma)/(G^2 T)}$$

The numerator is the qubit loss over the whole transfer: going slowly costs. The denominator is the bus loss: going fast costs too. No pulse escapes both. For small losses, $1 - P \gtrsim \gamma T + (\kappa - \gamma)/(G^2 T) \ge 2\sqrt{\gamma(\kappa - \gamma)}/G$.

Why it holds. In the frame where the qubits are lossless the bus still loses at rate $\kappa - \gamma$. To move the excitation in time T with couplers of strength G, the bus must hold it for an integrated time of at least $\tilde{P}/(G^2 T)$, with $\tilde{P} = e^{\gamma T} P$ the population in that frame (Cauchy–Schwarz), and what the bus holds it loses.

Status. Proof in Section 4. Machine-checked for derivatives outside a countable set (piecewise-smooth controls): Isabelle/HOL, theorem `loss_floor` in theory `Relay_Loss_Floor` (entry `Quantum_Speed_Limits`, public in the repository karlesmarin/formalizations).

Data. The best loss found by optimal control lies above the bound in every case tested (Figure 1); at the loss ratio of an IBM link ($\kappa/\gamma = 4.2$) the best pulse found is within 4–12% of it. Scripts `suelo_exacto.py`, `suelo_error.py`.

Result 2 — the end qubits cannot be bypassed (Theorem 5.2)

What it says. In the frame in which the qubits are lossless, every transfer that empties A satisfies, for every T,

$$G \int_0^T \big(|a(t)|^2 + |b(t)|^2\big)\, dt \ge \frac{\pi}{4} + \Phi\big(|b(T)|\big) \xrightarrow{\,|b(T)|\to 1\,} \frac{\pi}{2}$$

where a, b are the amplitudes in A and B. Without bus loss a complete transfer is possible and swap–wait–swap attains $\pi/2$. Protocols such as STIRAP keep the bus nearly empty; this is the other half of the picture: the excitation must sit in the end qubits for an integrated time whose lower bound tends to $\pi/(2G)$ as the transfer becomes complete, however slowly or cleverly it is done.

Why it holds. The function $\Phi(u) = \frac{1}{2}(\arcsin u - u\sqrt{1-u^2})$ cannot change faster than G times the occupation u^2, because the flow out of a qubit is limited by G times the bus amplitude, which is at most $\sqrt{1-u^2}$. Emptying A and filling B each move Φ by $\Phi(1) = \pi/4$.

Status. Proof in Section 5. Machine-checked in Isabelle/HOL (theory `Relay_Endpoints`) for absolutely continuous amplitudes with exceptions in a null set (every measurable bounded control): the window lemma `endpoint_window_ac` (Lemma 5.1), of which the inequality above is a two-line corollary, and the complete-transfer form `endpoint_occupation_ac`.

Data. 300 random controls (1–3 modes, complex phases, losses, partial transfers): left side over right side of the general inequality at least 2.11 (Figure 3); $\Phi(1) = \pi/4$ and the value $\pi/2$ of the swap protocol in exact arithmetic (Sage). Scripts `falsar_extremos.py`, `verificar_v201.py`, `certificado_sage.sage`.

Result 3 — the known constant becomes a bound, and is attained (Theorems 6.2 and 6.3)

What it says. Write $\ell_0 = 2\sqrt{2\gamma(\kappa-\gamma)}/G$, $\alpha = \sqrt{\gamma(\kappa-\gamma)}/G$ and $\beta = \gamma/(\kappa-\gamma)$. For small losses ($\alpha, \beta \to 0$) every pulse has

$$1 - P \geq \ell_0 \left(1 - o(1)\right)$$

a factor $\sqrt{2}$ above Result 1; and for $\alpha, \beta \leq 0.01$, with a bus mode whose loss is exactly κ, an explicit pulse reaches $1 - P \leq \ell_0[1/(1-\alpha) + 1.5\beta]$. The constant was known from perturbative and quasi-adiabatic optimisation [2, 3]; here it is the asymptotically exact coefficient of the minimal loss over all controls. At finite parameters the best loss can lie slightly below ℓ_0 (Figure 1).

Why it holds. The excitation must cross from A to B; at the crossing both qubits hold the same amplitude, and reaching that point and leaving it doubles the bus occupation needed when GT is large (Lemma 6.1). The pulse that attains it keeps the rescaled bus amplitude fixed at $\sqrt{\beta}$ (Figure 4), where the two losses balance.

Status. The two inequalities of the lemma are machine-checked (Isabelle/HOL, theorem `partial_frontier` in `Relay_Loss_Floor`, from the reduced inequalities, exceptions in a countable set); the function Ψ and its limit, and the lower bound, are proved on paper, without explicit error terms. That the constant is attained is proved on paper with explicit error terms (Theorem 6.3).

Data. Lemma checked in 10/10 cases; the explicit pulse, integrated in the full lossy dynamics, reaches 1.00004–1.0044 times the constant, below its bound in 6/6 cases inside the domain of the theorem. Scripts `suelo_frontera_parcial.py`, `alcanzable_explicito.py`.

Result 4 — dephasing and IBM's links (Section 7, Table 1)

What it says. Under Markovian dephasing the coherence of B with the vacuum is itself an amplitude of the same problem, with rates $2/T_2$, so Results 1–3 bound the channel fidelity F_B. For the four links of Ref. [4] at $g/2\pi = 9$ MHz, with the coherence times of each element: $1 - F_B \geq 6.03$–10.05×10^{-4} for any pulse (certified lower bounds, Markovian dephasing at the measured T_2^*; for any classical dephasing noise, Markovian or not, the relaxation-only floor 1.5–1.8×10^{-4} still holds); a 39-ns minimal-time pulse comes within a factor 1.4–1.5 of this floor; the measured errors are 71–114×10^{-4}, 11–16 times the floor. A floor computed with IBM's symmetric comparison model (worse times of the pair on both qubits) is not a floor for the device: on Q241–Q246 the 39-ns pulse beats it. By reciprocity (Proposition 7.3) the optimal error is the same in both directions. For design: doubling G halves the floor, the T_2^* of the worse qubit is the next lever, g could rise to 22.5 MHz with this cable (a scenario from other samples, with the coherence times kept fixed), and couplers that ramp in about 3 ns raise the floor by only about 10%; a 26-ns trapezoid at 18 MHz gives 5.6–9.3×10^{-4} on all four links (Section 8.1). Three other platforms sit 4.5–7.6 times above their floors (Section 8.2).

Why it holds. No quantum jump feeds the coherence with the vacuum (Lemma 7.1). With a different rate for each element, the floor comes from the dissipative version of Result 2 (Proposition 7.2), which keeps the emitter, the bus and the receiver apart.

Status. Proved on paper. The two window comparisons of Proposition 7.2 (excess rate $d > 0$) and a one-rate version of its norm chain are machine-checked in Isabelle/HOL (theory `Relay_Windows`; the receiver's barrier as a hypothesis); the bus term and the three-rate balance are not. The four IBM floors are certified with ball arithmetic (Arb). Reciprocity is proved on paper (sketch). Assumes one effective bus mode and Markovian dephasing.

Data. Coherence lemma checked against full Lindblad simulations to 4×10^{-16}; window integrals reproduced by an independent integrator to 0.1%; with equal qubits the bound reproduces the symmetric floors exactly; the floor lies below the 39-ns pulse in all 8 link-directions. Scripts `coherencia_efectiva.py`, `suelo_asimetrico.py`, `tabla_ibm_asimetrica.py`, `verificar_v203.py` (coherence times of each element), `replica_ventanas.py`, `certificar_suelo.py` (Arb), `verificar_v206.py` (reciprocity, 30 random pulses, to 10^{-15}).

The scripts and the Isabelle theories named in the boxes are in the supplementary package (`scripts/`, `isabelle/`).

3. Model and hypotheses

One excitation; emitter qubit A, receiver qubit B, bus C with $M \geq 1$ modes and an arbitrary Hermitian internal Hamiltonian $H_C(t)$ (detunings allowed). The qubits couple to the bus through vectors $g_A(t), g_B(t) \in \mathbb{C}^M$ with

$$\|g_A(t)\| \leq G, \qquad \|g_B(t)\| \leq G \qquad \text{for all } t, \tag{1}$$

each bounded separately (a "box"; Remark 9.1 compares a joint bound). The hypotheses are:

- H1 Pure loss at zero temperature: jumps take the excitation to the vacuum, which does not repopulate B. Qubits lose at rate γ, bus modes at rates collected in $K(t) \succeq \kappa \mathbb{1}$.
- H2 No direct A–B coupling; no drive on B.
- H3 (Section 7 only.) Markovian dephasing: Lindblad operators $\sqrt{1/T_{1,j}}\,|0\rangle\langle j|$ and $\sqrt{2\gamma_{\varphi,j}}\,|j\rangle\langle j|$ with $\gamma_{\varphi,j} = 1/T_{2,j} - 1/(2T_{1,j})$.
- H4 (Section 8 only.) One effective bus mode with $G = 2\pi \times 9$ MHz. G bounds the norm of the coupling vector, so a multimode bus with a per-mode bound of 9 MHz is a different constraint.

Under H1–H2 the population in B is given exactly by the no-jump evolution of the unnormalised amplitudes $a(t)$, $c(t) \in \mathbb{C}^M$, $b(t)$:

$$\dot a = -i\, g_A^\dagger c - \tfrac{\gamma}{2} a, \quad \dot b = -i\, g_B^\dagger c - \tfrac{\gamma}{2} b, \quad \dot c = -i H_C c - i(g_A a + g_B b) - \tfrac{1}{2} K c, \tag{2}$$

with $a(0) = 1$, $b(0) = 0$, $c(0) = 0$, and $P(T) = |b(T)|^2$. Controls are measurable and bounded, and $H_C(t)$ and $K(t)$ measurable and locally integrable in norm; the amplitudes are then absolutely continuous.

Rescaled frame. When $\kappa \geq \gamma$, put $\tilde a = e^{\gamma t/2} a$, $\tilde c = e^{\gamma t/2} c$, $\tilde b = e^{\gamma t/2} b$. The qubits become lossless and the bus keeps the loss $K - \gamma \succeq \delta \mathbb{1}$, $\delta := \kappa - \gamma \geq 0$. Write

$$x = |\tilde a|, \quad y = \|\tilde c\|, \quad z = |\tilde b|, \quad \tilde n = x^2 + y^2 + z^2, \quad \tilde I = \int_0^T y^2\, dt, \quad \tilde P = z(T)^2 = e^{\gamma T} P.$$

From (2) and (1), almost everywhere,

$$\dot{\tilde n} \leq -\delta\, y^2, \qquad |\dot x| \leq G\, y, \qquad |\dot z| \leq G\, y. \tag{3}$$

The lower bounds of Sections 4–6 use only (3) and $\tilde n(0) = 1$ (Section 7 uses the analogous inequalities (6) with a rate per element); this is why the number of modes, the phases and H_C never enter. The Isabelle proofs start from (3): the step from (2) to (3) is done on paper. The construction of Theorem 6.3, which shows that a loss is reachable, uses instead the exact equations of one mode.

