ROBUST QUBIT–MEMORY TRANSFER AGAINST QUASI-STATIC FREQUENCY NOISE COSTS TIME AND EXPOSURE: EXACT BOUNDS FROM DUBINS CURVES

CARLES MARÍN
Independent researcher

Abstract. A quasi-static offset of the qubit or memory frequency, the slowest part of the noise in a superconducting quantum memory, is described by no Markovian rate. For the exchange of one excitation between a qubit and a memory mode at a coupling bounded by G, we characterise every complete transfer by the times A and B it spends in the qubit and in the memory and by its first-order leakage ζ. For a coupling of any phase and any detuning the duration obeys $T \geq T_0 + \arccos(|\zeta|/z_0)/G$ for $|\zeta| \leq z_0$ and the exposures obey $(|\zeta|-z_0)^2 \leq (A-A_0)(B-A_0)$, with $T_0 = \pi/(2G)$, $A_0 = \pi/(4G)$, $z_0 = 1/(2G)$; both are sharp, and for a coupling of fixed phase and no detuning they follow from known results on paths of bounded curvature. Cancelling the leakage with a pulse of fixed phase and no detuning costs exactly $(\gamma_q+\gamma_m)/(2G)\cdot\arccos(|\gamma_q-\gamma_m|/(\gamma_q+\gamma_m))$ of first-order population loss above that of the square pulse (a third of it in average infidelity), through the shortest Dubins U-turn; for any pulse, with relaxation only, cancelling it costs at least $\max(\sqrt{\gamma_q\gamma_m}, \frac{\pi}{2}\min(\gamma_q,\gamma_m))/G$ in the same units; curves made of pieces in planes through the axis do no better in space. At second order in the offsets every pulse has error $[2\,\mathbb{E}(\xi_q-\xi_m)^2|\zeta|^2+\mathbb{E}(\xi_q A+\xi_m B)^2]/6$: with noise on the qubit no pulse goes below $\pi^2\sigma^2/(96G^2)$, against $(8+\pi^2)\sigma^2/(96G^2)$ for the square pulse; with noise on both elements the square pulse is optimal for $0.137 \leq \sigma_m^2/\sigma_q^2 \leq 7.30$ and explicit shaped pulses beat it outside $(0.1264, 7.91)$ (certified by interval arithmetic); for pulses of fixed phase and no detuning, a planar exposure bound derived from the reachable set of the Dubins car (through our reading of Patsko and Fedotov's description of it) replaces the relaxation, and the threshold ratio is exactly the crossing of the bang–bang family, $0.12645352248541\ldots$; with memory loss the square pulse stops being optimal (for fixed phase and no detuning, exactly at $\sigma^2 = 7.9078851606875\ldots\,\gamma_m G$, the crossing of a short-wait turn–wait–turn pulse), and for large $\sigma^2/(\gamma_m G)$ a pulse that waits in the memory attains the optimal error $\pi^2\sigma^2/(96G^2)+\sigma\sqrt{\pi\gamma_m}/(6G^{3/2})$ to leading order; positively correlated offsets on the two elements favour the square pulse, with an elementary sufficient condition that also covers relaxation of both elements, and the band of optimality of the square pulse can only widen as the correlation grows; for fixed phase and no detuning we certify the exact band with correlated offsets and the exact threshold with qubit relaxation at several parameter values, prove, given the planar bound, that this threshold grows as $2\sqrt{r}+\pi/2+4/\sqrt{3}+O(r^{-1/2})$ in the ratio r of qubit to memory relaxation, and above the lossy threshold the optimum is a turn–wait–turn pulse. The robustness inequalities for the space curve, the minimisation of the cost, the unitarity step of the second-order expansion and the algebra of the bounds are machine-checked in Isabelle/HOL; Section 4 says where each formal statement ends and the written proof begins.

Date: November 2026.

Companion site: https://karlesmarin.github.io/blueprints/robust-offset/: every statement with its dependencies and its checks, and labs in the browser; each figure of this paper is a link to the lab that reproduces it with its data. This paper: doi:10.5281/zenodo.24250503 (Zenodo), with its supplementary material.

Results at a glance

- **Robustness costs time and exposure:** the frontier between duration and leakage and an inequality between the times in the qubit and in the memory, for a coupling of any phase and any detuning (Theorem 1; for a coupling of fixed phase and no detuning they follow from Souères–Balluchi–Bicchi and Patsko–Fedotov, and the minimum time at zero leakage is due to Van Damme et al.). formalised for the space curve
- **The exact cost of cancelling the leakage** with a pulse of fixed phase and no detuning (Proposition 3); pieces in planes through the axis do not shorten the U-turn. written proof, audited, comparison formalised
- **The second-order error of every pulse**, a floor $\pi^2\sigma^2/(96G^2)$ for noise on the qubit, the square pulse optimal or beaten in certified ranges, and the gain of waiting in a lossy memory (Propositions 5 and 8). written proof, audited, ranges certified
- **Exact thresholds for pulses of fixed phase and no detuning:** the square pulse loses its optimality exactly at the crossing of a competing family, $\varrho_c = 0.12645352248541\ldots$ and $\Lambda_c = 7.9078851606875\ldots$, given a planar exposure bound from the reachable set of the Dubins car (Lemma 6); above Λ_c the optimum among these pulses is a member of the turn–wait–turn family (Proposition 8(iii)). certified, written proof, audited
- **Relaxation and correlation:** with qubit relaxation the exact threshold $\Lambda_c(\gamma_q/\gamma_m)$, certified at nine ratios, grows from 7.908 to 67.245 (Remark 9); a common offset leaves the square pulse optimal, opposite offsets are cancelled by a closed curve, and in between the exact band $[\varrho_c(\kappa), 1/\varrho_c(\kappa)]$ widens with the correlation but never covers every ratio (Remark 10); one elementary condition covers both cases (Remark 11). certified, formalised
- **Scope:** two states, one excitation, a static offset, first and second order; the limits of the exchange under Markovian relaxation and dephasing are in the companion paper [10].

formalised machine-checked in Isabelle/HOL

certified numbers enclosed by interval arithmetic

written proof, audited proof in the text, audited step by step and tested numerically

Contents

1 Introduction . . . 3
2 Robustness costs time and exposure . . 3
2.1 Exposures, leakage and the two robustness bounds . . . 3
2.2 The exact cost of cancelling the leakage . . . 6
2.3 Slow noise at second order . . . 9
2.4 Fixed phase and no detuning: the planar exposure bound . . . 11
2.5 Slow noise together with memory loss 13
2.6 Correlated offsets . . . 17
3 Limits of this paper . . . 19
4 What backs each statement . . . 20
Appendix A Formula sheet . . . 23
Appendix B The companion site . . . 24
Acknowledgement . . . 24
References . . . 25

1. Introduction

A quantum memory attached to a superconducting qubit is written by exchanging one excitation between the qubit and a long-lived mode at a coupling rate bounded by G [10]:

$$i\dot{a} = \big(D(t) - \tfrac{i}{2}\gamma_q\big)a + g(t)\,b, \qquad i\dot{b} = -\tfrac{i}{2}\gamma_m\,b + \overline{g(t)}\,a, \qquad |g(t)| \le G, \tag{1}$$

with a complex coupling g and a detuning D as controls. The companion paper [10] gives the exact limits of this exchange under Markovian relaxation and dephasing and shows that, there, no pulse transfers more population than the resonant swap. **Slow noise is different: a frequency offset that stays constant during one transfer is described by no rate, and pulse shaping can then help.** The error of the write is measured, as in [10, Eq. (6)], by the average infidelity $r = (1-p)/6 + (1-c)/3$ of the channel with transferred population p and transferred coherence c.

This paper treats a static offset at first and second order. Section 2 sets up the exposures A, B and the leakage ζ of a lossless transfer, proves the frontier between duration and leakage and the inequality between exposures (Theorem 1), computes the exact cost of cancelling the leakage with a pulse of fixed phase and no detuning (Proposition 3), and gives the second-order error of every pulse, alone and with memory loss (Propositions 5 and 8). Section 3 states what is not covered, Section 4 what backs each statement, and Appendix A collects the formulas.

Earlier work. Van Damme, Ansel, Glaser and Sugny found the minimum time π/G of an inversion robust to first order [18]. We use the space-curve picture of quantum control of Barnes and coworkers [20, 3, 19], the leakage whose vanishing Daems et al. [5] and Lu et al. [9] impose for robust inversion, and the geometry of curves of bounded curvature: the shortest paths of Dubins [6] and of Souères, Balluchi and Bicchi [15], the reachable sets of Patsko and Fedotov [13, 14], and the estimate on angles in the proof of Schur's comparison theorem [16]. Composite pulses with a detuning that alternates in sign [8] are robust to errors in the pulse area, a different requirement: their two-pulse inversion transfers completely but keeps a first-order leakage $|\zeta| \approx 1.11\,z_0$, so it is not robust to a static offset in our sense (our check, consistent with Theorem 1). What is new and what is not is said in detail in Section 2.

2. Robustness costs time and exposure

2.1. Exposures, leakage and the two robustness bounds. Under slow noise the relevant quantities of a complete transfer are time integrals of the projector on the qubit in the Heisenberg picture. Take no losses, a control with $|g| \le G$ and any detuning $D(t)$, and a duration T with $U(T)|a\rangle \propto |b\rangle$. Put

$$K = \int_0^T U^\dagger(t) P_a U(t)\,dt = \begin{pmatrix} A & \bar{\zeta} \\ \zeta & B \end{pmatrix}, \qquad A + B = T.$$

A and B are the times spent in the qubit and in the memory, and ζ is the first-order leakage: a static shift Δ_0 of the qubit level changes $U(T)$ into $U(T)(I - i\Delta_0 K) + O(\Delta_0^2)$, so the amplitude left in the qubit is $\Delta_0|\zeta|$ to first order and the phase of the transferred amplitude moves by $\Delta_0 A$. Write $U^\dagger P_a U = \frac{1}{2}(I + \mathbf{n}\cdot\boldsymbol{\sigma})$; then $\mathbf{n}(0) = \mathbf{e}_z$, $\mathbf{n}(T) = -\mathbf{e}_z$, $|\dot{\mathbf{n}}| = 2|g| \le \lambda := 2G$ (the detuning commutes with P_a and does not enter), $A = \frac{1}{2}\int(1 + n_z)$, $B = \frac{1}{2}\int(1 - n_z)$ and $|\zeta| = \frac{1}{2}|\int \mathbf{n}_\perp|$. Let $T_0 = \pi/(2G)$, $A_0 = \pi/(4G)$ and $z_0 = 1/(2G)$, the values of the square pulse. The curve $\mathbf{r}(t) = \int_0^t \mathbf{n}$ is the one of the space-curve picture of quantum control [20, 3]:

its curvature is $2|g(t)|$, its length is the duration, and minimum times become shortest curves under a curvature bound [19, 12]. For gates the first-order robustness asks the curve to close; for the population of a transfer only its lateral displacement $2|\zeta|$ counts (writing an arbitrary qubit state also keeps the longitudinal phase exposure A), and $\zeta = 0$ is the condition for robust inversion against a static detuning of Daems et al. [5] and Lu et al. [9].

Theorem 1 formalised for the space curve. *For every complete transfer without losses, with $|g| \le G$ and any detuning:*

(a) *$A \ge A_0$, $B \ge A_0$ and*

$$\boxed{(|\zeta| - z_0)^2 \le (A - A_0)(B - A_0)};$$

(b) *if $|\zeta| \le z_0$, then*

$$\boxed{T \ge T_0 + \frac{1}{G} \arccos \frac{|\zeta|}{z_0}}.$$

(b) is sharp for every $|\zeta| \le z_0$: equality holds for $g = +G$ during $(\pi + d)/(2G)$ followed by $g = -G$ during $d/(2G)$, $d = \arccos(|\zeta|/z_0)$, and for its time reverse. (a) is an equality on the branch $|\zeta| \ge z_0$: for $+G$ during $\alpha/(2G)$, no coupling and no detuning during w, and $+G$ again during $(\pi - \alpha)/(2G)$, for every $\alpha \in [0, \pi]$ and $w \ge 0$ (the square pulse is $w = 0$).

Proof. Choose a time t_e with $n_z(t_e) = 0$ and put $\mathbf{e} = \mathbf{n}(t_e)$, $x(t) = \mathbf{n}(t) \cdot \mathbf{e}$. Since $\mathbf{e}$ is equatorial, $\int x\,dt = \mathbf{e} \cdot \int \mathbf{n}_\perp \le 2|\zeta|$. The angle $\psi(t)$ between $\mathbf{n}(t)$ and $\mathbf{e}$ is at most $\lambda|t - t_e|$, and, because the end points are the poles, at most $\pi/2 + \lambda t$ and $\pi/2 + \lambda(T - t)$; in particular $\lambda t_e \ge \pi/2$ and $\lambda(T - t_e) \ge \pi/2$.

(b) Put $\alpha = \lambda t_e - \pi/2$ and $\beta = \lambda(T - t_e) - \pi/2$, so $\alpha + \beta = \lambda T - \pi =: \Theta$, and assume $\Theta \le \pi$ (otherwise there is nothing to prove). The bound $\bar\psi(t) = \min(\lambda|t - t_e|,\ \pi/2 + \lambda t,\ \pi/2 + \lambda(T - t))$ is at most $\pi/2 + \max(\alpha, \beta)/2 \le \pi$, where the cosine decreases, so $x = \cos\psi \ge \cos\bar\psi$; integrating this piecewise linear angle on each side of t_e gives

$$2|\zeta| \ \ge\ \int_0^T x\,dt \ \ge\ \frac{2}{\lambda}\Big(\cos\frac{\alpha}{2} + \cos\frac{\beta}{2} - 1\Big) \ \ge\ \frac{2}{\lambda}\cos\frac{\Theta}{2},$$

the last step because $\alpha \mapsto \cos(\alpha/2) + \cos((\Theta - \alpha)/2)$ is concave on $[0, \Theta]$ and so is smallest at an end point. With $\lambda = 2G$ this reads $\cos(G(T - T_0)) \le |\zeta|/z_0$, which is (b).