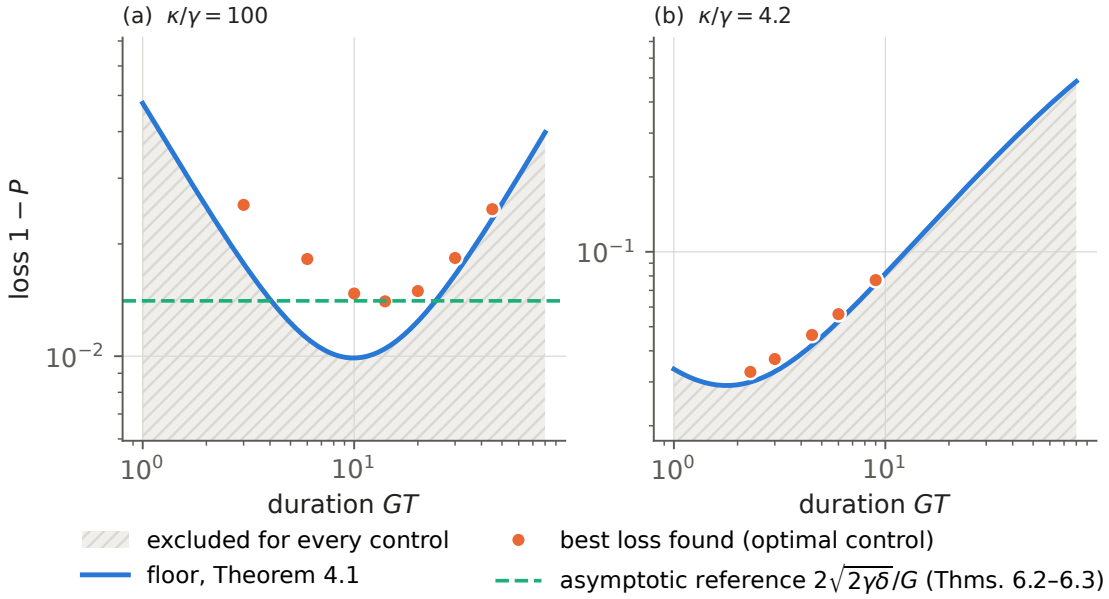


FIGURE 1. The floor of Theorem 4.1 against the duration, with $G = 1$. Hatched: losses that no control reaches. Orange: best losses found by optimal control (one mode, 40 piecewise-constant steps, 4 starts), not proven optima. (a) $\gamma = 5 \times 10^{-4}$, $\kappa = 0.05$: the best loss found is 44–63% above the floor at short durations and approaches it at long ones; near T_{asym} it reaches the asymptotic reference $2\sqrt{2\gamma\delta}/G$ of Theorems 6.2–6.3 (dashed), a factor $\sqrt{2}$ above the floor's minimum; the point at $GT = 14$ lies 0.3% below it, which the asymptotic statement allows. (b) $\gamma = 0.00826$, $\kappa = 0.0347$ ($\kappa/\gamma = 4.2$, the ratio of IBM's device): the floor is within 4–12% of the best loss found.

4. The floor

Theorem 4.1 (Loss floor)**.** *Under H1–H2 with $\kappa \geq \gamma$, for every control satisfying (1) and every $T > 0$,*

$$P(T) \leq \frac{e^{-\gamma T}}{1 + (\kappa - \gamma)/(G^2 T)} \tag{4}$$

Proof. By (3), $\tilde{P} \leq \tilde{n}(T) \leq 1 - \delta\tilde{I}$. Since $z(0) = 0$ and $|\dot{z}| \leq Gy$, $\sqrt{\tilde{P}} = z(T) \leq G\int_0^T y \leq G\sqrt{T\tilde{I}}$ by Cauchy–Schwarz, so $\tilde{I} \geq \tilde{P}/(G^2T)$. Hence $\tilde{P}\,(1 + \delta/(G^2T)) \leq 1$, and $P = e^{-\gamma T}\tilde{P}$. □

How close it is. Numerically (optimal control with piecewise-constant pulses, one mode) the best loss found lies above the floor $1 - e^{-\gamma T}/[1 + \delta/(G^2T)]$ given by (4) in all twelve cases of Figure 1. The ratio of best loss found to floor is 1.06–1.63 for $\kappa/\gamma = 100$ (largest at short durations) and 1.04–1.12 for $\kappa/\gamma = 4.2$, the ratio of IBM's device. In the regime of present devices the one-line bound is therefore within 4–12% of the best loss found; the factor $\sqrt{2}$ that separates it from the optimum when $\kappa \gg \gamma$ is the subject of Section 6.

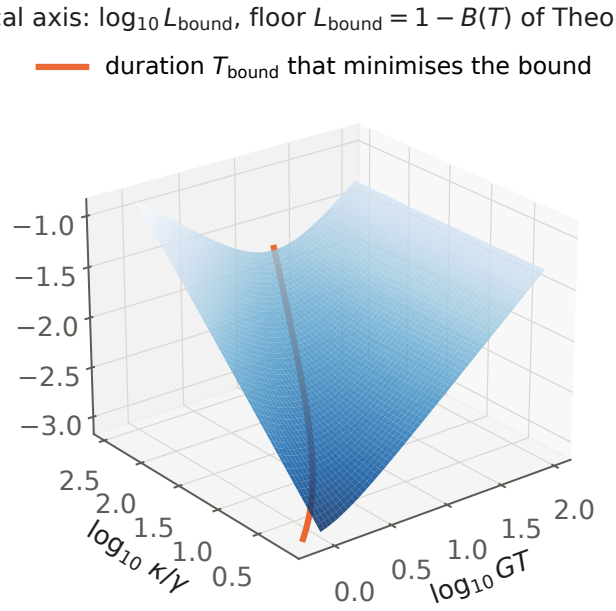


FIGURE 2. The floor $L_{\text{bound}} = 1 - B(T)$ of Theorem 4.1 as a function of the duration GT and of the loss ratio κ/γ, at $\gamma/G = 5 \times 10^{-4}$ (logarithmic axes). Short transfers pay the bus loss, long ones the qubit loss; the orange curve marks the exact duration T_{bound} that minimises the bound (Section 4).

Two durations. Write $B(T)$ for the right side of (4) and $a = \delta/G^2$. The duration that minimises the floor $L_{\text{bound}}(T) = 1 - B(T)$ solves $\gamma = a/(T(T+a))$:

$$T_{\text{bound}} = \frac{\sqrt{a^2 + 4a/\gamma} - a}{2} = \frac{\sqrt{\delta/\gamma}}{G}\Big(\sqrt{1 + \tfrac{\alpha^2}{4}} - \tfrac{\alpha}{2}\Big), \qquad \alpha = \frac{\sqrt{\gamma\delta}}{G}.$$

It is not the duration of the best pulse: in the asymptotic regime of Section 6 the best pulses last $T_{\text{asym}} \sim \sqrt{2\delta/\gamma}/G$, a factor $\sqrt{2}$ longer. Minimising a bound gives a floor, not a pulse. For a quick estimate by hand, $e^{-u} \le 1/(1+u)$ and the inequality of arithmetic and geometric means turn (4) into a bound valid at every T:

$$1 - P \ge 1 - \Big(1 + \frac{\sqrt{\gamma\delta}}{G}\Big)^{-2} \approx \frac{2\sqrt{\gamma\delta}}{G}.$$

Remark 4.2 (Ramped couplings). If in addition $\|\dot{g}_i(t)\| \le S$ (Euclidean norm of the coupling vector) and $g_i(0) = g_i(T) = 0$, then $\|g_i(t)\| \le m(t) := \min(St, S(T-t), G)$, and Cauchy–Schwarz with weight m gives (4) with G^2T replaced by $J(T) = \int_0^T m^2$ ($S^2T^3/12$ for $T \le 2G/S$, $G^2T - 4G^3/(3S)$ for $T \ge 2G/S$). With $A = \pi/\sqrt{2}$, reparametrising time with $s = \int_0^t m$ gives the minimal time $2\sqrt{A/S}$ when $S \le G^2/A$ (triangular envelope), $A/G + G/S$ otherwise (trapezoidal), attained without losses by $g_A = g_B = m$. (Proved on paper; not formalised.)

5. The end qubits must be occupied

Let $\Phi(u) = \frac{1}{2}(\arcsin u - u\sqrt{1-u^2})$ on $[0,1]$. Then $\Phi'(u) = u^2/\sqrt{1-u^2} \ge 0$, $\Phi(0) = 0$ and $\Phi(1) = \pi/4$.

Lemma 5.1. *Let w be absolutely continuous and $y \geq 0$ continuous on $[0,T]$, with $w \geq 0$, $w^2+y^2 \leq 1$, and $|\dot{w}| \leq G\,y$ outside a null set. Then $|\Phi(w(T)) - \Phi(w(0))| \leq G \int_0^T w^2\,dt$.*

Proof. Where $w < 1$, $|\frac{d}{dt}\Phi(w)| = w^2|\dot{w}|/\sqrt{1-w^2} \leq G\,w^2\,y/\sqrt{1-w^2} \leq G\,w^2$, because $y \leq \sqrt{1-w^2}$. At $w = 1$ the derivative of Φ does not exist, so apply the estimate to $\Phi(\lambda w)$ with $0 < \lambda < 1$: Φ is Lipschitz on $[0,\lambda]$, hence $\Phi(\lambda w)$ is absolutely continuous, the fundamental theorem of calculus for absolutely continuous functions applies, and almost everywhere

$$\left|\tfrac{d}{dt}\Phi(\lambda w)\right| = \frac{\lambda^3 w^2|\dot{w}|}{\sqrt{1-\lambda^2 w^2}} \leq G\,w^2\,\frac{\lambda^3\sqrt{1-w^2}}{\sqrt{1-\lambda^2 w^2}} \leq G\,w^2.$$

Let $\lambda \to 1$; Φ is continuous at the end points. □

Theorem 5.2 (Endpoint occupation)**.** *Let x, z be absolutely continuous and y continuous on $[0,T]$, with $x,y,z \geq 0$, $x^2+y^2+z^2 \leq 1$ and $|\dot{x}|, |\dot{z}| \leq G\,y$ outside a null set. Then*

$$G\int_0^T (x^2+z^2)\,dt \;\geq\; [\Phi(x(0)) - \Phi(x(T))] + [\Phi(z(T)) - \Phi(z(0))].$$

In particular, for a complete transfer ($x : 1 \to 0$, $z : 0 \to 1$) and every T,

$$\boxed{G\int_0^T (x^2+z^2)\,dt \geq \frac{\pi}{2}.}$$

Proof. Apply Lemma 5.1 to (x,y) and to (z,y); $x^2+y^2 \leq 1$ and $z^2+y^2 \leq 1$ follow from the hypothesis. □

By (3) the theorem applies to the relay in the rescaled frame, with any number of modes, any phases and any bus loss. With bus loss ($\delta > 0$) a complete transfer is impossible in that frame, since $\tilde{P} \leq \tilde{n}(T) \leq 1 - \delta\tilde{I}$; what the theorem gives is the partial form: a transfer that empties A and delivers $\tilde{P} = z(T)^2$ has $G\int_0^T (x^2+z^2) \geq \pi/4 + \Phi(\sqrt{\tilde{P}})$, which tends to $\pi/2$ as $\tilde{P} \to 1$. Without bus excess loss ($\delta = 0$) the bound $\pi/2$ is attained: swap $A \to C$ with $g_A = G$ for a time $\pi/(2G)$, wait, swap $C \to B$; each swap contributes exactly $\pi/(4G)$, for every $T \geq \pi/G$. Random controls (300 instances, 1–3 modes, complex phases, losses; these transfers are partial) satisfy the general inequality with a ratio of left to right side of at least 2.11 (Figure 3).

Scope of the formal proof. For measurable bounded controls the moduli $x = |\tilde{a}|$ and $z = |\tilde{b}|$ are absolutely continuous, with $|\dot{x}| \leq |\dot{\tilde{a}}| \leq Gy$ almost everywhere; the exceptional set is null but need not be countable. The Isabelle theory `Relay_Endpoints` covers exactly this case (absolutely continuous functions, a negligible exceptional set), through the fundamental theorem of calculus for absolutely continuous functions of Isabelle/HOL-Analysis: Lemma 5.1 is `endpoint_window_ac`, and the complete-transfer form of Theorem 5.2 is `endpoint_occupation_ac`; the general inequality is the lemma applied twice. Earlier versions (`endpoint_window`, `endpoint_occupation`), for exceptions in a countable set, cover piecewise-smooth controls.

6. Partial transfer and the asymptotic constant

The step $\tilde{I} \geq \tilde{P}/(G^2T)$ in Theorem 4.1 loses a factor 2 when GT is large: the excitation must cross from A to B, and the crossing forces more bus occupation.

Lemma 6.1 (Frontier for partial transfer)**.** *Under (3) with $\tilde{n}(0) = 1$, $x(0) = 1$, $y(0) = z(0) = 0$, let $\tilde{P} = 1 - \varepsilon$ with $\varepsilon \leq 1/2$. There is $y_c \in [0,1]$ with*

$$G^2T\,\tilde{I} \geq \left(\sqrt{2(1-\varepsilon-y_c^2)_+} - \sqrt{\varepsilon}\right)_+^2 \quad \text{and} \quad G\,\tilde{I} \geq \tfrac{\sqrt{2}}{6}\,y_c^3.$$

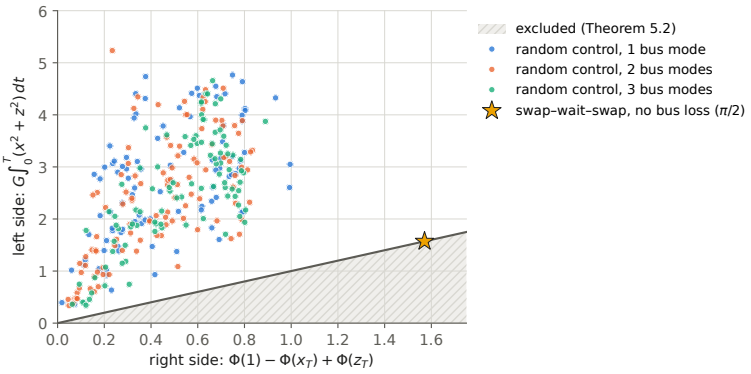


FIGURE 3. Theorem 5.2 on 300 random controls (piecewise-constant complex couplings with $\|g\| \le G$, random bus Hamiltonian and losses, 1–3 bus modes; qubits lossless). Each point is one control: horizontal, the right side $\Phi(1) - \Phi(x_T) + \Phi(z_T)$; vertical, the occupation $G\int_0^T (x^2+z^2)$. The theorem excludes the hatched region. The star is swap–wait–swap without bus loss, a complete transfer that lies on the boundary at $\pi/2$.

Consequently $G^2 T\,\tilde{I} \ge \Psi(GT, \varepsilon)$, *where*

$$\Psi(\tau,\varepsilon) := \min_{y\in[0,1]} \max\Big(1-\varepsilon,\ \big(\sqrt{2(1-\varepsilon-y^2)_+} - \sqrt{\varepsilon}\big)_+^2,\ \tau\tfrac{\sqrt{2}}{6}y^3\Big),$$

and $\Psi(\tau,\varepsilon) \to 2$ *as* $\tau\to\infty$ *and* $\varepsilon \to 0$.

Proof. $x(T)^2 \le \tilde{n}(T) - z(T)^2 \le \varepsilon$, so $x(T) \le \sqrt{\varepsilon} \le \sqrt{1-\varepsilon} = z(T)$; hence $x - z$ goes from 1 to ≤ 0 and vanishes at some t_c, where $x = z = s$ and $y = y_c$. There $2s^2 + y_c^2 = \tilde{n}(t_c) \ge \tilde{n}(T) \ge 1-\varepsilon$. Then $G\int_0^T y \ge \int_0^{t_c} |\dot{z}| + \int_{t_c}^T |\dot{x}| \ge s + (s - x(T)) \ge \sqrt{2(1-\varepsilon-y_c^2)_+} - \sqrt{\varepsilon}$, and Cauchy–Schwarz gives the first inequality. For the second, the coupling changes y^2 at rate $-2(x\dot{x} + z\dot{z})$, of modulus at most $2G\,y\,(x+z) \le 2\sqrt{2}\,G\,y$, and the bus loss only decreases it; so $\dot{y} \le \sqrt{2}\,G$ where $y > 0$. As $y(0) = 0$, $t_c \ge y_c/(\sqrt{2}G)$ and $y(t) \ge y_c - \sqrt{2}G(t_c - t)$ on $[t_c - y_c/(\sqrt{2}G), t_c]$, whence $\tilde{I} \ge y_c^3/(3\sqrt{2}\,G)$. The term $1-\varepsilon$ is the step $\tilde{I} \ge \tilde{P}/(G^2T)$ of Theorem 4.1. Finally, if $\Psi < 2-\eta$ then $y \le (6(2-\eta)/(\sqrt{2}\tau))^{1/3} \to 0$ and the middle term tends to 2, a contradiction for large τ and small ε. □

Theorem 6.2 (Asymptotic floor)**.** *Let* $L^* = \inf_T \inf_{\text{controls}}(1 - P(T))$ *and* $\ell_0 = 2\sqrt{2\gamma\delta}/G$. *If* $\gamma/\delta \to 0$ *and* $\sqrt{\gamma\delta}/G \to 0$, *then*

$$\boxed{L^* \ge \frac{2\sqrt{2\gamma\delta}}{G}\,(1 - o(1)).}$$

Proof. Fix a control and T, and let $L = 1 - P$; assume $L \le \ell_0$ (otherwise there is nothing to prove). Since $\tilde{P} = e^{\gamma T}P \ge P$, $\varepsilon = 1 - \tilde{P} \le L \le \ell_0 \to 0$. By Theorem 4.1, $L \ge \delta/(G^2T + \delta)$, so $G^2T \ge \delta(1/\ell_0 - 1)$ and $GT \ge \sqrt{\delta/\gamma}/(2\sqrt{2})\,(1 - o(1)) \to \infty$. Lemma 6.1 gives $\tilde{I} \ge (2-\eta)/(G^2T)$ with $\eta \to 0$ uniformly. From $P = e^{-\gamma T}\tilde{P} \le e^{-\gamma T}(1 - \delta\tilde{I}) \le e^{-s}$ with $s = \gamma T + \delta\tilde{I}$,

$$L \ge 1 - e^{-s}, \qquad s \ge \gamma T + \frac{(2-\eta)\delta}{G^2T} \ge \frac{2\sqrt{(2-\eta)\gamma\delta}}{G} = \ell_0(1 - o(1)),$$

and $1 - e^{-s} \ge s(1 - s/2)$, where $s \to 0$ because $1 - e^{-s} \le L \le \ell_0$ gives $s \le -\log(1 - \ell_0)$. □

The constant is attained, with explicit error terms.

Theorem 6.3 (Attaining the constant)**.** *Let $G > 0$, $\gamma > 0$, $\delta > 0$, $\alpha = \sqrt{\gamma\delta}/G$ and $\beta = \gamma/\delta$, with $\alpha \le 0.01$ and $\beta \le 0.01$. Suppose the bus has a mode, an eigenmode of H_C that the couplers can address alone, whose loss rate is exactly $\kappa = \gamma + \delta$ (resonant, or with its detuning compensated by the phase of the couplings). There is a control acting on that mode, with $|g_A|, |g_B| \le G$, such that*

$$1 - P \le \frac{2\sqrt{2\gamma\delta}}{G}\Big(\frac{1}{1-\alpha} + 1.5\,\beta\Big).$$

Together with Theorem 6.2, $L^* = \ell_0\,(1 + o(1))$ when $\alpha, \beta \to 0$, for this realisable one-mode problem. (For a bus whose loss is only bounded below by κ, Theorem 6.2 still holds, but the upper bound need not.)

Proof. Work in the rescaled frame with real amplitudes $\tilde a = x$, $\tilde c = -iy$, $\tilde b = -z$, so that $\dot x = -g_A y$, $\dot z = g_B y$, $\dot y = g_A x - g_B z - \frac{\delta}{2} y$. Put $y_0 = \sqrt{\gamma/\delta}$, so that $\delta y_0 = \alpha G$ and $y_0^2 = \beta$. The control has three stages. (i) Ramp: $g_A = G$, $g_B = 0$ until $y = y_0$. (ii) Arc: writing $x = r\cos\theta$, $z = r\sin\theta$, take $g_A = h\sin\theta + d\cos\theta$, $g_B = h\cos\theta - d\sin\theta$ with $d = \delta y_0/(2r)$ and $h = (G - d)/\max(\sin\theta, \cos\theta)$, until $\theta = \pi/2$. Then $g_A x - g_B z = rd = \delta y_0/2$, so $\dot y = 0$ and $y \equiv y_0$; $r\dot\theta = y_0 h$; and $|g_A|, |g_B| \le G$ because $h\max(\sin\theta, \cos\theta) + d \le G$. (iii) Unload: $g_A = 0$, $g_B = G$ until $y = 0$.

During all three stages $y \le y_0$, so $\delta\tilde I \le \delta y_0^2 T = \gamma T$, and therefore $1 - P = 1 - e^{-\gamma T}\tilde P \le \gamma T + (1 - \tilde P) \le \gamma T + \delta\tilde I \le 2\gamma T$ (at the end $x = y = 0$, so $\tilde P = \tilde n(T) \ge 1 - \delta\tilde I$). It remains to bound T. Let $t_* = 0.02/\gamma$ and run the protocol until it ends or until t_*, whichever comes first. On that interval $\gamma t \le 0.02$, and the exact norm balance of the one-mode problem, $\tilde n(t) = 1 - \delta\int_0^t y^2 \ge 1 - \gamma t$, gives $x^2 + z^2 \ge 1 - \beta - \gamma t \ge 0.97$ in stages (i)–(ii) and $z \ge 0.97$ in stage (iii); so r stays away from zero and the control is well defined. In (i), $\dot y \ge Gx - \frac{\delta}{2} y_0 \ge G(0.97 - \alpha/2) \ge 0.96\,G$, so this stage lasts at most $y_0/(0.96G)$; in (iii), $\dot y \le -Gz \le -0.97G$, at most $y_0/(0.97G)$. In (ii), $r \le 1$ and $d \le \delta y_0/1.9 \le \alpha G$, so $\dot\theta \ge y_0 G(1 - \alpha)/\max(\sin\theta, \cos\theta)$ and the arc lasts at most $\int_0^{\pi/2} \max(\sin\theta, \cos\theta)\,d\theta/(y_0 G(1-\alpha)) = \sqrt{2}/(y_0 G(1-\alpha))$. These speed bounds limit the time spent in each stage even before it ends. Hence the protocol ends by $\overline{T} = \sqrt{2}/(y_0 G(1-\alpha)) + 2.1\,y_0/G$, and $\gamma\overline{T} \le \sqrt{2}\,\alpha/(1-\alpha) + 2.1\alpha\beta \le 0.0145 < 0.02$; if it had not ended by t_* we would have $t_* \le \overline{T} < t_*$, a contradiction. So $T \le \overline{T}$ and

$$2\gamma T \le \frac{2\sqrt{2}\,\sqrt{\gamma\delta}}{G(1-\alpha)} + \frac{4.2\,\gamma y_0}{G} = \ell_0\Big(\frac{1}{1-\alpha} + \frac{4.2}{2\sqrt{2}}\,\beta\Big),$$

and $4.2/(2\sqrt{2}) < 1.5$. □

Figure 4 shows the pulse. Three practical remarks: g_B becomes slightly negative at the end of the arc ($g_B = -d$, of order αG), the couplings switch instantaneously between stages (so the construction does not cover the slope limit of Remark 4.2), and the control depends on the state only through the known trajectory, so it can be computed in advance and applied in open loop. Integrating it in the original (not rescaled) lossy dynamics, with $G = 1$ and $(\gamma, \kappa) = (10^{-4}, 0.0101)$, $(10^{-5}, 0.0101)$, $(10^{-6}, 0.002)$, $(10^{-7}, 0.001)$, $(5\times10^{-6}, 0.000505)$ and $(10^{-5}, 0.00101)$, all inside the domain of the theorem, gives $(1 - P)/\ell_0 = 1.0035$, 1.0002, 1.0002, 1.00004, 1.0044 and 1.0043, all below the bound, with $\max|g| = G$.