(a) On $|t - t_e| \le \pi/(2\lambda)$ the speed bound gives $x(t) \ge \cos(\lambda|t - t_e|) \ge 0$, so $\int x_+ \ge 2/\lambda$. With $q = |\mathbf{n}_\perp| \ge |x|$ and $Q = \int q$, this gives $Q + 2|\zeta| \ge \int(q + x) \ge 2\int x_+ \ge 4/\lambda$, and with $2|\zeta| \le Q$ we get $|2|\zeta| - 2/\lambda| \le Q - 2/\lambda$. Let $\theta = \arccos n_z$, the spherical distance of n from the north pole, which is λ-Lipschitz and hence absolutely continuous (the formal proof avoids θ near the poles with a comparison argument, so that curves through a pole are covered), and $h = 1 - \dot\theta/\lambda \ge 0$; the fundamental theorem of calculus gives $\int(1 \pm \cos\theta)\dot\theta = \pi$ and $\int \sin\theta\,\dot\theta = 2$ whether or not θ is monotone, so $A - \pi/(2\lambda) = \frac{1}{2}\int(1 + \cos\theta)h$, $B - \pi/(2\lambda) = \frac{1}{2}\int(1 - \cos\theta)h$ and $Q - 2/\lambda = \int \sin\theta\, h$. Cauchy–Schwarz gives $Q - 2/\lambda \le 2\sqrt{(A - \pi/(2\lambda))(B - \pi/(2\lambda))}$, and the two inequalities give (a).

The equality cases are direct computations: for the bang–bang pulse of (b), $T = (\pi + 2d)/(2G)$, $|\zeta| = \cos d/(2G)$, $A = A_0 + (d - \sin d)/(2G)$ and $B = A_0 + (d + \sin d)/(2G)$; its time reverse exchanges the two values of $A - A_0$ and $B - A_0$. □

Together, (a) and (b) give the whole frontier between duration and leakage: the smallest $|\zeta|$ at duration $T \in [T_0, \pi/G]$ is $z_0 \cos(G(T - T_0))$, and $\zeta = 0$ is reachable for every $T \geq \pi/G$; at prescribed $|\zeta| \geq z_0$, (a) and $x + y \geq 2\sqrt{xy}$ give $T \geq T_0 + 2(|\zeta| - z_0)$, attained by the pulse of the theorem with $\alpha = \pi/2$. On the branch $|\zeta| < z_0$, (a) is not attained by these pulses: at $\zeta = 0$, among pulses of fixed phase and no detuning and at fixed $A - B$, the smallest product $(A - A_0)(B - A_0)$ is $(2\pi/3)^2 z_0^2 \approx 4.39\, z_0^2$ for $A = B$ and $(\pi^2/4 - 1) z_0^2 \approx 1.47\, z_0^2$ for the pulses of least duration, and tends to z_0^2 only as $|A - B|$ grows without bound.

What is known. With $|g| \leq G$ the curve $\mathbf{r}$ has unit speed and curvature at most λ: it is a Dubins curve of minimum radius $z_0 = 1/\lambda$, and a complete transfer is a U-turn, with lateral displacement 2ζ and axial displacement $A - B$. Zeng and Barnes found the fastest robust phase gates as shortest closed plane curves of bounded curvature [19]; the two ends of (b) are known values, T_0 for $|\zeta| = z_0$ (the π pulse, the elementary polar-angle speed limit for population inversion under $|g| \leq G$; related gate speed limits under a bounded spectral width of the Hamiltonian are derived in [12]) and π/G for $\zeta = 0$ (below), and (b) is the exact curve between them. For $g = e^{i\varphi} u(t)$ with u real, φ constant and no detuning (or a detuning compensated by the phase of g) the curve is planar; in general it leaves the plane. In the planar case (b) is the case of reversed heading of the shortest path of a bounded-curvature vehicle to a directed line with free arrival point, synthesised by Souères, Balluchi and Bicchi [15] (their paths $C_a C_b$ at the edge $a = b + \pi R$, which give the length in closed form), and (a) is, in substance, a consequence of the description of the reachable set of the Dubins car by Patsko and Fedotov [13, 14]: at heading π the circle bounding (a) is their arc A_1, which lies on the boundary of the section for every duration, the other boundary arcs lie inside its disk [14, Property 8 and Theorem 11.1], and for headings $\pi + 2\pi k$, $k \geq 1$, the section is a smaller disk with the same centre; they do not write the inclusion as an inequality. The case $\zeta = 0$ of (b), $T \geq \pi/G$ with the bang–bang pulse $(3\pi/(4G), \pi/(4G))$, is the minimum time of a first-order robust inversion found by Van Damme et al. [18], already with two controls, by the maximum principle and a numerical search, and rederived for a real control by Ansel, Glaser and Sugny [1]. What the proof above adds is a complete elementary proof of that minimum, both statements for a coupling of any phase and any detuning, where the curve leaves the plane, (a) written as an inequality between exposures, and the reading of both in terms of leakage: (a) says how the extra time must be shared between the two subsystems. The bound on the angle to $\mathbf{n}(t_e)$ is the step of the classical proof of Schur's comparison theorem for curves of bounded curvature [16], so (b) is an adaptation of that argument to an antipodal U-turn. In the sources on bounded-curvature curves in space that we read, neither statement appears: for the reachable set with prescribed final direction in $\mathbb{R}^3$, Bae et al. [2, Theorem 3 of the PDF v3, Theorem 8 of the HTML version] give the possible boundary extremals (pieces CCC, CSC, their subsegments, or helicoidal arcs) but leave the explicit sections, the counterpart of [14], for future work, so the antipodal-section inequalities used here are not there; and Sussmann's classification of the shortest paths in space, a helicoidal arc or three pieces each a circle or a segment [17, Theorem 1], does not say when a helicoidal arc can be excluded, even for coplanar end data such as those of a U-turn. In space the planar paths need not be optimal: for displacements with equal end tangents a curve leaving the plane is shorter [11].

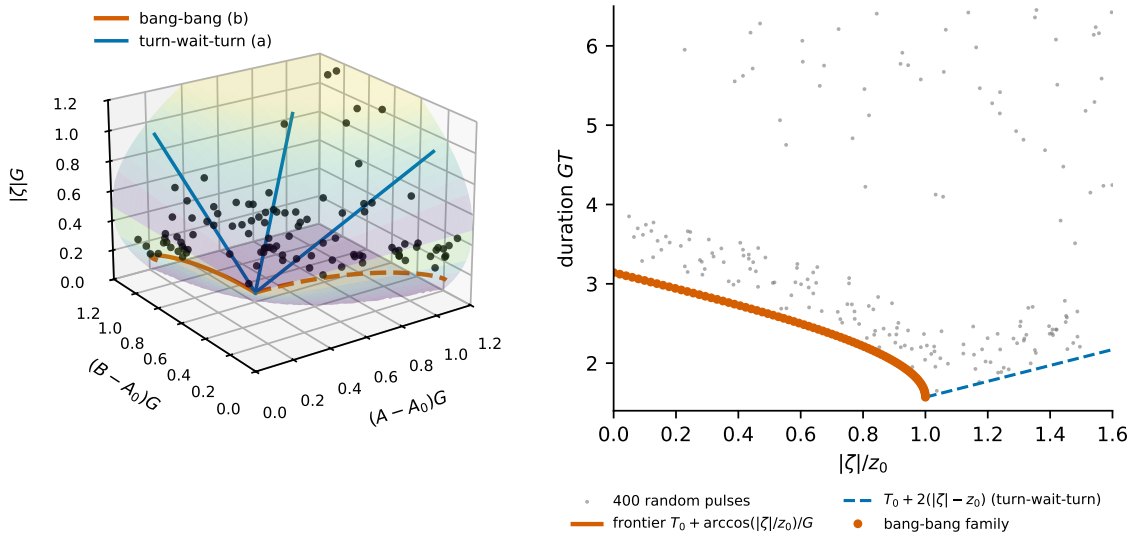


Figure 1. Theorem 1 as a picture. Left: in the coordinates $(A-A_0, B-A_0, |\zeta|)$ (units of $1/G$) every lossless complete transfer lies between the two sheets $|\zeta| = z_0 \pm \sqrt{(A-A_0)(B-A_0)}$ (translucent); dots: 400 random piecewise-constant pulses with complex coupling $|g| \le G$ and detuning (those inside the box shown), none of which violates (a) or (b) (smallest margins $5\cdot10^{-4}$ and $2\cdot10^{-2}$). Orange: the bang–bang family, equality in (b); blue: turn–wait–turn pulses, equality in (a) for $|\zeta| \ge z_0$. Right: duration against $|\zeta|/z_0$, with the frontier $T_0 + \arccos(|\zeta|/z_0)/G$ of (b) and the turn–wait–turn line beyond z_0.

Corollary 2 written proof, audited. *A complete transfer with $\zeta = 0$ satisfies*

$$\gamma_q(A-A_0) + \gamma_m(B-A_0) \;\ge\; \frac{1}{G}\max\Big(\sqrt{\gamma_q\gamma_m},\; \frac{\pi}{2}\min(\gamma_q,\gamma_m)\Big).$$

With relaxation only (no dephasing), to first order in the rates the write channel loses $(\gamma_q A + \gamma_m B)/3$, and the square pulse, which does not cancel ζ, reaches the floor $(\gamma_q + \gamma_m)A_0/3$; so ***cancelling the leakage costs at least a third of the right-hand side above that floor, whatever the pulse.*** *For equal rates γ the bound is $\pi\gamma/(2G)$ and is attained by the bang–bang pulse $(3\pi/(4G), \pi/(4G))$; the branch $\sqrt{\gamma_q\gamma_m}/G$ is not attained by pulses of fixed phase and no detuning (see above).*

Proof. By (a) with $\zeta = 0$, $(A-A_0)(B-A_0) \ge 1/(4G^2)$, and $\gamma_q u + \gamma_m v \ge 2\sqrt{\gamma_q\gamma_m\, uv}$. By (b) with $\zeta = 0$, $(A-A_0) + (B-A_0) = T - 2A_0 \ge \pi/(2G)$, and $\gamma_q u + \gamma_m v \ge \min(\gamma_q,\gamma_m)(u+v)$. The bang–bang pulse has $T = \pi/G$. At first order $1-p = \gamma_q A + \gamma_m B$, since $A = \int |a|^2$ and $B = \int |b|^2$ along the lossless transfer, and $1-c = \frac{1}{2}(\gamma_q A + \gamma_m B)$, so $r = (3-p-2c)/6 = (\gamma_q A + \gamma_m B)/3$. □

2.2. The exact cost of cancelling the leakage. For pulses of fixed phase and no detuning the cost of cancelling the leakage can be computed exactly. With $\zeta = 0$ the curve $\mathbf{r}$ is a planar U-turn back to its own axis, and by Dubins' theorem [6] the shortest one with axial displacement s is a path of type CSC or CCC. In units of z_0, with $q = \arcsin(\sqrt{4+s^2}/4)$

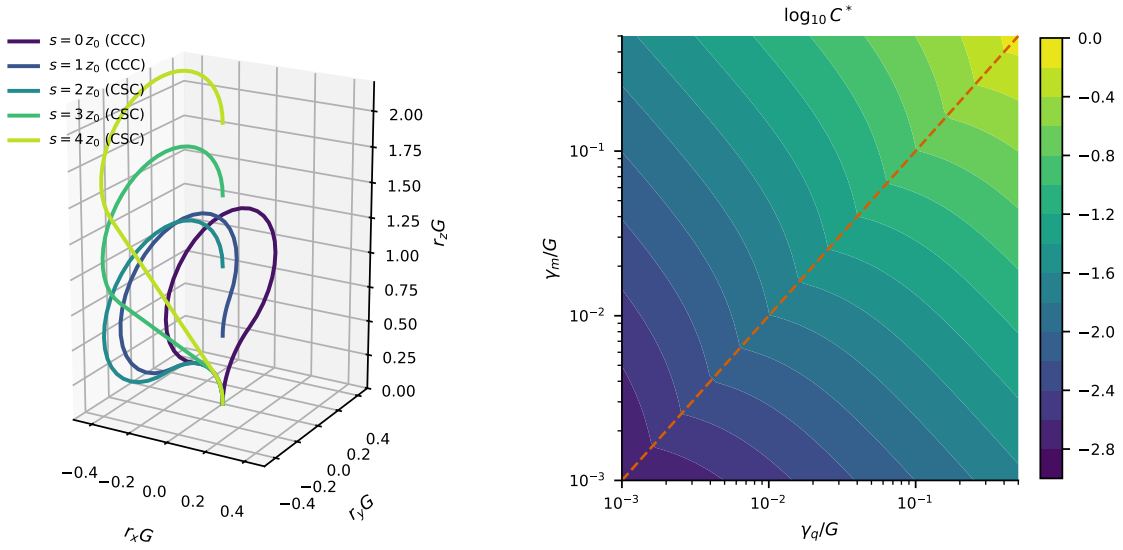


Figure 2. Proposition 3. Left: the curve $\mathbf{r}(t) = \int \mathbf{n}$ of the shortest planar U-turns with $\zeta = 0$ for axial displacement $s = 0, 1, \ldots, 4\, z_0$ (CCC for $s < 2z_0$, CSC beyond); each ends on the axis at height s. Right: the minimum cost $C^* = \frac{\gamma_q+\gamma_m}{2G} \arccos \frac{|\gamma_q-\gamma_m|}{\gamma_q+\gamma_m}$ for pulses of fixed phase and no detuning (colours, $\log_{10} C^*$, with C^* the dimensionless cost of Proposition 3; dashed: equal rates). Its ratio to the bound of Corollary 2 for every pulse depends only on γ_m/γ_q: it is 1 for equal rates and has maximum $(\pi/2 + 2/\pi) \arctan(2/\pi) \approx 1.25140958$, attained at $\gamma_m/\gamma_q = 4/\pi^2$ and $\pi^2/4$, where the two branches of that bound meet.

and $s \geq 0$ by symmetry, the triangle of centres (sides 2, 2, $\sqrt{4+s^2}$) and the condition on the final heading give the eight candidates: LSL and RSR of length $3\pi + \sqrt{4+s^2}$; LSR and RSL, for $s \geq 2$, of length $2\pi - 2\arctan(\sqrt{s^2-4}/2) + \sqrt{s^2-4}$; and, for $s^2 \leq 12$, two pairs of CCC paths of lengths $3\pi \mp 4q$ when $s < 2$, and $5\pi - 4q$ and $\pi + 4q$ when $s > 2$ (each checked against the integrated path for $s \neq 2$). At $s = 2$ the end arcs degenerate and the shortest CCC representatives have length 2π; this end point is treated separately from the open branches. For $s < 2$ the shortest is $3\pi - 4q < 3\pi$. For $s > 2$, LSR has length at most $2\pi + \sqrt{s^2-4} < 3\pi$ when $s^2 \leq 12$, which excludes LSL and $5\pi - 4q$; and $f(s) = 4q - \pi + 2\arctan(\sqrt{s^2-4}/2) - \sqrt{s^2-4}$ tends to 0 as $s \to 2$ with $f'(s) = 4s/\sqrt{(12-s^2)(4+s^2)} - \sqrt{s^2-4}/s > 0$, since the first term is at least 1 exactly when $s^2 \geq 4$; so $\pi + 4q \geq LSR$. Beyond $s^2 = 12$ only the CSC paths exist. Hence CCC for $|s| \leq 2z_0$ and CSC beyond, and the length is the least duration whose section at heading π in [14] contains the point $(s, 0)$, which they do not compute. In units of z_0 it is

$$L(s) = \begin{cases} 3\pi - 4\arcsin\dfrac{\sqrt{4+s^2}}{4}, & |s| \leq 2, \\ 2\pi - 2\arctan\dfrac{\sqrt{s^2-4}}{2} + \sqrt{s^2-4}, & |s| \geq 2. \end{cases}$$

Proposition 3 written proof, audited. *Among complete transfers with $\zeta = 0$ by pulses $g = e^{i\varphi}u(t)$, u real, φ constant, without detuning,*

$$\min[\gamma_q(A - A_0) + \gamma_m(B - A_0)] = \frac{\gamma_q + \gamma_m}{2G} \arccos \frac{|\gamma_q - \gamma_m|}{\gamma_q + \gamma_m},$$

attained for $\gamma_q \neq \gamma_m$, both positive, by a turn of $\pi + \phi$, a wait without coupling, and a turn of ϕ back, with $|A - B| = (\gamma_q + \gamma_m)/(2G\sqrt{\gamma_q\gamma_m})$, and for equal positive rates by the bang–bang pulse. If exactly one rate vanishes the value 0 is an infimum, approached only by ever longer waits and not attained; if both vanish the cost is identically zero.

Proof. With $A - B = s$ and $A + B \geq L(s)$, the cost is at least $f(s) = \frac{1}{2}[(\gamma_q + \gamma_m)L(s) + (\gamma_q - \gamma_m)s] - (\gamma_q + \gamma_m)A_0$ (in units $z_0 = 1$). L is continuous, and differentiable except at $|s| = 2$: at $s = 2$ its slope is -1 from the left and 0 from the right, and symmetrically at $s = -2$. On $|s| > 2$, $L'(s) = \operatorname{sgn}(s)\sqrt{s^2 - 4}/|s|$, so f is convex there with its minimum where $\sqrt{s^2 - 4}/|s| = |\gamma_q - \gamma_m|/(\gamma_q + \gamma_m)$, i.e. $|s| = (\gamma_q + \gamma_m)/\sqrt{\gamma_q\gamma_m}$; at that point the algebraic terms cancel and $f = (\gamma_q + \gamma_m)\arctan(2/\sqrt{s^2 - 4})$, which is the stated arccosine. On $|s| < 2$, $|L'(s)| = 4|s|/\sqrt{(12 - s^2)(4 + s^2)} \leq 1$ (equivalent to $s^2 \leq 4$), so on the favourable side, where $(\gamma_q - \gamma_m)s \leq 0$, f decreases towards $|s| = 2$; and since L is even, $f(s) - f(-s) = (\gamma_q - \gamma_m)s$, so the other side is no better. □

For equal rates this is $\pi\gamma/(2G)$, the value of Corollary 2, which holds for every pulse. For unequal rates we do not know whether a coupling of varying phase, or a detuning, can do better. The question is whether, with $\zeta = 0$, a non-planar path, for instance one with a helicoidal arc in Sussmann's sense [17], can make a shorter U-turn than $L(s)$. Two numerical searches over curves that leave the plane (24 rotations of free axis with 30 starts, and a direct minimisation of the length at fixed s for $0 \leq s \leq 4$) found no improvement beyond their tolerances: the first accepts end-state and leakage residuals up to 10^{-4}, compares costs to 10^{-3}, and for equal rates lands slightly below the formula within them; the second falls below $L(s)$ by the same amount, $3 \cdot 10^{-3}$ or less, as at the control $s = 2$, where the planar path is proved optimal. This is evidence, not a proof; and in space planar paths are not always the shortest [11]. Mittenhuber's mechanism, two planar pieces turned about the common tangent at their junction, does not apply here. The data of the U-turn lie on the axis, so pieces in distinct planes through the axis can only meet on the axis with tangent $\pm e_z$. Each piece can then be turned about the axis into one common plane without changing length, curvature or end data: a curve made of such pieces is no shorter than $L(s)$. The unfolding is machine-checked (formalised) for a finite partition in which, on each piece, both the position $\mathbf{r}$ and the tangent $\mathbf{n}$ lie in one plane containing the axis, with the junction conditions, the final displacement on the axis and the pointwise differentiability hypotheses of the curve; that only $\mathbf{n}$ be meridian on each piece is not enough, since turning the pieces could destroy a transverse cancellation that occurs only at the end. The comparison with $L(s)$ then follows from the written planar Dubins argument. Pieces in planes that do not contain the axis, and helicoidal arcs, are not covered.

Remark 4. Two representative pulses cancel ζ in the least time π/G (controls related to them by a constant phase of g, or by a diagonal gauge transformation with a compensating detuning, have the same curve and exposures): the bang–bang pulse ($+G$ during $3\pi/(4G)$, $-G$ during $\pi/(4G)$), with $A = (\pi - 1)/(2G)$ and $B = (\pi + 1)/(2G)$, and its time reverse,

with the two values exchanged. Both exceed A_0 in each exposure; the better one, when the qubit is the noisier element, spends $(\pi-1)/(2G)$ in the qubit, more than the square pulse, and so worsens the dephasing of the superposition with the vacuum. **Robustness of the population is not the same as robustness of the write channel.**

2.3. Slow noise at second order. The exposures also give the effect of slow noise on the write channel at second order, for every pulse. Let the offsets ξ_q, ξ_m (static shifts of the qubit and memory levels, as in [10, Proposition 14]) be centred, with σ_q^2, σ_m^2 their second moments.

Proposition 5 written proof, audited. *For every complete transfer without losses, with $|g| \le G$ and any detuning, to second order in the offsets*

$$r = \frac{2\,\mathbb{E}(\xi_q-\xi_m)^2\,|\zeta|^2 + \mathbb{E}(\xi_q A+\xi_m B)^2}{6}.$$

(i) *For the square pulse this is the bound of* [10, Proposition 14(a)], $(8+\pi^2)(\sigma_q^2+\sigma_m^2)/(96G^2)$ *for uncorrelated offsets: that bound is exact at this order.*

(ii) *With noise on the qubit only, $r \ge \pi^2\sigma_q^2/(96G^2)$ for every pulse. The square pulse is* 1.81 *times this floor; the better of the two pulses of least duration that cancel ζ gives $(\pi-1)^2\sigma_q^2/(24G^2)$,* 2.66 % *above the square pulse. The floor is the infimum of the second-order coefficient: zero-leakage pulses that spend their extra time in the memory approach it only as $T\to\infty$, outside the regime $\sigma_q T \ll 1$ where the expansion holds (with memory loss, see Proposition 8).*

(iii) *With independent offsets on both elements and $\varrho = \sigma_m^2/\sigma_q^2$: for $\varrho \le 0.1264$ the bang–bang pulse of Theorem 1(b) with $d = 0.78$ does better than the square pulse (by* 3.3 % *at $\varrho = 0.1$ with $d = 0.848$), and for $\varrho \ge 7.91$ its time reverse does (certified by interval arithmetic, directly at* 7.91*). The square pulse is optimal at this order among all pulses for $0.137 \le \varrho \le 7.30$ (the upper end certified through* 10/73*, whose reciprocal is* 7.3*) (certified by interval arithmetic over the region allowed by Theorem 1; a floating-point minimisation over that relaxed region suggests that this certificate is lost near* 0.1357*, which is not established as a physical threshold). (The problem is symmetric: F/ϱ is the same problem with $1/\varrho$ and A, B exchanged.)*

(iv) *For pulses of fixed phase and no detuning the set of ratios at which the square pulse is optimal at this order is a closed interval $[\varrho_c, 1/\varrho_c]$, and ϱ_c is the crossing of the bang–bang family:*

$$\varrho_c = \varrho_{\mathrm{bb}} := \inf\{\varrho : \min_d [F_{\mathrm{bb}}(\varrho,d) - F_{\mathrm{sq}}(\varrho)] \ge 0\} = 0.12645352248540975\ldots$$

(enclosed in an interval of width 10^{-12}; certified by interval arithmetic, using the planar exposure bound of Lemma 6). Below ϱ_{bb} the bang–bang pulse with d near $d_* = 0.7717755295532\ldots$ is strictly better, and its time reverse for the reciprocal ratios. For pulses of varying phase or with detuning, $0.1264 < \varrho < 0.137$ and $7.30 < \varrho < 7.91$ remain undecided.

Proof. A static shift changes $U(T)$ into $U(T)(I - i\Xi_1 - \Xi_2 + \ldots)$, with $\Xi_1 = \xi_q K + \xi_m(T-K)$, Ξ_2 quadratic in the offsets and a remainder of third order, bounded uniformly for small offsets. With $U(T)|a\rangle = e^{i\phi}|b\rangle$ the transferred amplitude is $t = e^{i\phi}(1 - i(\xi_q A + \xi_m B) - (\Xi_2)_{aa} + \ldots)$ and the leaked one is $(\xi_q - \xi_m)\zeta$ to first order; unitarity, $|t|^2 = 1 - (\xi_q-\xi_m)^2|\zeta|^2$, fixes $2\,\mathrm{Re}(\Xi_2)_{aa} =$

$(\xi_q-\xi_m)^2|\zeta|^2+(\xi_qA+\xi_mB)^2$. For each fixed pulse and duration, write $\xi=\epsilon X$ with X of fixed law and finite second moments: the averaged formula below then holds with a remainder $o(\epsilon^2)$; a cubic averaged remainder needs a further moment assumption. Averaging, the first-order terms vanish, $p=1-\mathbb{E}(\xi_q-\xi_m)^2|\zeta|^2$ and $c=|\mathbb{E}t|=1-\frac{1}{2}\mathbb{E}\big[(\xi_q-\xi_m)^2|\zeta|^2+(\xi_qA+\xi_mB)^2\big]$, and $r=(3-p-2c)/6$ gives the formula. (i) is $|\zeta|=z_0$, $A=B=A_0$. (ii): $A\geq A_0$ by Theorem 1(a); the minimum-time pulses have $A=(\pi\mp1)/(2G)$; and the zero-leakage pulses $C_{\pi+\phi}SC_\phi$ (a turn of $\pi+\phi$, a wait, a turn of ϕ back) have $A-A_0=O(1/(G^2T))$. (iii): the lower bound minimises $2(\sigma_q^2+\sigma_m^2)|\zeta|^2+\sigma_q^2A^2+\sigma_m^2B^2$ over the region allowed by Theorem 1, $A,B\geq A_0$, (a), (b) and $T\geq T_0+2(|\zeta|-z_0)$. The difference g between a point of this region and the square pulse is linear in ϱ, so it suffices to show $g\geq0$ at $\varrho=0.137$ and at $\varrho=1$; the region is symmetric in A and B, and dividing by ϱ gives $1\leq\varrho\leq1/0.137$. For $|\zeta|\geq z_0$ every term of g is non-negative; g increases with $A-A_0$ and $B-A_0$, so the latter can be taken at its least admissible value; near $|\zeta|=z_0$, $\arccos(|\zeta|/z_0)\geq\sqrt{2(1-|\zeta|/z_0)}$ gives $g\geq0$ by a one-line estimate; the rest of the region, a rectangle in $(|\zeta|,A-A_0)$, is covered by 551 accepted leaf boxes (1101 nodes of the subdivision evaluated; 554 leaves at $\varrho=10/73$) on each of which an interval lower bound of g is non-negative (at $\varrho=0.137$ computed twice, with mpmath and with Arb in SageMath; at the exact rational $\varrho=10/73$, which gives the decimal endpoint 7.30, with the Python certificate). The other side evaluates one bang–bang pulse, $d=0.78$: its difference with the square pulse is linear in ϱ and negative at $\varrho=0$ and $\varrho=0.1264$ (enclosed by intervals), and its time reverse covers $\varrho\geq1/0.1264\approx7.9114$; at $\varrho=7.91$ the reversed pulse is checked directly (difference $-7.6\cdot10^{-4}$, slope -0.74, both enclosed). The analytic patch and the subdivided region share the same exactly represented endpoint, rational parameters are enclosed directly, and every coefficient used as a lower bound is rounded downwards.