The lower bound of Theorem 6.2 converges slowly: $\Psi(\tau, 0) = 1.66$, 1.92, 1.98, 1.996 at $\tau = 10^2, 10^3, 10^4, 10^5$, and $\Psi(\tau, 0)$ exceeds 1, the value used in Theorem 4.1, only for $\tau > 12$ (at $y^2 = 1/2$ the third term is $\tau/12$), and in our numerical comparisons the improvement becomes noticeable from $GT \approx 20$. For present devices ($GT \approx 2$–3 at the optimum) Theorem 4.1 is the bound to use; Theorems 6.2 and 6.3 fix the constant.

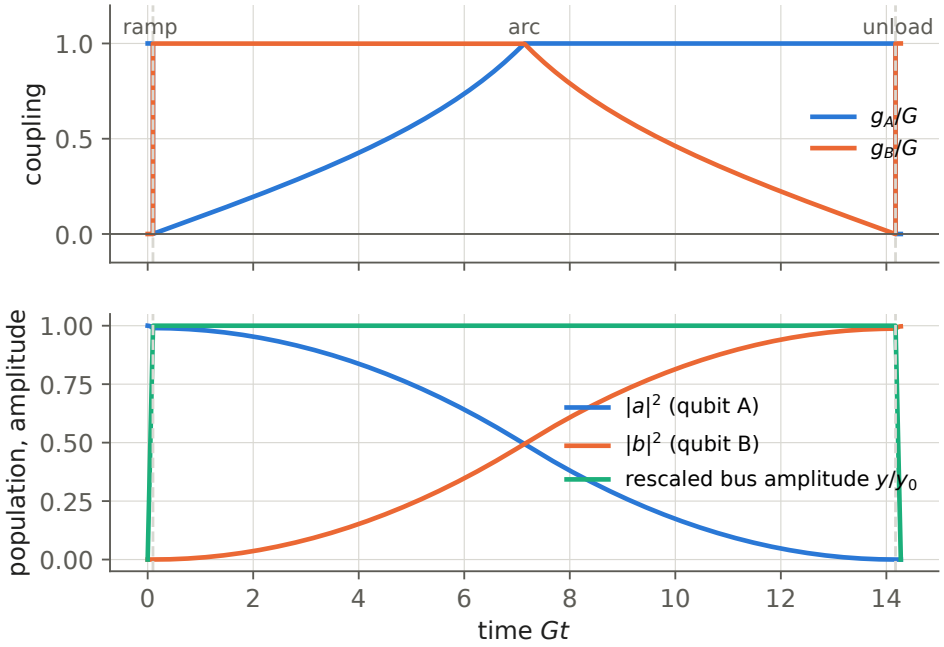


FIGURE 4. The pulse of Theorem 6.3 for $G = 1$, $\gamma = 10^{-4}$, $\kappa = 0.0101$ ($\alpha = 0.001$, $\beta = 0.01$), integrated in the original lossy dynamics. Top: the couplings in the three stages (ramp, arc, unload; dashed lines). Bottom: the populations of the end qubits and the rescaled bus amplitude y/y_0, which stays at 1 during the arc; the physical amplitude is $e^{-\gamma t/2} y_0$. The final $1 - P$ is $1.0035\,\ell_0$.

7. Dephasing

Lemma 7.1 (Coherence is an amplitude)**.** *Under H1–H3, let ρ evolve from $(|0\rangle + |A\rangle)/\sqrt{2}$ and set $v_j = 2\rho_{j0}$ for j in the single-excitation sector. Then v obeys the no-jump dynamics (2) with the loss rates replaced by $\Gamma_j = 1/T_{1,j} + 2\gamma_{\varphi,j} = 2/T_{2,j}$, and $v(0) = (1, 0, \dots, 0)$.*

Proof. For $j \neq 0$, every Lindblad operator L of H3 satisfies $\langle j|L = 0$ or $L^\dagger|0\rangle = 0$, so no jump term $L\rho L^\dagger$ feeds ρ_{j0}; the anticommutator contributes $-\frac{1}{2}(L^\dagger L)_{jj}\rho_{j0}$, since $(L^\dagger L)_{00} = 0$. The Hamiltonian conserves the excitation number. □

Let P and n be the population of B and the excitation remaining at T for the input $|A\rangle$, and $c = |v_B(T)|$ the coherence for the input $(|0\rangle + |A\rangle)/\sqrt{2}$ of Lemma 7.1. (For the latter input the vacuum component does not evolve, so its populations are $P/2$ and $n/2$.) With the phase of B corrected by a known local rotation, the six-state average fidelity of the channel from A to B (the global fidelity, which also requires A and the bus to end empty) is $F = (3 + 2P + 2c - n)/6$. The reduced channel at B, which is what network benchmarking measures, has

$$F_B = \frac{3 + P + 2c}{6}, \qquad 1 - F_B = \frac{1 - P}{6} + \frac{1 - c}{3} \;\leq\; 1 - F. \tag{5}$$

Both formulas agree with full Lindblad simulations to 4×10^{-16} in six random controls; with rates $1/T_1$ in place of $2/T_2$ the coherence is off by 7×10^{-3}, as it should be. Hence every bound on the amplitude c is a bound on $1 - F_B$: Theorems 4.1–6.2 apply to c (the bounds on P become bounds

on c^2) with $\gamma \to 2/T_{2q}$ and $\kappa \to 2/T_{2c}$. If the bus dephases more ($T_{2c} < T_{2q}$),

$$c^2 \le \frac{e^{-2T/T_{2q}}}{1 + (2/T_{2c} - 2/T_{2q})/(G^2 T)}.$$

The same theorems also bound P itself in the presence of dephasing, by an argument that does not use the no-jump evolution: in the rescaled single-excitation block $\tilde\rho$, put $x = \sqrt{\tilde\rho_{AA}}$, $y = \sqrt{\operatorname{Tr}\tilde\rho_{CC}}$, $z = \sqrt{\tilde\rho_{BB}}$. The coupling changes $\tilde\rho_{AA}$ at a rate $\pm 2\operatorname{Im}(g_A^\dagger \tilde\rho_{CA})$ (the sign depends on the convention), of modulus at most $2\|g_A\|\,\|\tilde\rho_{CA}\|$; positivity gives $\|\tilde\rho_{CA}\| \le \sqrt{\tilde\rho_{AA}\operatorname{Tr}\tilde\rho_{CC}}$, so $|\dot x| \le Gy$ (and likewise for z); dephasing that preserves the populations of A, B and the bus does not enter the norm balance, and (3) holds.

Real devices do not have equal qubits, and which element dephases fastest varies: in IBM's links the worse qubit dephases faster than the bus, but two qubits (Q242 and Q246) dephase more slowly than their bus modes (Section 8). We therefore keep a rate for each element. Let r_A, r_C, r_B be the decay rates of the amplitude problem ($2/T_2$ of each element for c, $1/T_1$ for P), $r_0 = \min(r_A, r_C, r_B)$ and $d_j = r_j - r_0 \ge 0$. In the frame rescaled by $e^{r_0 t/2}$ the elements keep the excess rates d_A, d_C, d_B, and with x, y, z the rescaled moduli,

$$\dot{\tilde n} \le -d_A x^2 - d_C y^2 - d_B z^2, \qquad \dot x \ge -Gy - \tfrac{d_A}{2}x, \qquad \dot z \le Gy - \tfrac{d_B}{2}z. \tag{6}$$

Proposition 7.2 (Windows with a rate per element)**.** *Let $\zeta = z(T)$. Let f_A solve $f' = -G\sqrt{1-f^2} - \frac{d_A}{2}f$, $f(0) = 1$, and w_ζ solve $w' = -G\sqrt{1-w^2} + \frac{d_B}{2}w$, $w(0) = \zeta$, both taken as the solutions that leave their initial value immediately and stopped at 0 ($w_0 \equiv 0$); equivalently $F_A(f_A(t)) = \min(t, \tau_A)$ and $K_\zeta(w_\zeta(s)) = \min(s, \tau_B(\zeta))$, with $\tau_A = F_A(0)$, $\tau_B(\zeta) = K_\zeta(0)$ and $F_A(u) = \int_u^1 dv/(G\sqrt{1-v^2} + \frac{d_A}{2}v)$ and $K_\zeta(u) = \int_u^\zeta dv/(G\sqrt{1-v^2} - \frac{d_B}{2}v)$. Let M_B solve $M' = G\sqrt{1-M^2} - \frac{d_B}{2}M$, $M(0) = 0$, and $J_B(T) = (1 - e^{-d_B T})/d_B$ ($J_B = T$ if $d_B = 0$). Put $a_A(T) = \int_0^T f_A^2$ and $b_B(T,\zeta) = \int_0^T w_\zeta^2$. Then*

$$\boxed{\zeta \le M_B(T), \qquad \zeta^2 + f_A(T)^2 + d_A\, a_A(T) + d_B\, b_B(T,\zeta) + \frac{d_C\,\zeta^2}{G^2 J_B(T)} \le 1,}$$

and the physical amplitude at B is at most $e^{-r_0 T/2}\zeta$.

Proof. Since $y^2 \le \tilde n - x^2 \le 1 - x^2$, (6) gives $\dot x \ge -G\sqrt{1-x^2} - \frac{d_A}{2}x$, so $\frac{d}{dt}F_A(x(t)) \le 1$, $F_A(x(t)) \le t$ and $x \ge f_A$; hence $\int_0^T x^2 \ge a_A(T)$. In the same way $\dot z \le G\sqrt{1-z^2} - \frac{d_B}{2}z$ gives, running time backwards from T, $K_\zeta(z(t)) \le T - t$, so $z(t) \ge w_\zeta(T-t)$ and $\int_0^T z^2 \ge b_B(T,\zeta)$; and forwards from $z(0) = 0$ it gives $z(t) \le M_B(t) \le M_B(T)$. If $d_B > 0$, M_B stays below the zero $v^* = G/\sqrt{G^2 + d_B^2/4} < 1$ of $G\sqrt{1-v^2} - \frac{d_B}{2}v$, so the denominator of K_ζ stays positive on $[0,\zeta]$; if $d_B = 0$, $K_\zeta(u) = (\arcsin\zeta - \arcsin u)/G$ and $w_\zeta(s) = \sin([\arcsin\zeta - Gs]_+)$ explicitly. If $d_A = 0$ the same holds with $F_A(u) = \arccos(u)/G$, the singular start $x(0) = 1$ being handled by the regularisation of Lemma 5.1. Next, $\frac{d}{dt}(e^{d_B t/2}z) \le G\, e^{d_B t/2}y$, so $\zeta \le G\int_0^T e^{-d_B(T-t)/2}y \le G\sqrt{J_B(T)\int_0^T y^2}$ by Cauchy–Schwarz, i.e. $\int_0^T y^2 \ge \zeta^2/(G^2 J_B(T))$. Finally, with $x(T) \ge f_A(T)$, $f_A(T)^2 + \zeta^2 \le x(T)^2 + \zeta^2 \le \tilde n(T) \le 1 - d_A\int x^2 - d_B\int z^2 - d_C\int y^2$. □

The term $f_A(T)^2$ comes from keeping $x(T)^2 \ge f_A(T)^2$ in the final balance $x(T)^2 + \zeta^2 \le \tilde n(T)$; it matters when the emitter is the element with the lowest rate ($d_A = 0$, $f_A(T) = \cos(\min(GT, \pi/2))$), where a_A drops out and the bound would otherwise not depend on the emitter's coupling. All floors of Section 8 use this form, with Proposition 7.3; without $f_A(T)^2$ and reciprocity they would be 1.3–3.7% lower. Each window uses only the coupling of its own end, so with different maximal couplings G_A and G_B the proposition holds with G_A in the emitter window and f_A, and G_B in the receiver window, M_B and J_B.

Proposition 7.3 (Reciprocity)**.** *Under H1–H3, with local loss and diagonal local dephasing at rates that do not depend on time, let a pulse $H(t)$ on $[0,T]$ transfer from A to B, and let $H_{\mathrm{rev}}(t) = H(T-t)^{\mathsf{T}}$ (time-reversed and transposed, the bus Hamiltonian H_C included; for real couplings and real H_C, just played backwards). Then $P_{B\leftarrow A}[H] = P_{A\leftarrow B}[H_{\mathrm{rev}}]$ and $c_{B\leftarrow A}[H] = c_{A\leftarrow B}[H_{\mathrm{rev}}]$, so F_B is the same. If each coupler keeps its own bound and the reversed and transposed pulse is realisable, H_{rev} is admissible whenever H is, and the optimal error is the same in both directions.*

Proof sketch. In the single-excitation block the generator is $-i(H - \frac{i}{2}\Gamma)$ with Γ diagonal, and the no-jump propagator of H_{rev} is the transpose of that of H, so $\langle A|U_{\mathrm{rev}}|B\rangle = \langle B|U|A\rangle$; this gives c by Lemma 7.1, and P under pure loss. With dephasing, the dissipative part of the Lindblad generator acts diagonally on the operator basis $|j\rangle\langle k|$, and the same transposition argument applies to the populations. Numerically, 30 random pulses with random relaxation and dephasing give equal P and c in both directions to 9×10^{-16}. □

By reciprocity a floor may take, at each T, the smaller of the upper bounds of the two directions before minimising over T; the direction of transfer is not a design choice in this model.

Proposition 7.2 contains the earlier cases. If only the bus has an excess ($d_A = d_B = 0$), it is the mechanism of Theorem 4.1 ($J_B = T$). If only the qubits have one and it is the same for both ($d_A = d_B = \Delta$, $d_C = 0$), it is the symmetric bound, which we evaluated earlier with the coherence model that Ref. [4] uses for comparison. As $d_A, d_B \to 0$ with enough time, the emitter window tends to $\pi/(4G)$ and the receiver window to $\Phi(\zeta)/G$, recovering Theorem 5.2 in its partial form; setting all $d_j = 0$ in the inequality alone would only give $\zeta^2 + f_A(T)^2 \le 1$. Linear windows $x \ge 1 - Gt$ would give $1/(3G)$ instead of $\pi/(4G)$, a factor $3\pi/4$ weaker. The proposition is proved on paper. Its two window comparisons, $F_A(x(t)) \le t$ and $K_\zeta(z(t)) \le T - t$ together with $\zeta \le M_B(T)$ (for $d_A, d_B > 0$; with $d = 0$ they are explicit), and a version of the final norm chain with one excess rate common to both qubits, are machine-checked in Isabelle/HOL (theory `Relay_Windows`), for derivatives outside a countable set and with the barrier $z \le c < v^*$ of the receiver as a hypothesis (v^* the zero of $G\sqrt{1-v^2} - \frac{d_B}{2}v$); the bus term with J_B and the general balance with three excess rates are not formalised. It is evaluated numerically (integration in the angle $f = \cos\theta$ with adaptive step, cross-checked with an independent DOP853 integrator to 0.1%); with $d_A = d_B$ and $d_C = 0$ it reproduces our earlier symmetric floors to 0.01×10^{-4}.

8. Application to IBM's relay

We use the coherence times of each element as reported in Ref. [4]: T_1 and T_2^* of each qubit (Extended Data Table 1) and of the cable mode of each link (Extended Data Table 2), with $G = 2\pi \times 9$ MHz (H4). IBM's own coherence model assigns the shorter T_1 and the shorter T_2 of the pair to both qubits; that is a comparison model, not the device, and a floor computed from it is not a floor for the device (with the individual times, a 39-ns pulse on Q241–Q246 reaches 11.35×10^{-4}, below the 12.77×10^{-4} that the symmetric model gives). The floor on $1 - F_B$ is

$$\min_T \Big\{ \frac{1 - c_{\max}(T)}{3} + \frac{1 - P_{\max}(T)}{6} \Big\},$$

with $c_{\max} = e^{-r_0 T/2}\zeta$ from Proposition 7.2 at rates $2/T_2$ and $P_{\max} = e^{-r_0 T}\zeta^2$ from the same proposition at rates $1/T_1$ (each with its own r_0 and ζ), where, by reciprocity (Proposition 7.3), $c_{\max}(T)$ and $P_{\max}(T)$ are the smaller of the bounds of the two directions at each T. The numerical evaluation integrates the windows in the angle θ ($f = \cos\theta$), where every window, reach and time has a closed-form primitive (in u, $\sin 2u$, $\sin^2 u$ and $\ln\sin u$); they agree with 30-digit quadrature (mpmath) to 10^{-13}, and the minimum over T is searched over a range that widens until the minimum is interior, with both tails closed by rigorous bounds. Besides evaluating this floor numerically, we certified it with ball arithmetic (Arb, in SageMath): in the angle θ of $f = \cos\theta$ every window, reach

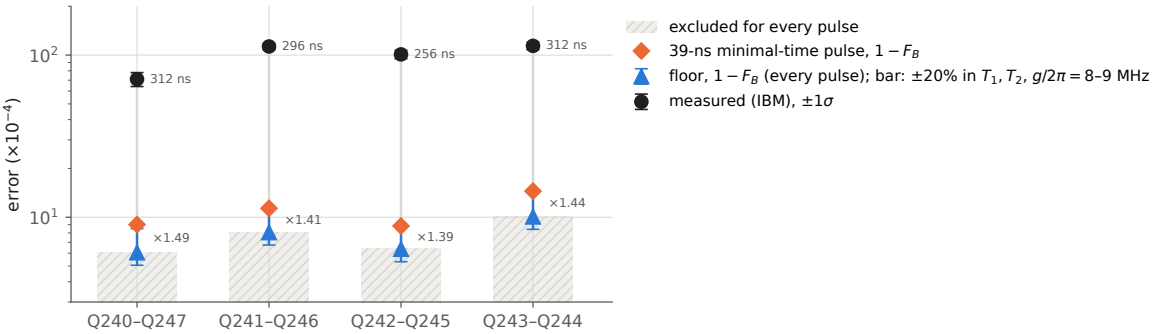


FIGURE 5. The four links of Table 1 on a logarithmic scale, with the coherence times of each element. Hatched: errors below the floor, which no pulse reaches under H1–H4. The bar on each floor is its range when all coherence times change by ±20% and $g/2\pi$ is 8 or 9 MHz. The factor next to each link is minimal-time error over floor; the duration next to the measured point is IBM's pulse length, and its bar the quoted uncertainty.

and time is an integral of an analytic function, which Arb encloses rigorously; each inversion (the angle reached at a given time, the largest admissible ζ) is found in floating point and then certified; and the minimum over T is bounded on 200 intervals $[T_a, T_b]$ using the monotonicity of each piece, plus the two tails $T \to 0$ and $T \to \infty$. The certified values lie 0.3–0.6% below the numerical ones (Table 1). The laboratory Check your link of the companion site computes this floor for any link, with the same code. For each link, the numerical floor, an independent implementation (another integrator, within 0.2%), an exact control (with excess loss only in the bus and $T \geq \pi/(2G)$ the bound reduces to Theorem 4.1) and the certificate are run together as one reproducible pipeline (`scripts/misiones/`).

TABLE 1. The four links (each name opens the laboratory with its values), in units of 10^{-4}, all in the metric $1 - F_B$. *Floor*: the bound for every pulse under H1–H4, with the coherence times of each element (Proposition 7.2 with $f_A(T)^2$, and reciprocity), evaluated numerically. *Certified*: a rigorous lower bound on the same floor, computed with ball arithmetic (Arb) as described in the text. *Range*: the floor when all coherence times change by ±20% together and $g/2\pi$ is 8 or 9 MHz. *Relaxation only*: the same floor with $T_2 = 2T_1$ for every element (no pure dephasing). *Minimal time*: $1 - F_B$ of the pulse $g_A = g_B = G$ for $\pi/(\sqrt{2}G) \approx 39$ ns, in a full Lindblad simulation with the same coherence times. *Ratio*: minimal time over floor. *Measured*: the mean transfer error of both directions reported in Ref. [4] (Extended Data Table 2, with its quoted uncertainty), at their duration T_g.

link	floor	certified	range	relaxation only	minimal time	ratio	measured (T_g)
Q240–Q247	6.06	6.03	5.06–8.51	1.71	9.02	1.49	71 ± 7 (312 ns)
Q241–Q246	8.07	8.05	6.74–11.31	1.51	11.35	1.41	113 ± 5 (296 ns)
Q242–Q245	6.37	6.34	5.31–8.95	1.84	8.83	1.39	101 ± 6 (256 ns)
Q243–Q244	10.09	10.05	8.42–14.11	1.83	14.49	1.44	114 ± 6 (312 ns)

Three readings. (a) Relaxation alone sets a floor of only 1.5–1.8 × 10^{-4} on $1 - F_B$. This floor does not depend on how the qubits dephase: it uses only the loss rates $1/T_1$, the flux inequalities hold for every realisation of a classical dephasing noise (random detunings of each

element, Markovian or not, slow or static), and the coherence obeys $c^2 \le P$ by Jensen's inequality, so $1 - F_B \ge (1 - P_U)/6 + (1 - \sqrt{P_U})/3$, which is the relaxation-only floor (checked with static Gaussian detunings, `prueba_robusta.py`). With Markovian pure dephasing the floor rises to 6.1–10.1 $\times 10^{-4}$ (certified: 6.03–10.05), so dephasing dominates the increase of the calculated floor (a difference of two bounds, not a decomposition of the measured error). (b) A simple minimal-time pulse comes within a factor 1.4–1.5 of the floor. Among the simple pulses we tried with the same coherence times (the simultaneous pulse, and two sequential swaps with a wait of 0, 5 or 20 ns in the bus), the simultaneous one is the closest: the sequential ones give 11.2–19.5 $\times 10^{-4}$. This is a statement about these links, not a rule. With no bus loss and dephasing only in the qubits ($G = 1$, $\gamma_\varphi = 10^{-3}$ in both qubits), two sequential swaps last 41% longer than the simultaneous pulse and give 7% less error (6.54 against 7.05 $\times 10^{-4}$): what matters is the time the excitation spends in the noisy elements, not the total duration. Where the qubits are the worse element, a dark-state protocol that keeps the bus empty is the wrong goal. (c) The measured errors are 11–16 times the floor, with pulses 6.5–8 times longer than the minimal time. The range column shows that this gap does not depend on the precise coherence times: even the largest floor in the range is 8–11 times below the measured error. At the measured duration the bound is higher: for Q241–Q246 it is 26.5 $\times 10^{-4}$ at 296 ns, 4.3 times below the measured error (the laboratory gives it for each link). Whether shorter pulses pay depends on whether dephasing stays Markovian when the coupling saturates. The exponential law is a modelling choice in Ref. [4], and $T_{2E} > T_2^*$ there suggests slow noise. The decisive measurement is Ramsey and echo at the flux points visited, a sweep of the duration, and readout of A and of the bus.