(iv): for every pulse P the difference $g_P(\varrho)$ with the square pulse is affine in ϱ, so the set $\bigcap_P\{\varrho>0:g_P(\varrho)\geq0\}$ is a closed interval; it contains 1 and is invariant under $\varrho\mapsto1/\varrho$, hence it is $[\varrho_c,1/\varrho_c]$. For pulses of fixed phase and no detuning the curve $\mathbf{r}$ is planar, and the region of Theorem 1 can be replaced by the bound of Lemma 6: for $|\zeta|<z_0$ and $A-A_0$ below its bang–bang value, $B-A_0$ is at least that of the family; symmetrically with A and B exchanged. Since F grows with A and with B, the minimum over that branch lies on the family, a two-parameter set covered by intervals (605 accepted leaves, in the variables (ϕ,w), which avoid the singularity at $\phi\to0$). The rest is covered by the relaxed region (22 322 leaves in all at $\varrho=0.12645395$). By linearity in ϱ it suffices to add $\varrho=1$, and the symmetry gives the upper end: $\varrho_c\leq0.12645395$. At 0.126453 the bang–bang witness $d=0.7717755$ is strictly better (difference $-3.0\cdot10^{-6}$, enclosed), and its difference increases with ϱ, so it is better on all of $[0,0.126453]$ and $\varrho_c>0.126453$; `certificar_umbral_plano.py`. The identification $\varrho_c=\varrho_{\mathrm{bb}}$ is a second certificate (`certificar_rho_c_igual_bb.py`). A Krawczyk operator in (ϱ,d) proves that $\{h=0,\ \partial_dh=0\}$, $h=F_{\mathrm{bb}}-F_{\mathrm{sq}}$, has a unique zero in a box, with enclosure widths below $6\cdot10^{-41}$ for ϱ_{bb} and $4\cdot10^{-40}$ for d_* (regular Jacobian: $\partial_\varrho h=5.80$, $\partial_d^2h=2.54$). Since $h(\varrho,0)=0$ always, at the crossing the global minimum of h is zero and is attained both at $d=0$ and at the unique nonzero minimiser $d=d_*$; throughout the enclosing interval E of ϱ, the critical branch $d(\varrho)$ is the unique critical point in $U=[d_*\pm0.01]$, and above the crossing it need not be the global minimiser of h. On $N=U\times[0,0.02]$ in the family variables (ϕ,w) the Hessian of g is positive definite (by intervals), uniformly for $\varrho\in E$. At $w=0$ the exact identities $g_\phi=h_d$ and $g_w=h_d/2$ make the full gradient vanish at $(d(\varrho),0)$; since N is convex, Taylor's formula gives $g\geq h(\varrho,d(\varrho))$ on N, including its edge $w=0$. Outside N, the analytic near-square patch and the large-wait tail are treated separately; on

the remaining retained branch boxes the interval lower bound is at least $5 \cdot 10^{-7}$ at both endpoints of E, hence throughout E by affinity (equality at the square pulse is allowed in the analytic patch). The reflected branch is certified independently. For the middle region put $x_0 = d - \sin d$ and $y_0 = d + \sin d$ (in the units of Lemma 6): if $x_0 \le a \le y_0$, monotonicity in b and Theorem 1(b) reduce the cost to $g(a, 2d-a)$, which increases in a for $\varrho \le \frac{1}{4}$; if $a \ge y_0$, then $g(a,b) \ge g(y_0, x_0) \ge g(x_0, y_0)$ because $\varrho \le 1$; so the middle region is bounded below by $h(\varrho, d)$ without subdivision. Hence the square pulse is optimal at ϱ_{bb}; since $h(\varrho_{\mathrm{bb}}, d_*) = 0$ and $\partial_\varrho h > 0$ at every positive d, the same nonzero bang–bang pulse is strictly better for every $\varrho < \varrho_{\mathrm{bb}}$. Controls (a neighbourhood too small, a ratio 10^{-4} lower, an off-centre Krawczyk box) fail. □

What was known. That a quasi-static shift of a level which the protocol must populate cannot be cancelled to first order is known. For the controlled-phase gate of Rydberg atoms with a bounded Rabi frequency, Jandura, Thompson and Pupillo [7] show that the first-order error is a phase equal to the detuning times the time spent in the Rydberg state, which is positive for every pulse, and that at second order the infidelity splits into a term fixed by those times and a term from the population left behind, which a longer composite pulse removes while the first stays; they minimise the first term numerically. Bergonzoni, Jandura and Pupillo [4] state the obstruction as general criteria, without bounds on the controls and without values. In the write of a memory the same two terms are those of the formula above, in A^2 and in $|\zeta|^2$; what (ii) adds is the closed floor $\pi^2\sigma_q^2/(96G^2)$ for every pulse with $|g| \le G$, through $A \ge A_0$.

2.4. Fixed phase and no detuning: the planar exposure bound.

Lemma 6 planar exposure bound written proof, audited. *For a pulse of fixed phase and no detuning, put* $a = 2G(A - A_0)$, $b = 2G(B - A_0)$, $z = |\zeta|/z_0 \in [0,1)$ *and* $d = \arccos z$, *and for* $0 < \phi \le d$

$$a_z(\phi) = \phi - (1+z)\tan\tfrac{\phi}{2}, \qquad b_z(\phi) = \phi + (1-z)\cot\tfrac{\phi}{2}.$$

Then $a_z' = (\cos\phi - z)/(1+\cos\phi) \ge 0$, $b_z' = (z - \cos\phi)/(1-\cos\phi) \le 0$, a_z *maps* $(0, d]$ *onto* $(0, d - \sin d]$, *and every such pulse with* $0 < a \le d - \sin d$ *has*

$$b \;\ge\; b_z(a_z^{-1}(a)),$$

and symmetrically with a *and* b *exchanged. This is the type-*$(1, 0, -1)$ *family of the remark below, written at fixed* $|\zeta|$.

Proof. With fixed phase and no detuning the Bloch direction moves on one meridian, $\mathbf{n} = (\sin\alpha, 0, \cos\alpha)$ with $|\dot\alpha| \le 2G$, and $\mathbf{r}$ is the path of a planar Dubins car of heading α; A and B are the times weighted by $\cos^2(\alpha/2)$ and $\sin^2(\alpha/2)$; the axial displacement is $A - B$, and the magnitude of the lateral displacement is $2|\zeta|$. After normalisation, $t_f = \pi + a + b$, $x = a - b$ and $y = \pm 2z$. The write ends at the south pole, so the final heading is an odd multiple of π. (1) Final headings $\pm(2k+1)\pi$ with $k \ge 1$: for $f \ge 0$ and an absolutely continuous angle, $2G\int_0^T f(\alpha)\,dt \ge \int_0^T f(\alpha)|\dot\alpha|\,dt \ge |\int_0^T f(\alpha)\dot\alpha\,dt|$, so with the primitives of $\cos^2(\alpha/2)$ and $\sin^2(\alpha/2)$, $A, B \ge A_0 + k\pi/(2G)$, i.e. $a, b \ge k\pi$, which exceeds the bang–bang values $d \mp \sin d$ at the same $|\zeta|$ in both exposures; such pulses are dominated by the bang–bang pulse (F grows with A and B). Equality $a = b = k\pi$ forces a full-speed monotone turn and gives $|\zeta| = z_0$, outside the sector $z < 1$; the square pulse has the same leakage and smaller exposures. (2) Heading π (by reflecting the control if necessary). Set $\theta = (t_f - \pi)/2 = (a+b)/2$; Theorem 1(b) gives $\theta \ge d$. At heading π the classification of [14, Lemma 3.1 and Theorem 11.1] gives their

Case 1 for $\pi < t_f < 3\pi$ and Case 3 for $t_f \geq 3\pi$; the non-simply connected Case 2 does not occur. The lower boundary on the negative-x side is the $U_3 = (1, 0, -1)$ arc

$$x_3 = -2\sin\phi - 2(\theta - \phi)\cos\phi, \qquad y_3 = 2\cos\phi - 2(\theta - \phi)\sin\phi.$$

In Case 1 the U_6 branch satisfies $a \geq \theta - \sin\theta \geq d - \sin d$ and does not enter the sector under consideration; in Case 3 it contributes no additional boundary arc. U_1 has $y \geq 2$, and U_2 lies on the opposite axial side. The retained U_3 arc includes $0 \leq \phi \leq d$: in Case 1 this follows from $\theta \geq d$, and in Case 3 its first intersection with its reflected arc occurs after $\phi = \pi/2$. Since x_3 is increasing on the relevant arc, every reachable point in the strict sector $a < d - \sin d$ lies above a unique U_3 point at the same x, with $0 < \phi < d$. (3) Both lateral signs are covered directly: $y = \pm 2z \leq 2z$ and $y \geq y_3$ imply

$$\theta \;\geq\; \theta_z(\phi) := \phi + \frac{\cos\phi - z}{\sin\phi},$$

and consequently $a \geq a_z(\phi)$ (along the arc the wait is $w = (\cos\phi - z)/(G\sin\phi)$ and the exposures are the closed forms $a_z(\phi)$, $b_z(\phi)$ of the remark, simplified). (4) Let $\beta = a_z^{-1}(a)$. The strict monotonicity of a_z gives $\phi \leq \beta$, while $\theta_z'(\phi) = \cos\phi\,(z - \cos\phi)/\sin^2\phi \leq 0$ for $0 < \phi \leq d$; therefore $b = 2\theta - a \geq 2\theta_z(\beta) - a = b_z(\beta)$. At $a = d - \sin d$ the same inequality follows from $\theta \geq d$, and $a = 0$ is excluded by the disk bound for $z < 1$. Equality is attained by U_3 with the stated non-negative wait. Finally, time reversal combined with $\alpha \mapsto \pi - \alpha$ exchanges a and b, giving the symmetric assertion. Steps (1), (3) and (4) are elementary; the passage from the sections of [14], which are stated at fixed duration and unwound final heading, to exposures is our written derivation, not a lemma of that source. The closed forms were checked against piecewise pulses with exact integrals and against the parametrisation of the remark (to $2 \cdot 10^{-13}$, `verificar_v289.py`), and the inequality with both lateral signs on 7061 random turn–wait–turn pulses in the sector (`verificar_v291.py`). □

Remark 7 (the exact planar section). For a pulse of fixed phase and no detuning, the duration, ζ and the axial displacement $A - B$ are the time, the lateral position and the axial position of a planar Dubins car of radius z_0 that turns back to heading π. The set of such end points at a given duration is the section of the reachable set at heading π, described exactly by Patsko and Fedotov [14, Lemma 3.1 and Theorem 11.1]: its boundary is made of arcs of extremal controls of four types. In exposures, the part with $|\zeta| < z_0$ and the smallest $A - A_0$ is the arc of their type $(1, 0, -1)$: a turn of $\pi + \phi$ at full coupling, a wait w, and a turn of ϕ back,

$$|\zeta| = z_0\big[\cos\phi - Gw\sin\phi\big],$$

$$A - A_0 = \tfrac{1}{2G}\big[\phi - \sin\phi + Gw(1 - \cos\phi)\big], \qquad B - A_0 = \tfrac{1}{2G}\big[\phi + \sin\phi + Gw(1 + \cos\phi)\big].$$

For prescribed $z = |\zeta|/z_0 \in [0, 1)$ the branch is parametrised by $0 < \phi \leq \arccos z$ and $w = (\cos\phi - z)/(G\sin\phi) \geq 0$. At $\phi = \arccos z$ ($w = 0$) it is the bang–bang pulse of Theorem 1(b); at fixed z, letting $\phi \downarrow 0$ sends $w \to \infty$ and gives $(A - A_0)(B - A_0) \to (z_0 - |\zeta|)^2$, so the disk of Theorem 1(a) is approached along the branch only in that limit. This explains why the relaxation in the proof of (iii) stops at 0.137: its minimum lies where the disk and the frontier are both active, a point no pulse reaches. For $0 \leq |\zeta|/z_0 < 1$ and $0 < \phi < \arccos(|\zeta|/z_0)$ the family satisfies both bounds of Theorem 1 strictly, $(A - A_0)(B - A_0) > (z_0 - |\zeta|)^2$ and $(A - A_0) + (B - A_0) > 2d/(2G)$ with $d = \arccos(|\zeta|/z_0)$, A increases and B decreases with ϕ, and the product tends to $(z_0 - |\zeta|)^2$; these facts are formalised from the closed forms above (the chain from the control of type $(1, 0, -1)$ to these forms is the written computation).

We conjecture that the same holds for every pulse, of varying phase and with detuning: a numerical search over three-dimensional pulses with exact integrals, aimed at going below

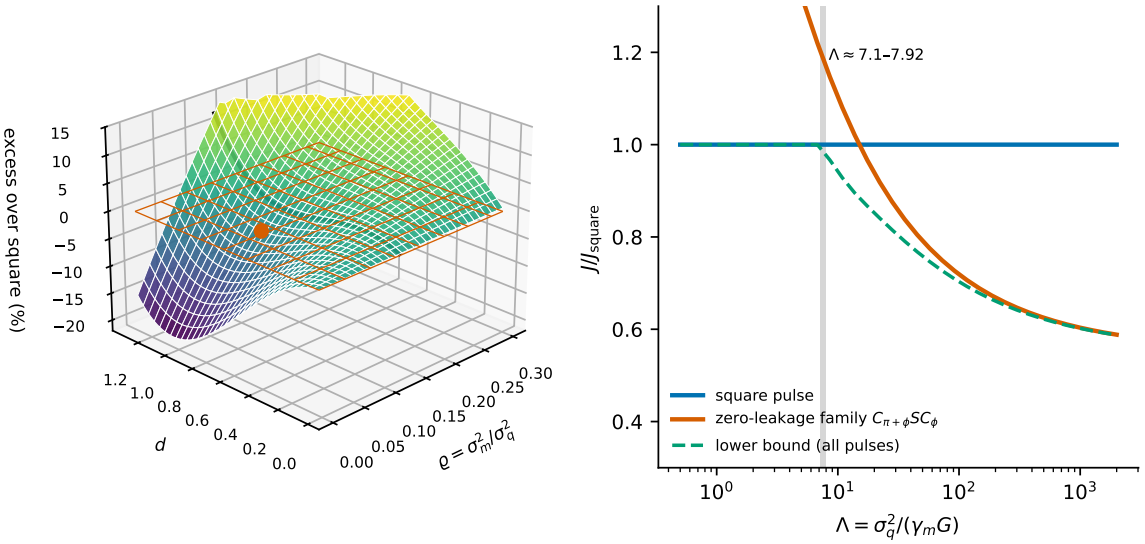


Figure 3. Propositions 5 and 8. Left: relative excess $(F - F_{\mathrm{sq}})/F_{\mathrm{sq}}$ of the bang–bang pulse of Theorem 1(b) with parameter d over the square pulse, with $F = 2(1+\varrho)|\zeta|^2 + A^2 + \varrho B^2$ (negative: better than the square pulse; orange plane: zero; dot: $-3.3\,\%$ at $\varrho = 0.1$, $d = 0.848$). Right: leading-order error J with a lossy memory, relative to the square pulse, against $\Lambda = \sigma_q^2/(\gamma_m G)$, for the zero-leakage pulses $C_{\pi+\phi}SC_\phi$ and for the lower bound of the proof of Proposition 8, valid for every pulse: the square pulse is optimal up to $\Lambda = 7.1$ and not beyond 7.92 (grey band; for pulses of fixed phase and no detuning the exact threshold is $\Lambda_c = 7.9079$, Proposition 8(iii)); the explicit family beats it from $\Lambda \approx 15$.

the family, lands on it in five values of $|\zeta|$ and finds nothing below it beyond its constraint tolerance: over 40 starts per value the lowest margin is $-8.7 \cdot 10^{-8}$, at $|\zeta|/z_0 = 0.664$ and $2G(A - A_0) = 0.019$, where the constraints are met to 10^{-8} and the frontier is steep enough for that error to move the margin by up to about 10^{-6} (`conjetura_adversario.py`). If it holds, the same planar exposure reduction, and hence the certified value of ϱ_c (and of Λ_c in Proposition 8), applies to all pulses.