8.1. **Which hardware change lowers the floor.** Because the floor depends only on G and on the coherence times of each element, it can rank hardware changes before any pulse is designed. Table 2 doubles one parameter at a time and recomputes the floor of each link.

TABLE 2. Floor on $1-F_B$ ($\times 10^{-4}$) after doubling one parameter, with the coherence times of each element and $g/2\pi = 9$ MHz otherwise. "Worse qubit" is the qubit of the pair with the shorter T_2^*; doubling the T_1 of either qubit changes the floor by less than 1% and is not shown. The last column removes the pure dephasing of the bus ($T_2 = 2T_1$); doubling T_2 alone would violate $T_2 \le 2T_1$ for every bus mode here. Each change keeps the other measured times fixed, which may change the inferred dephasing rate of the same element.

link	now	$2G$	$2T_2^*$ worse qubit	$2T_2^*$ other qubit	$2T_1$ bus	bus, no pure dephasing
Q240–Q247	6.06	3.04	4.44 (Q240)	5.39	5.74	6.03
Q241–Q246	8.07	4.05	4.88 (Q241)	7.99	7.83	7.94
Q242–Q245	6.37	3.19	4.30 (Q245)	6.27	5.96	6.19
Q243–Q244	10.09	5.08	7.19 (Q244)	8.56	9.71	10.05

Doubling the coupling halves the floor in every link (in this range it scales approximately as $1/G$). The next most effective change is the T_2^* of the qubit that dephases faster, which lowers the floor by 27–40%; the other qubit gives 1–15%, and the bus mode, whose coherence matters most in a dark-state picture, at most 7% (doubling its T_1, or removing its pure dephasing). For this relay, then, the design levers are the coupling strength and the dephasing of the worse qubit, in that order. Two cautions: these are changes of the floor, not of the error of a particular pulse, and a larger G may be limited by the free spectral range of the cable, which the next paragraphs take into account.

Which coupler should be strong. With different maximal couplings $G_A \neq G_B$ (Proposition 7.2 and the remark after it) the floor depends on how a fixed coupling budget is shared. For Q241–Q246, whose qubits differ most ($T_2^* = 7.1$ against 60 μs), we use the strengthened and reciprocal floor (term $f_A(T)^2$ and Proposition 7.3; the direction does not matter). It is lowest when the stronger coupler sits on the worse qubit Q241: with a fixed sum of the two maxima, 18 MHz, the floor is 7.71×10^{-4} at (10.9, 7.1) MHz for (Q241, Q246), against 8.07 with equal couplers (−4.4%); with a fixed sum of squares, 7.83 at (10.2, 7.6) MHz (−3.0%). The gain is modest, and it concerns the floor: within the family of two sequential swaps the optimal shares give 11.47 and 11.85×10^{-4}, and neither beats the simultaneous pulse with equal couplers, 11.35. A feasible pulse with the same G_A and G_B lies above the floor in all eight cases tested. (Without the term $f_A(T)^2$ the bound loses its dependence on the emitter's coupling when the emitter has the lowest rate, and the share appears to matter much more; that was an artefact of the weaker bound.)

How far the coupling can go. At 80 cm the free spectral range of IBM's cable is about 150 MHz, the links run at $(g/2\pi)/\mathrm{FSR} = 0.05$–$0.06$, and Ref. [4] reports that the relay scheme has been operated up to $(g/2\pi)/\mathrm{FSR} \approx 0.15$ in other samples, which with this cable would mean $g/2\pi \approx 22.5$ MHz. That is a scenario, not an improvement shown on these links: the qubit frequency and dephasing depend on the coupler flux, and the table keeps all coherence times fixed. Table 3 follows the floor up to that value: it falls approximately as $1/g$, to 2.4–4.1×10^{-4} at 22.5 MHz, a factor 0.40 of the present floor.

TABLE 3. Floor on $1 - F_B$ ($\times 10^{-4}$) against the coupling $g/2\pi$ (MHz), with the coherence times of each element kept fixed. 9 MHz is the present value; 22.5 MHz corresponds to $(g/2\pi)/\mathrm{FSR} = 0.15$ with the 80-cm cable.

link	9	12	15	18	22.5
Q240–Q247	6.06	4.55	3.65	3.04	2.43
Q241–Q246	8.07	6.07	4.86	4.05	3.25
Q242–Q245	6.37	4.79	3.83	3.19	2.56
Q243–Q244	10.09	7.59	6.09	5.08	4.07

Two readings of Table 3 rest on stated assumptions, not on data. *Cable length:* the free spectral range scales as $1/L$, so at a fixed ratio $(g/2\pi)/\mathrm{FSR} = 0.15$ the allowed coupling is about $22.5\,\mathrm{MHz} \times (0.8\,\mathrm{m}/L)$; this is a ceiling set by holding that ratio fixed, not a law of the coupling the hardware realises. If the coherence times of all elements (qubits and the T_1 and T_2 of the mode) do not depend on L, the floor at the largest allowed coupling grows roughly in proportion to L: from 2.4–4.1×10^{-4} at 0.8 m to the present 6.1–10.1×10^{-4} at 2 m, where the allowed coupling falls to 9 MHz. *Several modes:* G bounds the norm of the coupling vector, so coupling each qubit to M modes with g per mode gives at most $G = \sqrt{M}\, g$; with two modes at 9 MHz the floor can be no lower than its value at 12.7 MHz, between the 12 and 15 MHz columns (with the loss of the least lossy mode). This bounds the floor, not the gain of an actual protocol.

How fast the couplers must be. Real couplers cannot jump to full strength. If $\|\dot{g}_i\| \leq S$ and $g_i(0) = g_i(T) = 0$, then $\|g_i(t)\| \leq m(t) = \min(St, S(T-t), G)$ (Remark 4.2), and Proposition 7.2 holds with G replaced by the envelope $m(t)$: the emitter window compares with $f' = -m(t)\sqrt{1-f^2} - \frac{d_A}{2} f$; the receiver window, run backwards, with the non-autonomous $w'(s) = -m(T-s)\sqrt{1-w^2} + \frac{d_B}{2} w$, $w(0) = \zeta$ (which may rise at first, since $m(T) = 0$; in the numerics we use a solution clipped to be non-increasing, which lies below the true one and so gives a weaker but valid window); the reach with $M' = m(t)\sqrt{1-M^2} - \frac{d_B}{2} M$; and the bus term becomes $d_C \zeta^2 / J_B^m$ with $J_B^m = \int_0^T m(t)^2 e^{-d_B(T-t)}\, dt$, which already contains the coupling (the proof of Proposition 7.2 with G replaced by $m(t)$, using

comparison for non-autonomous equations; not formalised; evaluated numerically). Here S is the slope of the angular coupling, in rad μs^{-2}; tables quote $S_f = S/(2\pi \cdot 10^3)$, the slope of $g/2\pi$ in MHz/ns. Ref. [4] does not report the slew rate of its couplers, so Table 4 gives the floor for a range of them. With a coupling that reaches 9 MHz in about 3 ns (a slope of 3 MHz/ns), the floor is 9–10% above that of ideal couplers; at 0.3 MHz/ns, a 30-ns ramp, it doubles. These numbers compare floors with and without the slope limit, not the error of a pulse. Keeping the same relative penalty at a larger coupling needs a slope that grows as G^2: $3 \times (22.5/9)^2 \approx 19$ MHz/ns at 22.5 MHz. As a check, a feasible trapezoidal pulse with the same slope and maximum gives errors above the floor in all four cases tested: Q240→Q247 and Q241→Q246 (the direction written first), at 1 and 0.3 MHz/ns, with 10.8 and 15.2, and 13.9 and 19.8×10^{-4}, against the floors 7.9, 12.3, 10.5 and 16.3×10^{-4} (Table 4).

TABLE 4. Floor on $1 - F_B$ ($\times 10^{-4}$) when the coupling $g/2\pi$ rises and falls with a slope of at most S (MHz/ns), up to 9 MHz, with the coherence times of each element; in brackets, the factor over the floor without a slope limit (the first column; both directions, by reciprocity).

link	no limit	$S = 3$	$S = 1$	$S = 0.3$	$S = 0.1$
Q240–Q247	6.06	6.66 (1.10)	7.90 (1.30)	12.27 (2.02)	21.15 (3.49)
Q241–Q246	8.07	8.84 (1.10)	10.48 (1.30)	16.25 (2.01)	27.85 (3.45)
Q242–Q245	6.37	6.96 (1.09)	8.18 (1.28)	12.54 (1.97)	21.58 (3.39)
Q243–Q244	10.09	11.08 (1.10)	13.15 (1.30)	20.42 (2.02)	34.96 (3.46)

Towards 10^{-3} *and* 10^{-4}. The floors say what no pulse can beat; a design target also needs a pulse that reaches it. With the present coherence times, a trapezoidal pulse $g_A = g_B$ of maximum 18 MHz, slope 3 MHz/ns and area $\int g = \pi/\sqrt{2}$ lasts 25.6 ns and gives $1 - F_B = 5.7$, 7.4, 5.6 and 9.3×10^{-4} on the four links in a full Lindblad simulation: below 10^{-3} everywhere, if the coherence times stay as measured at the larger coupling and the one-mode reduction still holds for these shorter pulses. For 10^{-4}, raising the coupling alone is not enough within the scenario $(g/2\pi)/\mathrm{FSR} \leq 0.15$: the coherence times must improve too. One sufficient combination is qubits with $T_1 = T_2 = 150$ μs, a bus mode with $T_1 = T_2 = 75$ μs, $g/2\pi = 22.5$ MHz and a slope of 10 MHz/ns: the floor with that slope is 0.61×10^{-4} and a 18.0-ns trapezoid gives 0.82×10^{-4}, so $0.61 \times 10^{-4} \leq E_{\mathrm{opt}} \leq 0.82 \times 10^{-4}$. This is a calculated combination within the model, not a minimum of resources nor a demonstrated performance.

8.2. **Other platforms.** The same bounds apply to other published transfers through a bus mode, within an effective one-mode description. We took three experiments whose transfer runs through one lossy bus mode with bounded couplings, read their parameters from their tables (in the images, to avoid misaligned columns), and computed the floor with the coherence times of each element. For the bus, whose T_2 none of the three reports directly, we take the minimum of the floor over $T_2 \in [T_1/2, 2T_1]$; it is attained at $T_2 = 2T_1$, the rate $1/T_1$ of the amplitude problem without pure dephasing: since $2/T_2 = 1/T_1 + 2\gamma_\varphi \geq 1/T_1$ under H3, a lower rate only weakens the balance of Proposition 7.2, so the floor holds whatever the pure dephasing of the bus. (Chang *et al.* report $T_{2r} \approx 2T_{1r}$; the other two do not report it.) G is the coupling to that mode in the pulse used, so each row is a floor of the effective one-mode model for every pulse with the same maximal coupling. Two of the devices are described in their own papers with several line modes [14, 15]; since G bounds the norm of the coupling vector, extending the floors to all controls of the full device would need that reduction justified, or the multimode model with its own norm and rates.