2.5. Slow noise together with memory loss. When the memory also loses its excitation, waiting in it is no longer free. To leading order in both effects the error of a transfer with slow noise on the qubit and a memory decay rate γ_m is $J = \sigma_q^2(2|\zeta|^2 + A^2)/6 + \gamma_m B/3$.

Proposition 8 written proof, audited. *Let the offset act on the qubit only, centred with variance σ_q^2; let the memory decay at rate γ_m, with $\gamma_q = 0$ and no dephasing; let A, B, ζ be those of the lossless pulse, with $|g| \le G$ and any detuning, and J the leading-order error above. Put $\Lambda = \sigma_q^2/(\gamma_m G)$.*

(i) *The square pulse minimises J over all pulses for $\Lambda \le 7.1$, and for $\Lambda \ge 7.92$ it does not (both certified by interval arithmetic; a floating-point minimisation over the relaxed region suggests that the certificate is lost near 7.15, which is not established as a physical threshold).*

(ii) *As* $\Lambda \to \infty$,

$$\min J = \frac{\pi^2 \sigma_q^2}{96 G^2} + \frac{\sigma_q \sqrt{\pi \gamma_m}}{6\, G^{3/2}} + O\Big(\frac{\gamma_m}{G}\Big),$$

and the zero-leakage pulse $C_{\pi+\phi} S C_\phi$ *whose wait lasts about* $\sigma_q \sqrt{\pi}/(4 G^{3/2} \sqrt{\gamma_m})$ *attains the first two terms.*

(iii) *For pulses of fixed phase and no detuning the square pulse minimises* J *exactly for* $\Lambda \le \Lambda_c$*, and* Λ_c *is the crossing of the general turn–wait–turn family* $C_{\pi+\phi} S C_\phi$*, with* ϕ *and* w *free in the admissible domain:*

$$\Lambda_c = \sup\{\Lambda : \min_{\phi, w} g \ge 0\} = 7.9078851606875\ldots$$

(enclosed in an interval of width 10^{-12}*; certified by interval arithmetic, using Lemma 6); above it the member with* $\phi_* = 0.742069\ldots$ *and a short wait* $w_* = 0.257473\ldots$ *(units* $G = 1/2$*) beats it; this critical member has* $|\zeta|/z_0 \approx 0.6500698$*, so it is not a zero-leakage pulse. For pulses of varying phase or with detuning, optimality remains undecided for* $7.1 < \Lambda \le \Lambda_c$*; for every* $\Lambda > \Lambda_c$ *the explicit planar competitor already beats the square pulse. Moreover, for every* $\Lambda > \Lambda_c$ *the minimum of* J *over pulses of fixed phase and no detuning is attained, in* $(A, B, |\zeta|)$*, by a member of the family* $C_{\pi+\phi} S C_\phi$*.*

Proof. Lower bound: with $x = A - A_0$, $y = B - A_0$, (a) gives $xy \ge (z_0 - |\zeta|)^2$ when $|\zeta| \le z_0$, so $J \ge \sigma_q^2 A_0^2/6 + \gamma_m A_0/3 + \sigma_q^2 |\zeta|^2/3 + \sigma_q^2 A_0 x/3 + \gamma_m y/3$; minimising $\sigma_q^2 A_0 x/3 + \gamma_m y/3$ over $xy \ge (z_0 - |\zeta|)^2$ and then over $|\zeta|$ (optimum $|\zeta| = \sqrt{A_0 \gamma_m}/\sigma_q$) gives $J \ge \pi^2 \sigma_q^2/(96 G^2) + \frac{2}{3} z_0 \sigma_q \sqrt{A_0 \gamma_m}$, the term $\gamma_m A_0/3$ of the memory cancelling the one of the minimisation in $|\zeta|$; this is (ii) from below; when $|\zeta| > z_0$, $J \ge \sigma_q^2 (A_0^2 + 2 z_0^2)/6$, which is larger when Λ is large. Upper bound: the pulse $C_{\pi+\phi} S C_\phi$ with axial displacement $|s|$ has $A - A_0 = z_0^2/|s| + O(|s|^{-2})$ and $B = |s| + O(1)$, and the best $|s|$ is $z_0 \sigma_q \sqrt{A_0/\gamma_m}$. (i): the difference between a pulse and the square pulse is linear in Λ and non-negative at $\Lambda = 0$, so it suffices to show it at $\Lambda = 7.1$, which is done as in Proposition 5(iii): an estimate near $|\zeta| = z_0$ and 518 accepted leaf boxes elsewhere (1035 nodes evaluated). The bang–bang pulse with $d = 0.85$ has a difference negative at $\Lambda = 7.92$ and decreasing in Λ.

(iii): as in Proposition 5(iv), the set of $\Lambda \ge 0$ where the square pulse is optimal is an interval $[0, \Lambda_c]$ (affinity in Λ and optimality at $\Lambda = 0$), and the relaxed region is cut by Lemma 6, with the same reading of [14]. It suffices to certify $\Lambda_{\mathrm{lo}} = 7.90788$: the family in its parameters (ϕ, w) (2877 nodes; the minimum is interior, with lower bound $6 \cdot 10^{-9}$, since Λ_{lo} is $5 \cdot 10^{-6}$ below the family crossing), the symmetric branch and the rest of the relaxed region (4663 nodes in all, `certificar_umbral_plano_perdidas.py`); the same certificate at $\Lambda = 7.90789$, 7.91 and 7.92 fails, as it should. Above the enclosure the competitor is not the bang–bang pulse but the member of the family with $\phi = 0.742069$ and wait $w = 0.257473$ (units $G = 1/2$), whose difference with the square pulse is negative at $\Lambda = 7.90789$ (enclosed) and decreasing in Λ; integrated as a piecewise pulse it closes at the south pole to $3 \cdot 10^{-16}$ and reproduces the closed form to $5 \cdot 10^{-16}$. The bang–bang pulse crosses only at 7.9176. The identification with the family crossing is a second certificate (`certificar_lambda_c_identificacion.py`): a Krawczyk operator in (Λ, ϕ, w) proves that $\{g = 0,\ g_\phi = 0,\ g_w = 0\}$ has a unique zero in a box (the Λ coordinate of width below $2.7 \cdot 10^{-39}$, every coordinate below $3.3 \cdot 10^{-39}$). On the fixed rational rectangle $N = [\phi_* \pm 0.01] \times [w_* \pm 0.02]$, centred at rational

decimal approximations of (ϕ_*, w_*), the Hessian satisfies $H > 0.3\,I$ uniformly for every Λ in the enclosing interval E, as verified by interval arithmetic, and the Krawczyk-enclosed critical point lies strictly inside N and the physical family domain; at its parameter value, strict convexity makes this point the minimum over N, with value zero. Outside N, after the analytic near-square patch and the non-physical boxes are separated, the retained exterior family boxes have an interval lower bound at least $1.88 \cdot 10^{-7}$ at both endpoints of E (the square pulse remains a zero in the analytic patch); the symmetric branch and the rest of the region are certified as before, at both endpoints and extended by affinity; controls fail. Λ_c is close to, but distinct from, $1/\varrho_c$: approximately 7.90788516068753 against 7.90804384366106; no relation between the two is asserted.

The minimiser above Λ_c (units $G = 1/2$, so $z_0 = 1$ and $A_0 = \pi/2$; with $\gamma_m > 0$ we compare the normalised objective $\hat{J} = J/\gamma_m$, which has the same minimisers, and $12(\hat{J} - \hat{J}_{\mathrm{sq}}) = \Lambda Q + 4y$ with $Q = 2(z^2 - 1) + x(2A_0 + x)$, $z = |\zeta|$). At fixed z, J is strictly increasing in x and in y. Let $d = \arccos z$ and $x_0 = d - \sin d$. If $z \geq 1$, $J \geq J_{\mathrm{sq}}$. If $z < 1$ and $x \leq x_0$, Lemma 6 gives $y \geq Y(x; z)$, the y of the member of the family with the same x and z, so J is at least its value there. If $z < 1$ and $y \leq x_0$, the symmetric statement bounds J from below by a point of the mirrored branch, with $(x, y) = (y_f, x_f)$ for a member (x_f, y_f) of the family, where $12(\hat{J} - \hat{J}_{\mathrm{sq}}) = \Lambda Q_m + 4x_f$ and $Q_m = 2(z^2 - 1) + y_f(\pi + y_f)$. With $s = \sin\phi$ and $c = \cos\phi$, the closed forms of Remark 7 give $z \leq c$ and $1 - z^2 \leq (1 - z)(1 + c) = s^2 + (1 + c)ws/2$, hence

$$Q_m \;\geq\; \pi\phi + s(\pi - 2s) + w(1 + c)\left(\tfrac{\pi}{2} - s\right) \;\geq\; \pi\phi + s(\pi - 2s) \;>\; 0,$$

so the mirrored branch never beats the square pulse, for any $\Lambda \geq 0$. If $x, y > x_0$, Theorem 1(b) gives $x + y \geq 2d$, and by monotonicity the minimum lies on $x + y = 2d$, where $\frac{d}{dx}\hat{J}(x, 2d - x) = \frac{\Lambda}{6}(A_0 + x) - \frac{1}{3} > 0$ for $\Lambda > 2/A_0$; so it is the bang–bang point $x = x_0$, a member of the family. Since $\Lambda_c > 2/A_0$, every case is bounded by the family or by the square pulse, which the family beats above Λ_c; the family minimum exists because $J \to \infty$ as $w \to \infty$. This identifies the end point of the minimiser, not the control that reaches it. Numerically, a piecewise search with 8 to 24 segments and 252 starts per value lands on the family minimum to 10^{-8} for Λ from $\Lambda_c + 0.1$ to 1000, on the frontier of Remark 7 (`optimo_sobre_lambda_c/`). □

The expansion behind J needs $\sigma_q T \ll 1$, and the optimal wait makes T grow like $\sigma_q/\sqrt{\gamma_m}$: the regime of (ii) is $\gamma_m G \ll \sigma_q^2 \ll \sqrt{\gamma_m}\, G^{3/2}$, a memory that keeps its excitation much longer than the slow noise of the qubit lets it keep its phase. Inside it the exact dynamics (losses and a Gaussian offset averaged by quadrature) confirms the gain of waiting: with $\gamma_m/G = 2 \cdot 10^{-6}$ and $\Lambda = 60$ the error falls from $2.29 \cdot 10^{-5}$ for the square pulse to $1.76 \cdot 10^{-5}$, and with $\gamma_m/G = 2 \cdot 10^{-8}$ and $\Lambda = 2000$ from $7.45 \cdot 10^{-6}$ to $4.43 \cdot 10^{-6}$, within 1.1 % of J; with $\sigma_q T \approx 5$ the square pulse is better again.

Remark 9 (relaxation of the qubit; the bound and the sufficient condition formalised). With $\gamma_q > 0$ the error is $J = \sigma_q^2(2|\zeta|^2 + A^2)/6 + (\gamma_q A + \gamma_m B)/3$, and the same argument, with $a = \sigma_q^2 A_0 + \gamma_q$ in place of $\sigma_q^2 A_0$, gives

$$J \;\geq\; \frac{\sigma_q^2 A_0^2}{6} + \frac{(\gamma_q + \gamma_m)A_0}{3} + \frac{2}{3} z_0 \sqrt{a\gamma_m} - \frac{a\gamma_m}{3\sigma_q^2}$$

for every pulse (for $|\zeta| > z_0$ the difference is at least $(\sigma_q z_0 - \sqrt{a\gamma_m}/\sigma_q)^2/3 \geq 0$; for $\gamma_q = 0$ this is the bound of the proof above), while the zero-leakage pulses $C_{\pi+\phi} S C_\phi$ exceed $\frac{2}{3} z_0 \sqrt{a\gamma_m}$ only

by higher-order terms. So **the price of waiting in the memory is** $\frac{2}{3}z_0\sqrt{(\sigma_q^2 A_0+\gamma_q)\gamma_m}$**, to leading order**, when $\sqrt{a\gamma_m} \ll \sigma_q^2 z_0$. Conversely, the same disk argument shows that the square pulse is optimal at this order whenever $\sqrt{a\gamma_m} \geq \sigma_q^2 z_0$: minimising $\sigma_q^2 u^2/3 + \frac{2}{3}\sqrt{a\gamma_m}\,(z_0-u)$ over $0 \leq u = |\zeta| \leq z_0$ then gives $u = z_0$, the value of the square pulse, and $|\zeta| > z_0$ costs more. This is a sufficient condition, not a sharp threshold; for $\gamma_q = 0$ it reads $\Lambda \leq \pi$, inside the certified range of Proposition 8(i). (Checked against the grid minimisation and the explicit family in four parameter sets: the family is within 0.15 % of $\frac{2}{3}z_0\sqrt{a\gamma_m}$ and the lower bound within the stated correction.)

For pulses of fixed phase and no detuning the sufficient condition is not sharp, and the exact threshold can be certified as in Proposition 8(iii): with $r = \gamma_q/\gamma_m$ the difference $J - J_{\mathrm{sq}}$ stays linear in Λ, so $\Lambda_c(r)$ is again a minimum over the turn–wait–turn family, and interval arithmetic encloses it to 10^{-5}; a Krawczyk certificate identifies it, for each r below, with the zero, isolated uniquely in a certified box, of $\{g, g_\phi, g_w\}$ on that family, to 10^{-12} (as for $r = 0$; `umbral_relajacion_cubit/`):

$r = \gamma_q/\gamma_m$	0	0.1	0.3	1	3	10	30	100	1000
$\Lambda_c(r)$	7.908	7.956	8.050	8.366	9.170	11.346	15.521	24.260	67.245
sufficient	3.141	3.264	3.485	4.113	5.374	8.087	12.637	21.632	64.835

(the last row is the condition above, $\Lambda \leq (\pi+\sqrt{\pi^2+16r})/2$, rounded down). Given Lemma 6, $\Lambda_c(r)$ is the infimum over the members of the family with $Q < 0$, a set that does not depend on r, of $4(y+rx)/(-Q)$, affine and non-decreasing in r since $x \geq 0$; so Λ_c **is non-decreasing and concave in** r. **The threshold grows with the qubit relaxation like** $2\sqrt{r}$**, with an exact constant**: given Lemma 6, as $r \to \infty$,

$$\Lambda_c(r) = 2\sqrt{r} + \frac{\pi}{2} + \frac{4}{\sqrt{3}} + O(r^{-1/2}).$$

For the threshold $\Lambda_c^{\mathrm{all}}(r)$ over pulses of any phase and detuning two explicit bounds hold without Lemma 6: as $r \to \infty$,

$$2\sqrt{r} + \frac{\pi}{2} + O(r^{-1/2}) \;\leq\; \Lambda_c^{\mathrm{all}}(r) \;\leq\; 2\sqrt{r} + \frac{\pi}{2} + \frac{4}{\sqrt{3}} + O(r^{-1/2}),$$

the sharp constant of Λ_c^{all} is not identified here. The lower bound is the sufficient condition above, valid for every pulse. The upper bound needs no reading of [14]: the member of the family with $\phi = 2/\sqrt{r}$ and $w = 2/\sqrt{3}$ (units $G = 1/2$) has $Q < 0$ for large r ($Q = -2w\phi + O(\phi^2)$), so it beats the square pulse for every $\Lambda > 4(y+rx)/(-Q)$, and that explicit ratio, with $h = r^{-1/2}$, has a simple pole at $h = 0$ and a Laurent expansion $2\sqrt{r} + \pi/2 + 4/\sqrt{3} + O(r^{-1/2})$ ($\phi = 2/\sqrt{r}$ minimises the leading term $\sqrt{r}(u/2 + 2/u)$, $u = \phi\sqrt{r}$, and $w = 2/\sqrt{3}$ the next one, $w + \pi/2 + 4/(3w)$). Both bounds therefore hold for pulses of any phase and detuning too. At finite r the two explicit quantities, not their expansions, are the bounds: the certified values lie between $(\pi + \sqrt{\pi^2+16r})/2$ and $4(y+rx)/(-Q)$ for $10 \leq r \leq 1000$, and approach the upper one (67.2454 against 67.2788 at $r = 1000$, `cotas_asintoticas.py`).

The exact constant needs a lower bound over every member of the family that could compete, including waits that tend to zero or to infinity. Work in units $G = 1/2$, with $h = r^{-1/2}$, $a = 1 - z$, $D = -Q = 2(1-z^2) - \pi x - x^2$ and $R_r = 4(y+rx)/D$ on the members with $D > 0$, so that, given Lemma 6, $\Lambda_c(r) = \inf R_r$ (step (c)). (a) *Localisation.* With the closed forms of Remark 7, $xy - a^2 = E := \phi^2 - 2(1-\cos\phi) + w(\phi - \sin\phi) \geq 0$.

Since $D \le 2(1-z^2) = 4a(1-a/2)$, AM–GM gives $R_r \ge \frac{2}{h}\sqrt{1+E/a^2}/(1-a/2)$, and $D \le 4a \le 4\sqrt{xy}$ gives $R_r \ge (k+1/k)/h$ with $k = h\sqrt{y/x}$. The explicit member above has $R_r = 2/h + \frac{\pi}{2} + \frac{4}{\sqrt{3}} + O(h)$, so it is enough to bound the members with $R_r \le 2/h + C$ for a fixed C; for them $a = O(h)$, $E/a^2 = O(h)$ and $k = 1 + O(\sqrt{h})$. As $\phi \to 0$, $a = \frac{1}{2}\phi(\phi+w)[1+O(\phi^2)]$ and $E = \frac{1}{12}\phi^3(\phi+2w)[1+O(\phi^2)]$ uniformly in $w \ge 0$ (each is a sum of two non-negative terms), so $E/a^2 \ge \phi/(3(\phi+w))\,[1+O(\phi^2)]$; hence $\phi/w = O(h)$, then $x = \frac{1}{4}w\phi^2[1+O(h)]$, $y = w[1+O(h)]$, $u := \phi/h = 2 + O(\sqrt{h})$, and $m \le w \le M$ for constants independent of h. (b) *Expansion.* On that compact set, uniformly, $R_r = h^{-1}\left(\frac{u}{2} + \frac{2}{u}\right) + c(u,w) + O(h)$ with $c(u,w) = \frac{2-u^2/6}{w} + (u^2+4)\left(\frac{w}{8} + \frac{\pi}{16}\right)$. Since $\frac{u}{2} + \frac{2}{u} - 2 = \frac{(u-2)^2}{2u}$, $c(2,w) = w + \frac{4}{3w} + \frac{\pi}{2} \ge \frac{\pi}{2} + \frac{4}{\sqrt{3}}$ and c is Lipschitz in u there, completing the square in $u-2$ gives $R_r \ge 2/h + \frac{\pi}{2} + \frac{4}{\sqrt{3}} - O(h)$. (c) *From the family to the threshold.* The case analysis of Proposition 8(iii) carries over with the term $4rx$ in $12(\hat{J} - \hat{J}_{\mathrm{sq}}) = \Lambda Q + 4(y + rx)$: the cost still increases with y, the mirrored branch has difference $\Lambda Q_m + 4x_f + 4ry_f > 0$, and on the middle zone $\frac{d}{dx}\hat{J}(x, 2d-x) = \frac{\Lambda}{6}(A_0 + x) + \frac{r-1}{3} > 0$ for $\Lambda > 4/\pi$, while below $(\pi + \sqrt{\pi^2 + 16r})/2 \ge \pi > 4/\pi$ the sufficient condition decides. So the sufficient condition misses the threshold by exactly $4/\sqrt{3}$ asymptotically. (The identity for E and the coefficients $u/2 + 2/u$ and $c(u,w)$ checked symbolically; $\inf R_r$ equal to the certified $\Lambda_c(r)$ to $4 \cdot 10^{-15}$ at the six certified values $r = 1, 3, 10, 30, 100, 1000$; $\sqrt{r}\,(\Lambda_c - 2\sqrt{r} - \frac{\pi}{2} - \frac{4}{\sqrt{3}})$ equal to 3.77 at $r = 10^4$ and 3.76 at $r = 10^6$; with $w = 1$ fixed the constant is 3.907, larger, as it must be: `constante_asintotica.py`.)

Statements (i)–(iii) of Proposition 5 were checked against the exact dynamics averaged by Gauss–Hermite quadrature: the formula agrees to $4 \cdot 10^{-5}$ in relative terms on 300 random complete transfers with complex couplings and real detunings, none of which goes below the floor of (ii). Moving the noise from the qubit to the memory replaces A^2 by B^2 in the coefficient, as the same check confirms ($r/\sigma^2 = 22.70$ against $(2|\zeta|^2 + B^2)/6 = 22.70$ and $(2|\zeta|^2 + A^2)/6 = 0.47$ for one zero-leakage pulse); since $B \ge A_0$, the symmetric floor $\pi^2\sigma_m^2/(96G^2)$ holds as well.

2.6. Correlated offsets.

Remark 10 (correlated offsets written proof, audited). The formula of Proposition 5 holds for correlated offsets: with $C = \mathbb{E}(\xi_q \xi_m)$ and $V = \sigma_q^2 + \sigma_m^2 - 2C$, $6r = 2V|\zeta|^2 + \sigma_q^2 A^2 + 2CAB + \sigma_m^2 B^2$. At second order, the two limiting covariance structures simplify exactly. For a common offset, $\xi_q = \xi_m$, $6r = \sigma^2 T^2$, so the square pulse, of least duration, minimises this coefficient. For opposite offsets, $\xi_m = -\xi_q$ with $\sigma^2 > 0$, $6r = \sigma^2(8|\zeta|^2 + (A-B)^2)$ vanishes exactly when $\zeta = 0$ and $A = B$, i.e. when the curve $\mathbf{r}$ closes; among pulses of fixed phase and no detuning the shortest such transfer is the U-turn with $s = 0$, of duration $L(0) = \frac{7}{3}T_0$ (Proposition 3), the closed-curve condition of [19, 12]. In between, for $C \ge 0$, the square pulse is optimal at second order among all pulses whenever

$$\frac{\pi}{2}\sqrt{(\sigma_q^2 + C)(\sigma_m^2 + C)} \;\ge\; \sigma_q^2 + \sigma_m^2 - 2C.$$

Indeed, with $x = A - A_0$, $y = B - A_0 \ge 0$ and $u = |\zeta|$, the difference from the square pulse is $6\Delta = 2V(u^2 - z_0^2) + 2A_0[(\sigma_q^2 + C)x + (\sigma_m^2 + C)y] + \sigma_q^2 x^2 + 2Cxy + \sigma_m^2 y^2$; the quadratic part is non-negative for $C \ge 0$, and for $u \le z_0$ the disk of Theorem 1(a) and $ax + by \ge 2\sqrt{abxy}$ give $6\Delta \ge 2(z_0 - u)\big[2A_0\sqrt{(\sigma_q^2 + C)(\sigma_m^2 + C)} - 2Vz_0\big]$, with $A_0/z_0 = \pi/2$; for $u > z_0$ every term is non-negative. For $C = 0$ and $\sigma_q^2 + \sigma_m^2 > 0$ this sufficient condition cannot hold, since

$(\pi/2)\sigma_q\sigma_m \le (\pi/4)(\sigma_q^2+\sigma_m^2) < \sigma_q^2+\sigma_m^2$ (if both variances vanish it holds trivially and every pulse has zero noise cost); the content is the correlated case: for equal positive variances the condition is $\kappa \ge (4-\pi)/(4+\pi)$, approximately 0.1201983. For pulses of fixed phase and no detuning the band can be computed exactly. With $\kappa = C/(\sigma_q\sigma_m)$ the square pulse is optimal on $[\varrho_c(\kappa), 1/\varrho_c(\kappa)]$ and on no larger interval, where ϱ_c decreases from 0.12645 at $\kappa = 0$ through 0.0902 ($\kappa = 0.1$), 0.0476 (0.3), 0.0272 (0.5) to 0.00948 at $\kappa = 1$. At $\kappa = 0$, 0.1, 0.1202 ($= 601/5000$), 0.3, 0.5, 1 the threshold is identified exactly, under the same reading of [14]: it is the zero of the stationarity system of the competing family isolated uniquely in a certified box, enclosed to 10^{-12} by a Krawczyk certificate, the bang–bang family for the first three values ($\varrho_c = 0.12645352249$, 0.09019042063, 0.0843166534) and the turn–wait–turn family with a positive wait for the last three (0.047565938906, 0.027186606021, 0.009479169051; every digit printed is fixed by the enclosure); a strictly convex neighbourhood of the critical point, $\partial g/\partial s > 0$ there, and interval bounds on the rest of the region at both ends of the enclosure make the identification exact (`banda_correlada/`). Two facts reduce the certificate to three values of ϱ: with $s = \sqrt{\varrho}$ the difference from the square pulse is $Q_x + bs + Q_y s^2$ with $b \ge 0$ at each point of the region, so optimality at s_{hi} and at $s = 1$ gives it in between; and $\varrho \mapsto 1/\varrho$ exchanges the two elements. A rational turn–wait–turn witness beats the square pulse on $(0, \varrho_{\mathrm{lo}}]$. So **correlation widens the band but never makes the square pulse optimal for every ratio**: as $\varrho \to 0$ the correlation term $C = \kappa\sigma_q\sigma_m$ vanishes and the bang–bang pulse wins again. The competitor is the bang–bang pulse up to $\kappa \approx 0.2072$ and a turn–wait–turn pulse with a positive wait beyond it (numerical). The sufficient condition above holds at equal variances exactly for $\kappa \ge (4-\pi)/(4+\pi)$; at $\kappa = 0$ and 0.1 it holds for no ratio, and at $\kappa = 601/5000$, 0.3, 0.5, 1 its band lies inside the exact one, with lower end exceeding ϱ_c by factors 11.80, 3.79, 3.01 and 2.53 (interval arithmetic, with a failing control, `banda_suficiente.py`).

The band can only widen as the correlation grows, for pulses of any phase and detuning and without Lemma 6: fix positive variances and the relaxation rates; if the square pulse minimises the leading-order error among the pulses of an admissible class at $\kappa_1 \in [0,1]$, it does at every $\kappa_2 \in [\kappa_1, 1]$. Indeed, in the difference 6Δ above, $\partial(6\Delta)/\partial C = 4(z_0^2 - u^2) + 2[A_0(x+y) + xy]$, non-negative for $u \le z_0$ since $x, y \ge 0$ (Theorem 1(a)), while for $u > z_0$ every term of 6Δ is non-negative at every $C \ge 0$; relaxation adds terms that do not depend on C. So where the square pulse is optimal exactly on $[\varrho_c(\kappa), 1/\varrho_c(\kappa)]$ these bands are nested and ϱ_c is non-increasing (not shown to be strictly decreasing).

Remark 11 (one sufficient condition for both formalised). The sufficient conditions of Remarks 9 and 10 are cases of one. With slow offsets on both elements (second moments σ_q^2, σ_m^2, covariance $C \ge 0$, $V = \sigma_q^2 + \sigma_m^2 - 2C \ge 0$) and Markovian relaxation of both ($\gamma_q, \gamma_m \ge 0$, no pure dephasing), the leading-order error is $6J = 2V|\zeta|^2 + \sigma_q^2 A^2 + 2CAB + \sigma_m^2 B^2 + 2(\gamma_q A + \gamma_m B)$, and the square pulse minimises it over all pulses whenever

$$V z_0 \;\le\; \sqrt{\big((\sigma_q^2 + C)A_0 + \gamma_q\big)\big((\sigma_m^2 + C)A_0 + \gamma_m\big)}.$$

The proof is the disk of Theorem 1(a), $(z_0 - |\zeta|)^2 \le (A - A_0)(B - A_0)$, followed by the arithmetic–geometric mean inequality. Without relaxation it is the condition of Remark 10; with $\sigma_m = C = 0$ it is the condition of Remark 9; with $\gamma_q = \gamma_m = C = 0$ and $\sigma_q^2 + \sigma_m^2 > 0$ it cannot hold ($\frac{\pi}{2}\sqrt{\sigma_q^2\sigma_m^2} < \sigma_q^2 + \sigma_m^2$; if both variances also vanish it holds trivially and every pulse has zero leading-order cost), in agreement with the bang–bang pulse beating the square pulse outside $[\varrho_c, 1/\varrho_c]$. It is sufficient, not sharp: at $\gamma_q = \gamma_m = C = 0$ the exact band of

Proposition 5 is wider. (Checked on 3620 random parameter sets meeting the condition, over the region the disk allows: minimum of $6(J - J_{\mathrm{sq}})$ equal to $3.6 \cdot 10^{-4} \geq 0$; the control with $\gamma_q = \gamma_m = C = 0$, $\varrho = 0.05$ and the bang–bang pulse of $d = 0.8$ violates the condition and gives $6(J - J_{\mathrm{sq}}) = -0.115$; `condicion_unificada.py`.)