TABLE 5. Floors of effective one-mode models and measured errors of other transfers through a bus mode (each name opens the laboratory with its parameters). Heya *et al.*: SATD through a 60-cm coplanar waveguide; error per state transfer from network benchmarking, which estimates $1 - F_B$ under the assumptions of the benchmarking. Chang *et al.*: adiabatic transfer through a tunable lossy channel; readout-corrected process fidelity $F_p = 0.96 \pm 0.01$ from process tomography, converted to the average fidelity $(2F_p + 1)/3$. Burkhart *et al.*: transfer between two 3D cavities through a bus mode at the nominal coupling used; the measured loss $1-\eta$, with η the single-photon transfer efficiency, is compared with the floor on $1 - P$. *Ratio*: measured over floor, an upper bound on the factor by which a better pulse could lower the error within the model, not a measure of how far the device is from its optimum.

experiment	$G/2\pi$ (MHz)	metric	floor	measured	ratio
Heya *et al.* [14]	3.5	$1 - F_B$ (10^{-4})	26.8	120 ± 10	4.5
Chang *et al.* [15]	15 (peak)	$1 - F_B$ (10^{-4})	49.9	267 ± 67	5.4
Burkhart *et al.* [16]	0.56 (nominal)	$1 - P$	0.021	0.16 ± 0.01	7.6
IBM (Table 1)	9	$1 - F_B$ (10^{-4})	6.03–10.05 (cert.)	71–114	11–16

The measured errors are 4.5–16 times the computed floors, and IBM's meter-scale links have the largest ratio. A large ratio means that the bound leaves room for a better pulse, not that the device is the furthest from its own optimum: a weaker bound raises the ratio without any change in the device, and Burkhart *et al.* use a different metric. With the 4-MHz coupling that Heya *et al.* also calibrated, their floor falls to 23.5×10^{-4}. Three cautions. In Burkhart *et al.* the end elements are cavities and the conversion passes through pumped conversion transmons, so the qubit–bus–qubit model is an effective description of their beam-splitter transfer. In Chang *et al.* 15 MHz is the peak of the pulse used, not a limit of the device. Two other transfers, through itinerant photons (Kurpiers *et al.* [17], Magnard *et al.* [18]), involve cascaded resonators, propagation through a transmission line, an effective $|f0\rangle \leftrightarrow |g1\rangle$ drive and, in the first, leakage to $|f\rangle$ outside the qubit subspace; we have not verified a reduction of those systems to the model of Section 3 for the observables they report, and we leave them out.

9. **Prior work and what is new**

The constant. Wang *et al.* [2] optimised STIRAP-based transfer through a lossy bus perturbatively. With separately bounded couplings and the pulse shape optimised within their Fourier ansatz, Eq. (11) of their preprint gives the first-order loss $\kappa t_f + (\gamma/t_f)(1/G_1^2 + 1/G_2^2)$ and their Eq. (12) its minimum over t_f, $4\sqrt{1/C_1 + 1/C_2}$ with $C_i = 4G_i^2/(\gamma\kappa)$; in their notation κ is the qubit decay and γ the bus decay. For $G_1 = G_2 = G$ this is $2\sqrt{2\kappa\gamma}/G$: the constant of Theorem 6.2. Luo, Fang, Wang and Chesi [3] cast the leading-order loss as $\int[\kappa\dot\theta^2/G^2 + \gamma_1\cos^2\theta + \gamma_2\sin^2\theta]$ (plus dephasing terms) and obtain, for separately bounded couplings, $\Delta F_{\min} = 2\sqrt{\kappa\gamma_1}\sqrt{1/G_1^2 + 1/G_2^2}$ and the optimal time t_f^*; with $G_1 = G_2 = G$ this is $2\sqrt{2}\sqrt{\kappa\gamma}/G$ (this special form takes equal qubit relaxations and no dephasing). They state that their formalism does not apply when $\kappa \sim \gamma$, and their bound holds within the quasi-adiabatic approximation. Both results are optima within a perturbative treatment and a family of pulses. What Theorem 6.2 adds is the same constant (with $\kappa - \gamma$ in our notation) as a lower bound over all controls; Theorem 6.3 shows that an explicit pulse reaches it, with explicit error terms, in the full lossy dynamics.

The compensation law and the norm balance. The structure "loss $\propto x$ in one channel, $\propto 1/x$ in the other, minimum $\propto 2\sqrt{\text{rate}\cdot\text{rate}}/\text{coupling}$" appears in Goto and Ichimura [5] for cavity-mediated

atom–atom transfer in the adiabatic limit and for a family of pulses (as described by Suenaga *et al.* [6]; we could not access the original text), and in Goto *et al.* [7] as a bound for every control on single-photon generation, obtained from a balance of integrated occupations. The norm balance itself is classical: in the cavity model of Gorshkov *et al.* [8] the efficiency of complete retrieval is $C/(1+C)$, independent of the retrieval technique, and optimised storage can reach it with a suitable control and input mode. Closest in method is Hughes, Goodwin and Horak [9]: a bound at finite T for every pulse, obtained by a similar relaxation and of the same form $1/(1+m)$; the physical task differs (irreversible extraction of a photon through the cavity loss, at fixed coupling, rather than transfer between two lossy qubits with bounded couplings). Yuan *et al.* [10] solve a Hamilton–Jacobi–Bellman problem with one bounded and one unbounded control and no qubit loss. Stefanatos and Paspalakis optimise with a joint amplitude bound: in Ref. [11] after adiabatic elimination of a strongly dissipative intermediate level, with constant total intensity; in Ref. [12] for the full dynamics, within a family with a constant total amplitude and a monotone mixing angle. We found no earlier statement of Theorem 4.1 (finite T, lossy qubits, separately bounded couplings, any number of modes) or of Theorem 5.2.

Dephasing. That dephasing depopulates the dark state is known [13]; Wang *et al.* give the first-order term $\bar{\gamma}_\varphi \int \sin^2 2\theta$ in their supplement [2]. Hardware analyses [4, 14, 19] use Lindblad models with T_2; Ref. [4] uses an additive budget, time in each element times $(1/3T_1 + 1/3T_\varphi)$, as a first-order approximation. Non-Markovian $(1/f)$ dephasing changes the picture qualitatively [20]. Lemma 7.1 is elementary; we did not find it used to carry loss bounds over to the fidelity.

Remark 9.1 (Joint bound). With a joint bound $|g_A|^2 + |g_B|^2 \le G^2$ the constant changes: in the adiabatic-elimination limit the optimum of Ref. [11] gives a loss $\approx (\pi^2/4)\,\Gamma/(g^2 T)$ in our units, against $2\delta/(G^2 T)$ for the box (Lemma 6.1). The box is the natural constraint when each coupler has its own maximal strength.

10. Limitations

The bounds assume zero-temperature pure loss (H1) and no direct qubit–qubit coupling (H2). The fidelity bounds assume Markovian dephasing (H3); with $1/f$ or quasi-static noise the law of c changes. The IBM numbers reduce the bus to one effective mode (H4) and use the reported coherence times of each element; the range in Table 1 shows how the floors move with them, not a statistical uncertainty. The floors use only the rates of each element: correlated or non-Markovian noise and the pulse-dependent change of the qubit frequency are outside the model. The lower bound of Theorem 6.2 has no explicit error terms (the upper bound of Theorem 6.3 does), and Theorem 6.3 is proved on paper only. The floor with ramped couplings of Remark 4.2 is not shown to be tight. Of Proposition 7.2, only the window comparisons (for $d > 0$) and a one-rate version of the norm chain are formalised, with the receiver's barrier as a hypothesis and exceptions in a countable set; its version with ramped couplers follows the same proof with $G \to m(t)$ and is evaluated numerically. The formalised floor and frontier lemma assume exceptions in a countable set (piecewise-smooth controls). All fidelity floors assume Markovian dephasing at the measured T_2^*; slow noise would change the dephasing part, which dominates the IBM floors, but not the relaxation-only floor of 1.5–1.8×10^{-4}, which holds for any classical dephasing noise. Proposition 7.3 is proved on paper (sketch) and checked numerically; it assumes that the time-reversed pulse, phases included, is realisable, and does not cover non-reciprocal buses or memory between transfers. The pulses quoted as targets (10^{-3}, 10^{-4}) assume that the coherence times stay as measured at larger couplings and faster ramps. The design comparisons of Section 8.1 rank floors, not the errors of particular pulses, and the readings in cable length and number of modes rest on the stated assumptions. For the other platforms (Section 8.2) the floors are those of effective one-mode models, with the published parameters as given and the coupling of the pulses they used; two of those

devices are multimode in their own description, and the model is an effective description of the cavity transfer of Burkhart *et al.* The four floors of IBM's links are certified with interval arithmetic (column *certified* of Table 1); all other floors are numerical evaluations of rigorous bounds, without interval enclosures.

11. Verification and supplementary files

The supplementary package contains: (i) `isabelle/Quantum_Speed_Limits/`, the Isabelle entry, public in the repository karlesmarin/formalizations, which also holds the speed-limit inequalities of Part I [1]; among them `Relay_Loss_Floor.thy` with `loss_floor` (Theorem 4.1) and `partial_frontier` (Lemma 6.1), `Relay_Endpoints.thy` with `endpoint_window_ac` (Lemma 5.1) and `endpoint_occupation_ac` (Theorem 5.2) and their earlier versions for countable exceptions, and `Relay_Windows.thy` with `emitter_window`, `receiver_window` and `norm_chain` (Proposition 7.2); the entry builds with no `sorry` and without `quick_and_dirty`; (ii) `scripts/`, the Python and Sage scripts behind every number above (optimal control, Lindblad simulation, the exact Sage certificate of $\Phi(1) = \pi/4$, the attaining pulse of Theorem 6.3, the floor with a rate per element of Proposition 7.2 with its symmetric control, the IBM table with its range, the simple pulses compared in Section 8, the hardware changes of Tables 2–4 with the check of the ramped floor against feasible pulses, the other platforms of Table 5 with their transcribed data (`scripts/datos/plataformas.toml`, each value with its source), the checks of the symmetric-model counterexample and of the sequential-swap example, and the IBM calculator with its independent replica), with the IBM data in `scripts/datos/`; (iii) `figures/`, the scripts and data that draw Figures 1–5. Each script prints the value together with its verdict. The companion site, https://karlesmarin.github.io/blueprints/relay-floors/, links each result to its proof, Isabelle source and scripts, and runs the laboratory Check your link in the browser.