3. Limits of this paper

Questions of the companion paper addressed here. [10] gives the error of the square pulse under a quasi-static offset, $(8 + \pi^2)\sigma^2/(96G^2)$, and leaves open which pulse is best under slow noise, alone or together with loss. The companion questions are addressed through exact second-order floors, certified square-pulse optimality ranges, and explicit improving protocols (Propositions 5 and 8): the floor for every pulse is $\pi^2\sigma^2/(96G^2)$, and the square pulse is optimal in certified ranges of σ_m^2/σ_q^2 and of $\sigma_q^2/(\gamma_m G)$. For fixed-phase, zero-detuning pulses, the two square-pulse thresholds are identified (ϱ_c and Λ_c). With memory loss, the general optimal pulse at every finite noise-to-loss ratio is not determined here; the waiting family attains the stated asymptotic optimum in the large-ratio regime, and for fixed phase and no detuning it contains the end point of the optimum for every $\Lambda > \Lambda_c$ (given Lemma 6; the control that reaches it is not shown to be unique). The exact thresholds with qubit relaxation, $\Lambda_c(\gamma_q/\gamma_m)$, and with correlated offsets, $\varrho_c(\kappa)$, are certified at the listed parameter values only; for fixed phase and no detuning, $\Lambda_c(r) = 2\sqrt{r} + \pi/2 + 4/\sqrt{3} + O(r^{-1/2})$ is proved given Lemma 6, and Λ_c is non-decreasing and concave in r; for arbitrary phase and detuning the two explicit bounds with constants $\pi/2$ and $\pi/2 + 4/\sqrt{3}$ hold, but the sharp constant is not identified; ϱ_c is proved non-increasing in κ; the change of competitor at $\kappa \approx 0.2072$ is numerical.

The results are for two states, one excitation and a static offset: exactly at first order in it (Theorem 1) and at second order (Propositions 5 and 8), with relaxation of either or both elements included at the leading order used in Propositions 3 and 8 and Remarks 9 and 11, and in the regime $\sigma T \ll 1$. Offsets that change during a transfer, and Markovian relaxation and dephasing beyond that order, are outside; the latter are treated in [10]. The exact cost of Proposition 3 is proved for pulses of fixed phase and no detuning only, and whether a pulse of varying phase, or a detuning, can do better is open; curves made of pieces in planes through the axis are excluded. For pulses of varying phase or with detuning, undecided ranges remain in Proposition 5(iii) (0.1264–0.137 and 7.30–7.91; for fixed phase and no detuning (iv) identifies the threshold with the bang–bang crossing, given our reading of [14] in Lemma 6) and in Proposition 8(i) ($7.1 < \Lambda \leq \Lambda_c$; for every $\Lambda > \Lambda_c$ the explicit planar competitor already beats the square pulse, and for fixed phase and no detuning (iii) identifies the threshold with the crossing of a short-wait turn–wait–turn pulse, under the same reading). Those bounds come from a relaxation of the reachable set; the conjectured spatial exposure boundary of Remark 7 would extend the planar constraints to all pulses and could sharpen them, but determining the resulting optima would still require the remaining minimisation. The reduction from the two-level dynamics to the space curve is a written proof, and the paper has not been refereed by a journal.

4. What backs each statement

A formal label refers to the Isabelle statement named in each entry, not to the physical claim as a whole. Two bridges are written proofs throughout. (i) *Reduction:* the passage from the two-level dynamics to the space curve $\mathbf{n}$ and to the exposures A, B, ζ, and the Dyson expansion that gives the first-order term Ξ_1. (ii) *Inputs:* where a lemma takes a formula as a hypothesis (an integrated exposure, the candidate lengths of Dubins' classification), its derivation is written. The Isabelle theories are public, in https://github.com/karlesmarin/formalizations, folder `Isabelle/MemoryAccessFloors`; each certificate is listed with its script, and the commands and archived outputs are in the supplement.

Theorem 1 and Corollary 2 (robustness): *Backing:* proof in the text; numerical minima of the duration at prescribed leakage fall on the frontier to five digits; the bang–bang family checked by integrating the two-level dynamics, the turn–wait–turn family on the Bloch curve (2000 cases, slack below $5 \cdot 10^{-13}$); for a real coupling at $\zeta = 0$ the exact shortest U-turns (Dubins paths, checked by integrating each path and by an independent solver) lie above the bound. *Status:* the elementary steps formalised in Isabelle/HOL (`concavity_step`, `tent_integral`, `frontier_step`, `disk_from_two_bounds`, both equality families, `robust_cost_amgm`); (a) and (b) formalised for continuous unit curves n that are differentiable at every interior time, with the speed bound $2G$ imposed pointwise, through the poles included (`Robust_Disk.robust_disk`, `Robust_Core.frontier_general`); the written proof establishes the same statements for Lipschitz curves, including those of switching controls; the equality checks are formalised conditional on the explicit exposure formulas, whose derivation from the named controls is written; the reduction from the two-level dynamics to the curve (A, B, ζ as integrals of n) is the written proof.

Proposition 3 (exact cost for fixed-phase, zero-detuning pulses): *Backing:* proof in the text from Dubins' theorem, with the comparison of the eight candidates done analytically; the closed forms of the candidates agree to $7 \cdot 10^{-15}$ with their geometry for $0 \le s \le 4$, each path checked by integration and by an independent solver, and the optimal paths, run as pulses through the two-level dynamics, transfer completely with $\zeta = 0$; the minimum over s agrees with the formula to 10^{-7} and falls at the predicted $|s|$; curves out of the plane found nothing lower. *Status:* the minimisation over s formalised in Isabelle/HOL, with $L(s)$ as a definition (`csc_minimum`, `ccc_branch`, `closed_form_value`, `total_minimum`); the comparison of the eight candidates formalised too (`Dubins_U_Turn`), with their closed forms as definitions; the derivation of the candidate lengths from the paths is written and checked numerically for $s \neq 2$, with the degenerate end point $s = 2$ treated separately (at $s = 2$ the candidate `Lccc_b` keeps an extra-turn representative of length 4π; the shortest degenerate representative has length 2π, and the minimum at $s = 2$ is 2π); that the shortest is CSC or CCC is Dubins' theorem, cited; unequal rates with varying phase or detuning open (pieces in planes through the axis excluded, `Meridian_Unfolding`).

Proposition 5 (slow noise at second order): *Backing:* proof in the text; exact dynamics averaged by Gauss–Hermite on 300 random transfers (relative error $< 4 \cdot 10^{-5}$), control with the noise in the memory; thresholds by interval certificates (`certificar_cuadrado_optimo`, `certificar_robustez`). *Status:* the unitarity step of the expansion (`Unitarity_Second_Order`: a family $I - i\epsilon X_1 - \epsilon^2 X_2 + O(\epsilon^3)$ unitary for small ϵ has X_1 Hermitian and $X_2 + X_2^\dagger = X_1^2$, hence the formula for $2\,\mathrm{Re}(\Xi_2)_{aa}$; the formal result proves the coefficient identities and the truncated-modulus algebra, while the dynamical construction, the identification of X_1, the stochastic averaging and the averaged remainder are written proofs), the algebra, (i), (ii) given Theorem 1(a), and the bang–bang example at $\varrho = 0.1$ formalised in Isabelle/HOL, with the kernel only (`Slow_Noise_Second_Order`, `Slow_Noise_Kernel`); in (iii) the square pulse beaten for $\varrho \le 0.1264$ and $\varrho \ge 7.91$ certified by interval arithmetic (one pulse, the

difference being linear in ϱ), its optimality on $[0.137, 7.30]$ certified by branch and bound with intervals, with a control at $\varrho = 0.12$ that fails (the certificate at $\varrho = 0.137$ is implemented with both mpmath intervals and Arb; the extension to the decimal endpoint 7.30 uses the Python certificate at the exact rational $\varrho = 10/73$); (iv) the enclosure $0.126453 < \varrho_c \leq 0.12645395$ certified by intervals given Lemma 6 (`certificar_umbral_plano`: optimality over the cut region at 0.12645395 and 1, 55 567 nodes and 22 322 accepted leaves; strict witness $d = 0.7717755$ at 0.126453); the identification $\varrho_c = \varrho_{\mathrm{bb}}$ by a second certificate (`certificar_rho_c_igual_bb`: Krawczyk in (ϱ, d), enclosure widths below $6 \cdot 10^{-41}$ for ϱ_{bb} and $4 \cdot 10^{-40}$ for d_*; Hessian positive definite on N uniformly for ϱ in the enclosure; margin at least $5 \cdot 10^{-7}$ on the retained branch boxes outside N, the near-square patch and the tail being treated separately; five controls that fail; 1.3 s). The displayed root values are rounded approximations; the existence and uniqueness tests use the full internal intervals. The gradient and Hessian formulas were checked against finite differences (agreement $2 \cdot 10^{-8}$; a sign error gives 10^{-2}), which detects errors but is not part of the proof.

Lemma 6 (planar exposure bound): *Backing:* written derivation from the sections of [14, Lemma 3.1 and Theorem 11.1] (branch selection and the passage to exposures, both lateral signs covered directly); the inequality checked with both lateral signs on 7061 random turn–wait–turn pulses in the sector (`verificar_v291`: the only point below the bound is a member of the equality family itself at $a = 5.6 \cdot 10^{-8}$, short by $2.8 \cdot 10^{-9}$, the rounding of inverting a_z near $a = 0$); the closed forms a_z, b_z against the family of the remark to $2 \cdot 10^{-13}$ (`verificar_v289`). *Status:* the elementary parts formalised in Isabelle/HOL (`Memory_Floors_Pieces`: the derivatives and signs of a_z, b_z, their values at d, the monotonicity and the range of a_z, and the integral of step (1): `a_z_deriv`, `b_z_deriv`, `a_z_mono`, `b_z_antimono`, `a_z_range`, `odd_turns_integral`); the properties of the family are formalised too (Remark 7); the cut used in step (2), that the arc $U_3 = (1, 0, -1)$ is the lower envelope of the section in the sector, is certified within the extremal families U_1, U_2, U_3, U_6 by interval arithmetic for $z \leq 1 - 2^{-12}$ (`envolvente_U3`: 67 687 nodes, minimum margin $4 \cdot 10^{-4}$, six controls that fail), with closed-form identities checked to 60 digits rather than proved symbolically; the identification of these closed forms with the trajectories and the end region $z > 1 - 2^{-12}$ rest on the written argument, so this is not a complete certificate of the lemma; that these families form the boundary of the reachable set is taken from the classification of Dubins extremals [14, Lemma 3.1 and Theorem 11.1], and steps (2) and (3) are otherwise written.

Remark 7 (the exact planar section): *Backing:* the identification with the type-$(1, 0, -1)$ arc of [14, Lemma 3.1 and Theorem 11.1] is a written argument; the closed forms checked against piecewise pulses with exact integrals (`hueco_17iii_b`); the conjecture by an adversarial search over three-dimensional pulses (`conjetura_adversario`). *Status:* the properties of the family (strict bounds of Theorem 1, monotonicity, asymptote, bang–bang end point) formalised in Isabelle/HOL from the closed forms (`U_Turn_Exposure_Family`); the conjecture open.

Remark 9 (relaxation of the qubit): *Backing:* the bound and the sufficient condition are written proofs from Theorem 1(a), checked against grid minimisation and the explicit family in four parameter sets; the exact threshold $\Lambda_c(r)$ enclosed by interval arithmetic at nine ratios $r = \gamma_q/\gamma_m$ from 0 to 1000 (`cert_relajacion`; the control $r = 0$ reproduces the enclosure of Proposition 8(iii)) and identified with the crossing of the turn–wait–turn family by Krawczyk certificates (`certificar_identificacion_r`: unique zero to 10^{-12}, controls that fail; at $r = 1000$ with a narrower neighbourhood N), given Lemma 6; the large-r bounds are written proofs (the sufficient condition below; an explicit member of the family above, whose ratio is expanded in $r^{-1/2}$), compared with the certified values at $r = 10, 30, 100, 1000$ and evaluated at $r = 10^4, 10^6, 10^8$ as high-precision numerical checks, with no threshold certified there (`cotas_asintoticas.py`; with $w = 1$ instead of $2/\sqrt{3}$ the constant is larger,

as it must be). *Status:* the bound and the sufficient condition formalised in Isabelle/HOL (`Noise_Loss_Relax`, `Square_Condition`); the thresholds certified; the sharp constant and the $O(r^{-1/2})$ remainder established by the written localisation and uniform-expansion argument, given Lemma 6, with its elementary steps formalised (`Dubins_Family_Identity.dubins_identity`, `dubins_quantity_nonneg`, `dubins_D_bound`; `Threshold_Localisation.threshold_lower_bound`, `threshold_lower_bound_k`, `complete_square_bound`) and the uniform expansion not formalised; non-decrease and concavity in r from the infimum over a fixed domain, the lemma formalised (`Inf_Affine_Concave.inf_affine_concave`, `inf_affine_monotone`); the symbolic checks confirm the identities and coefficients and do not replace the uniform argument.