Appendix A. Formula sheet: which bound to use, and when

Notation: G maximal coupling (norm of the coupling vector), γ qubit loss rate, κ bus loss rate, $\delta = \kappa - \gamma$, T duration, P population delivered to B, c coherence of B with the vacuum, F_B fidelity of the channel at B. "Every pulse" means every control with $\|g_A\|, \|g_B\| \le G$, any shape, phase and number of bus modes, under H1–H2.

formula	valid when	use it to	status, where
$P(T) \le B(T) = \dfrac{e^{-\gamma T}}{1 + \delta/(G^2 T)}$	every pulse, every T, $\kappa \ge \gamma$	a lower bound $1 - B(T)$ on the loss at a given duration; minimise $1 - B(T)$ over T for the floor of a link	Isabelle; Thm. 4.1
$T_{\mathrm{bound}} = \dfrac{\sqrt{a^2 + 4a/\gamma} - a}{2}$, $a = \delta/G^2$	as the row above, with $\gamma > 0$ and $\delta > 0$	the duration that minimises the bound (not the duration of the best pulse)	exact; Sec. 4
$1 - P \ge 1 - \left(1 + \dfrac{\sqrt{\gamma\delta}}{G}\right)^{-2}$	same as the first row, every T	a rigorous estimate by hand; to first order $2\sqrt{\gamma\delta}/G$	from Thm. 4.1
$1 - P \ge \dfrac{2\sqrt{2\gamma\delta}}{G}\,(1 - o(1))$	$\alpha = \sqrt{\gamma\delta}/G \to 0$, $\beta = \gamma/\delta \to 0$	the asymptotically exact constant; not for present devices ($GT \approx 2$–3)	paper; Thm. 6.2

formula	valid when	use it to	status, where
$1-P \le \ell_0\big(\frac{1}{1-\alpha}+1.5\beta\big)$, $\ell_0 = 2\sqrt{2\gamma\delta}/G$	$\gamma,\delta>0$, $\alpha,\beta\le 0.01$; one bus mode with loss exactly κ	a guarantee for the explicit ramp–arc–unload pulse (Figure 4)	paper; Thm. 6.3
$T_{\mathrm{asym}} \sim \sqrt{2\delta/\gamma}\,/G$	$\alpha,\beta\to 0$	the leading time scale of the asymptotically best pulses, a factor $\sqrt{2}$ above T_{bound}	perturbative [2, 3]; Thm. 6.3
$G\int_0^T (x^2+z^2)\,dt \ge \frac{\pi}{4}+\Phi(z(T))$	rescaled moduli $x=\lvert\tilde{a}\rvert$, $z=\lvert\tilde{b}\rvert$, A emptied; $\to \pi/2$ for a complete transfer, possible only without bus excess loss	the occupation of the end qubits in the frame where they are lossless; for lossy end qubits use the windows of Prop. 7.2	Isabelle; Thm. 5.2
$T_{\min} = \dfrac{\pi}{\sqrt{2}\,G}$	no losses, box bound	the shortest complete transfer; the minimal-time pulse $g_A = g_B = G$	Isabelle (`Relay_Minimal_Time`); Table 1
$G^2T \to J(T) = \int_0^T m^2$, $m=\min(St, S(T-t), G)$	hypotheses of the first row, plus slew rate $\lVert\dot{g}\rVert \le S$ and $g=0$ at both ends	the floor with ramped pulses ($J=S^2T^3/12$ for $T\le 2G/S$, $G^2T-4G^3/(3S)$ for $T\ge 2G/S$)	paper; Rem. 4.2
the row of Prop. 7.2 with $G\to m(t)$, $J_B=\int_0^T m^2 e^{-d_B(T-t)}dt$	rates per element and slew rate S	the floor of a real device with ramped couplers; how fast the couplers must be	paper; Sec. 8.1; Table 4
$\gamma \to \dfrac{2}{T_{2q}}$, $\quad \kappa \to \dfrac{2}{T_{2c}}$	Markovian dephasing (H3); each bound keeps its own hypotheses, including $\kappa \ge \gamma$	the lower bounds of Theorems 4.1–6.2 applied to c^2 instead of P; for one pulse $P\ge c^2$, so a guarantee $1-c^2\le L$ (Thm. 6.3 at the rates $2/T_2$, with all its hypotheses) gives $1-F_B \le L/6+(1-\sqrt{1-L})/3$	paper; Lemma 7.1; checked to 10^{-16}
$1-F_B = \dfrac{1-P}{6}+\dfrac{1-c}{3}$	one excitation, Lindblad local; P from the input $\lvert A\rangle$, c from $(\lvert 0\rangle+\lvert A\rangle)/\sqrt{2}$, phase of B corrected (Sec. 7)	turning bounds on P and c into the error that network benchmarking measures	exact; Eq. (5)
$c^2 \le \dfrac{e^{-2T/T_{2q}}}{1+\Delta_2/(G^2T)}$, $\Delta_2 = 2/T_{2c}-2/T_{2q}$	H3, equal qubits, bus dephases more ($T_{2c}<T_{2q}$; with several modes, the longest T_2 among all accessible modes)	the dephasing bound when the bus is the worse element	from Thm. 4.1 and Lemma 7.1

formula	valid when	use it to	status, where
$\zeta_U(T) = \sup\{0 \le \zeta \le M_B(T) : \zeta^2 + f_A^2 + d_A a_A + d_B b_B(\zeta) + \frac{d_C \zeta^2}{G^2 J_B} \le 1\}$, amplitude $\le e^{-r_0 T/2} \zeta_U(T)$	any rates per element $(r_A, r_C, r_B$; $d_j = r_j - \min r)$; $2/T_2$ for c, $1/T_1$ for P; f_A, a_A, b_B, M_B, J_B as in Prop. 7.2; G_A for the emitter terms, G_B for the receiver terms	the floor of a real device, with each element's own coherence times; evaluated numerically (the tables include f_A^2 and reciprocity; the four IBM floors are also certified)	paper, windows in Isabelle; Prop. 7.2; Table 1
$P_{B\leftarrow A}[H] = P_{A\leftarrow B}[H_{\mathrm{rev}}]$, same c	local loss and diagonal dephasing; each coupler with its own bound	take, at each T, the smaller bound of the two directions; the direction of transfer is not a design choice	paper; Prop. 7.3; checked to 10^{-15}

Rule of thumb for a design. (1) Use the coherence times of each element, not the worst of a pair applied to both qubits: that is a comparison model, not a floor. (2) Compute upper bounds $c_U(T)$ from the row of Proposition 7.2 at rates $2/T_2$ and $P_U(T)$ from the same row at rates $1/T_1$ (or from the first row if the qubits are equal and the bus is the worse element). (3) The floor of the link is the rigorous bound

$$L_{\mathrm{cert}} = \inf_{T>0} \Big[\frac{1 - P_U(T)}{6} + \frac{1 - c_U(T)}{3} \Big],$$

taken over both directions of transfer; it bounds the error of every pulse from below. Its value can be certified: the window integrals become integrals of analytic functions in the angle θ ($f = \cos\theta$), which ball arithmetic encloses rigorously; each piece is monotone in T, so on $[T_a, T_b]$ one uses T_a for the windows and the decay factor and T_b for the reach, J_B and the term $f_A(T)^2$ (which, with $d_A a_A$, decreases in T), and two tails close the minimum (column *certified* of Table 1). The other values in this paper are numerical evaluations of that bound. (4) The floor does not choose the pulse: evaluate feasible pulses with the same coherence times and constraints, and compare. Their errors give the other side, and the gap tells how much a better pulse could still gain. The duration that minimises a bound is not the duration of the best pulse, and the shortest pulse is not always the best: what matters is the time the excitation spends in the noisy elements (Section 8). (5) To rank hardware changes, recompute L_{cert} with one parameter changed at a time (Table 2).

Use of AI tools

This work was carried out with an AI assistant. The author checked all statements, proofs and computations and takes full responsibility for the content.

References

[1] C. Marín, Certified speed-limit diagnostics for deployed superconducting quantum processors, Zenodo (2026), doi:10.5281/zenodo.23783382.

[2] Y.-D. Wang, R. Zhang, X.-B. Yan and S. Chesi, Optimization of STIRAP-based state transfer under dissipation, New J. Phys. **19**, 093016 (2017), doi:10.1088/1367-2630/aa7f5d; preprint (Y.-D. Wang, X.-B. Yan and S. Chesi) arXiv:1603.01701, whose equation numbers we quote.

[3] S. Luo, Y. Fang, Y. Wang and S. Chesi, Principle of least action for quasi-adiabatic state transfers with dissipation, arXiv:2402.12807 (2024).

[4] N. M. Sundaresan *et al.*, Meter-scale quantum interconnects between large quantum processors, arXiv:2610.29986 (2026).

[5] H. Goto and K. Ichimura, Phys. Rev. A **77**, 013816 (2008), doi:10.1103/PhysRevA.77.013816.

[6] A. Suenaga *et al.*, Speed limit of efficient cavity-mediated adiabatic transfer, arXiv:2401.00384 (2024).

[7] H. Goto, S. Mizukami, Y. Tokunaga and T. Aoki, Figure of merit for single-photon generation based on cavity quantum electrodynamics, Phys. Rev. A **99**, 053843 (2019).
[8] A. V. Gorshkov, A. André, M. D. Lukin and A. S. Sørensen, Photon storage in Λ-type optically dense atomic media. I. Cavity model, Phys. Rev. A **76**, 033804 (2007).
[9] W. J. Hughes, J. F. Goodwin and P. Horak, Optimising finite-time photon extraction from emitter-cavity systems, arXiv:2403.10355 (2024).
[10] H. Yuan, C. P. Koch, P. Salamon and D. J. Tannor, Controllability on relaxation-free subspaces: on the relationship between adiabatic population transfer and optimal control, Phys. Rev. A **85**, 033417 (2012), doi:10.1103/PhysRevA.85.033417.
[11] D. Stefanatos and E. Paspalakis, Optimal shape of STIRAP pulses for large dissipation at the intermediate level, Quantum Inf. Process. **20**, 391 (2021), doi:10.1007/s11128-021-03352-1.
[12] D. Stefanatos and E. Paspalakis, Optimal shortcuts of stimulated Raman adiabatic passage in the presence of dissipation, Phil. Trans. R. Soc. A **380**, 20210283 (2022), doi:10.1098/rsta.2021.0283.
[13] N. V. Vitanov, A. A. Rangelov, B. W. Shore and K. Bergmann, Stimulated Raman adiabatic passage in physics, chemistry, and beyond, Rev. Mod. Phys. **89**, 015006 (2017).
[14] K. Heya *et al.*, Randomized benchmarking of a high-fidelity remote CNOT gate over a meter-scale microwave interconnect, arXiv:2502.15034 (2025).
[15] H.-S. Chang *et al.*, Remote entanglement via adiabatic passage using a tunably dissipative quantum communication system, Phys. Rev. Lett. **124**, 240502 (2020), doi:10.1103/PhysRevLett.124.240502.
[16] L. D. Burkhart *et al.*, Error-detected state transfer and entanglement in a superconducting quantum network, PRX Quantum **2**, 030321 (2021), doi:10.1103/PRXQuantum.2.030321.
[17] P. Kurpiers *et al.*, Deterministic quantum state transfer and remote entanglement using microwave photons, Nature **558**, 264 (2018), doi:10.1038/s41586-018-0195-y.
[18] P. Magnard *et al.*, Microwave quantum link between two superconducting circuits housed in spatially separated cryogenic systems, Phys. Rev. Lett. **125**, 260502 (2020), doi:10.1103/PhysRevLett.125.260502.
[19] M. Malekakhlagh *et al.*, Enhanced quantum state transfer and Bell state generation over long-range multimode interconnects via superadiabatic transitionless driving, arXiv:2401.09663 (2024).
[20] G. Falci *et al.*, Phys. Rev. B **87**, 214515 (2013), doi:10.1103/PhysRevB.87.214515.

Independent researcher, Spain
Email address: karlesmarin@gmail.com