Remark 10 (correlated offsets): *Backing:* the two exact limits and the sufficient condition are short written proofs (AM–GM on the disk of Theorem 1(a)); the exact band certified by interval arithmetic at six values of κ given Lemma 6 (`cert_banda`: optimality at s_{hi} and $s = 1$, a rational witness on $(0, \varrho_{\mathrm{lo}}]$, enclosures of relative width $1.2 \cdot 10^{-3}$; the control $\kappa = 0$ reproduces ϱ_c; twelve controls that fail); the identification with the family crossing at the same six values by Krawczyk certificates (`certificar_rho_c_bb_kappa` for the bang–bang competitor at $\kappa = 0$, 0.1, 601/5000, with the middle zone bounded analytically; `certificar_identificacion_kappa` for the turn–wait–turn competitor at $\kappa = 0.3$, 0.5, 1; enclosures of ϱ_c of width 10^{-12}; the control $\kappa = 0$ reproduces the original certificate to 30 digits; Krawczyk off the zero, a tiny neighbourhood, a shifted end point and a dropped correlation term all fail); the change of competitor ($\kappa_w \approx 0.2072$, where the Hessian determinant of the family vanishes) is numerical; that ϱ_c is non-increasing in κ is a short written proof valid for any phase and detuning (the derivative of 6Δ in C), its algebra formalised (`Correlation_Monotone.six_delta_difference`, `delta_monotone_in_C`, `delta_nonneg_transfer`); the comparison of the sufficient band with ϱ_c by `banda_suficiente.py`. *Status:* the interval lemma behind the band formalised (`Band_Quadratic.band_between`); the sufficient condition formalised in Isabelle/HOL with the disk of Theorem 1(a) as a hypothesis (`Memory_Floors_Pieces`: `square_corr`, `square_corr_far`, `square_corr_uncorrelated_fails`, `square_corr_equal_variances`); the two limits are written.

Remark 11 (one sufficient condition): *Backing:* written proof (disk of Theorem 1(a) and AM–GM); 3620 random parameter sets meeting the condition, minimum of $6(J - J_{\mathrm{sq}})$ equal to $3.6 \cdot 10^{-4}$, and a control that violates it and gives -0.115 (`condicion_unificada`). *Status:* formalised in Isabelle/HOL with the disk as a hypothesis, kernel only (`Square_Unified_Condition`: `square_unified` and its three special cases).

Proposition 8 (slow noise and memory loss): *Backing:* lower bound in closed form from Theorem 1(a); upper bound in the large-Λ regime by the explicit zero-leakage family; their difference relative to the excess $J_{\mathrm{lb}} - \pi^2\sigma_q^2/(96G^2)$ falls from 7.9 % at $\Lambda = 200$ to 0.24 % at $\Lambda = 2 \cdot 10^5$ (`brecha_cotas_noiseloss.py`); the thresholds by interval certificates (`certificar_ruido_perdidas`: $\Lambda \le 7.1$ and $\Lambda \ge 7.92$, with a control at $\Lambda = 8.5$ that fails; for fixed phase and no detuning `certificar_umbral_plano_perdidas`: $7.90788 \le \Lambda_c < 7.90789$, 4663 nodes, strict witness at 7.90789, controls at 7.90789, 7.91, 7.92 that fail, given Lemma 6; the identification with the crossing 7.9078851606875… of the turn–wait–turn family by `certificar_lambda_c_identificacion`: Krawczyk in (Λ, ϕ, w), Λ coordinate of width below $2.7 \cdot 10^{-39}$ and every coordinate below $3.3 \cdot 10^{-39}$; Hessian $> 0.3\,I$ on N uniformly for every Λ in the enclosure (57 nodes, 29 leaves) and the isolated critical point strictly inside N; margin at least $1.88 \cdot 10^{-7}$ on the retained exterior boxes at both endpoints, the near-square patch and the non-physical boxes being separated; five controls that fail); exact dynamics with losses and averaged offsets confirms the gain inside the regime $\sigma_q T \ll 1$ and its loss outside. The optimum above Λ_c: written case analysis given Lemma 6 and Theorem 1(b); each step checked on $1.7 \cdot 10^6$ points with a control that fails (`prueba_Qm`); a piecewise search (8–24 segments, 252

starts per value) agrees with the family minimum to 10^{-8} for Λ up to 1000 (`optimo_sobre_lambda_c`). *Status:* the lower bound formalised in Isabelle/HOL (`Noise_Loss_Bound`), given Theorem 1(a), and its extension to $\gamma_q > 0$ of the remark that follows (`Noise_Loss_Relax`); the bound $Q_m \geq \pi\phi + s(\pi - 2s)$ on the mirrored branch formalised (`Symmetric_Branch.symmetric_branch_bound`) and the middle-zone step (`Middle_Zone.middle_zone_min`); the threshold enclosures and the crossing identification in (i) and (iii) certified by interval arithmetic, (iii) given Lemma 6, while the optimality of the turn–wait–turn family for every $\Lambda > \Lambda_c$ is a written proof with the mirrored-branch and middle-zone inequalities formalised as above; for varying phase or detuning, $7.1 < \Lambda \leq \Lambda_c$ open.

Appendix A. Formula sheet

quantity	formula
robustness (complete lossless transfer, $\lvert g\rvert \leq G$, any detuning)	$\big(\lvert\zeta\rvert - z_0\big)^2 \leq (A - A_0)(B - A_0)$, $A, B \geq A_0$; $T \geq T_0 + \dfrac{1}{G}\arccos\dfrac{\lvert\zeta\rvert}{z_0}$ for $\lvert\zeta\rvert \leq z_0$; $\quad T_0 = \dfrac{\pi}{2G}$, $A_0 = \dfrac{\pi}{4G}$, $z_0 = \dfrac{1}{2G}$
cost of $\zeta = 0$, fixed phase, no detuning, $\gamma_q, \gamma_m > 0$	$\min\big[\gamma_q(A - A_0) + \gamma_m(B - A_0)\big] = \dfrac{\gamma_q + \gamma_m}{2G}\arccos\dfrac{\lvert\gamma_q - \gamma_m\rvert}{\gamma_q + \gamma_m}$
slow noise, second order (static offsets, no losses)	$r = \dfrac{2\,\mathbb{E}(\xi_q - \xi_m)^2\lvert\zeta\rvert^2 + \mathbb{E}(\xi_q A + \xi_m B)^2}{6}$; qubit only: $r \geq \dfrac{\pi^2\sigma_q^2}{96G^2}$
noise on both elements, fixed phase, no detuning	square pulse optimal exactly for $\varrho_c \leq \sigma_m^2/\sigma_q^2 \leq 1/\varrho_c$, $\varrho_c = 0.12645352248541\ldots$ (the bang–bang crossing)
correlated offsets, $C = \mathbb{E}(\xi_q\xi_m) \geq 0$	$6r = 2V\lvert\zeta\rvert^2 + \sigma_q^2 A^2 + 2CAB + \sigma_m^2 B^2$, $V = \sigma_q^2 + \sigma_m^2 - 2C$; square pulse optimal if $\dfrac{\pi}{2}\sqrt{(\sigma_q^2 + C)(\sigma_m^2 + C)} \geq V$; fixed phase, no detuning: exactly on $[\varrho_c(\kappa), 1/\varrho_c(\kappa)]$, $\varrho_c(0.5) = 0.0272$, $\varrho_c(1) = 0.00948$
offsets and relaxation of both elements	square pulse optimal if $Vz_0 \leq \sqrt{\big((\sigma_q^2 + C)A_0 + \gamma_q\big)\big((\sigma_m^2 + C)A_0 + \gamma_m\big)}$

quantity	formula
slow noise on the qubit and memory decay	$\min J = \frac{\pi^2\sigma_q^2}{96G^2} + \frac{\sigma_q\sqrt{\pi\gamma_m}}{6G^{3/2}} + O(\gamma_m/G)$ for $\gamma_m G \ll \sigma_q^2 \ll \sqrt{\gamma_m}G^{3/2}$; square pulse optimal for $\sigma_q^2 \le 7.1\,\gamma_m G$, not for $\sigma_q^2 \ge 7.92\,\gamma_m G$; fixed phase, no detuning: optimal exactly for $\sigma_q^2 \le \Lambda_c\gamma_m G$, $\Lambda_c = 7.9078851606875\ldots$; with qubit relaxation $r = \gamma_q/\gamma_m$: $\Lambda_c(1) = 8.366$, $\Lambda_c(1000) = 67.245$, $\Lambda_c(r) = 2\sqrt{r} + \pi/2 + 4/\sqrt{3} + O(r^{-1/2})$ given Lemma 6, non-decreasing and concave in r; any phase and detuning: $2\sqrt{r} + \pi/2 \lesssim \Lambda_c^{\mathrm{all}}(r) \lesssim 2\sqrt{r} + \pi/2 + 4/\sqrt{3}$ up to $O(r^{-1/2})$

Appendix B. The companion site

This paper has a companion website, its *blueprint*, at https://karlesmarin.github.io/blueprints/robust-offset/. It shows what the paper proves, how its results depend on each other and what exactly has been checked, and lets the reader go from any statement to its proof, its dependencies and its checks:

- *The statements.* Every numbered statement of the paper, copied verbatim with its number (1 theorem, 3 propositions, 1 lemma, 1 corollary and 5 remarks), each with its proof.
- *The dependencies.* What each statement uses and where it is used, as a list and as a graph, including the links to the companion paper.
- *The map.* The main results, the tools they rest on and the fields they feed.
- *What was checked, and how.* For each statement, the backing and status of Section 4: 10 statements have named pieces machine-checked in Isabelle/HOL (94 theorems, linked by name to the public repository), and each numerical claim names the interval certificate of the supplementary material that backs it.
- *Labs.* Next to Theorem 1, Propositions 3 and 5 and Remarks 9, 10 and 11, a form runs the paper's own computations in the reader's browser (Python, nothing to install) for values the reader chooses.
- *Figures in motion.* Figures 1–3 are drawn live by the same computations, with the parameters of the paper; each figure of this PDF is a link to its lab.
- *The paper itself.* Its PDF, opened at the right page from every statement.

The site is a reading aid, not a part of the proofs: the statements and proofs are those of this paper, whose text prevails, and a formal or computational check is reported for exactly the part it covers.

Acknowledgement

We worked with an AI assistant.

References

[1] Q. Ansel, S. J. Glaser and D. Sugny, Selective and robust time-optimal rotations of spin systems, J. Phys. A 54, 085204 (2021); we used arXiv:2010.12454v2, §4.

[2] J. Bae, J. H. Bai, B.-Y. Lee, J.-Y. Lee and C.-H. Lee, On the reachability of 3-dimensional paths with a prescribed curvature bound, arXiv:2403.18707v3 (2025); Theorems 3 and 5, and p. 10.

[3] E. Barnes, F. A. Calderon-Vargas, W. Dong, B. Li, J. Zeng and F. Zhuang, Dynamically corrected gates from geometric space curves, Quantum Sci. Technol. 7, 023001 (2022); we used arXiv:2103.16015v1, Eqs. (6)–(8) and §II.B.

[4] M. Bergonzoni, S. Jandura and G. Pupillo, iSWAP gate with polar molecules: robustness criteria for entangling operations, Phys. Rev. A 112, 032621 (2025); we used arXiv:2502.21238v2.

[5] D. Daems, A. Ruschhaupt, D. Sugny and S. Guérin, Robust quantum control by a single-shot shaped pulse, Phys. Rev. Lett. 111, 050404 (2013); we used arXiv:1304.4016, Eqs. (7)–(10).

[6] L. E. Dubins, On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents, Amer. J. Math. 79, 497–516 (1957).

[7] S. Jandura, J. D. Thompson and G. Pupillo, Optimizing Rydberg gates for logical-qubit performance, PRX Quantum 4, 020336 (2023); we used arXiv:2210.06879v3, Sections II–III and Appendix A.

[8] E. Kyoseva, H. Greener and H. Suchowski, Detuning-modulated composite pulses for high-fidelity robust quantum control, Phys. Rev. A 100, 032333 (2019); we used arXiv:1811.02623v1.

[9] X.-J. Lu, X. Chen, A. Ruschhaupt, D. Alonso, S. Guérin and J. G. Muga, Fast and robust population transfer in two-level quantum systems with dephasing noise and/or systematic frequency errors, Phys. Rev. A 88, 033406 (2013); we used arXiv:1305.6127v1.

[10] C. Marín, Exact limits on writing a superconducting quantum memory: when can pulse shaping beat the resonant swap?, companion paper, Zenodo (2026), doi:10.5281/zenodo.24250501.

[11] D. Mittenhuber, Dubins' problem is intrinsically three-dimensional, ESAIM: Control Optim. Calc. Var. 3, 1–22 (1998); Theorem 1.1 and the planar and spatial shifts, pp. 2–3.

[12] H. Nelson and E. Barnes, Exact quantum speed limits for quantum gates from space curve geometry, Phys. Rev. A 114, 052426 (2026), doi:10.1103/28zn-ch8y; we used arXiv:2604.23031v1.

[13] V. S. Patsko and A. A. Fedotov, Analytic description of a reachable set for the Dubins car, Trudy Inst. Mat. Mekh. UrO RAN 26(1), 182–197 (2020), in Russian; Eq. (2.6), §§3.3–3.4 and Theorem 1.

[14] V. S. Patsko and A. A. Fedotov, Three-dimensional reachable set for the Dubins car: foundation of analytical description, Commun. Optim. Theory 2022, Art. 23, 1–42, doi:10.23952/cot.2022.23; Eq. (3.5), §3.4, Property 8 (§3.7), Theorem 11.1.

[15] P. Souères, A. Balluchi and A. Bicchi, Optimal feedback control for route tracking with a bounded-curvature vehicle, Proc. IEEE Int. Conf. Robotics and Automation (ICRA 2000), vol. 3, 2473–2478; Property 5 and the synthesis.

[16] J. M. Sullivan, Curves of finite total curvature, in *Discrete Differential Geometry*, Birkhäuser (2008); we used arXiv:math/0606007, §5 (Schur's comparison theorem and its proof).

[17] H. J. Sussmann, Shortest 3-dimensional paths with a prescribed curvature bound, Proc. 34th IEEE CDC (1995) 3306–3312; Theorem 1 and the proof on p. 3312.

[18] L. Van Damme, Q. Ansel, S. J. Glaser and D. Sugny, Robust optimal control of two-level quantum systems, Phys. Rev. A 95, 063403 (2017); we used arXiv:1704.07653.

[19] J. Zeng and E. Barnes, Fastest pulses that implement dynamically corrected single-qubit phase gates, Phys. Rev. A 98, 012301 (2018); we used arXiv:1802.10241v2.

[20] J. Zeng, C. H. Yang, A. S. Dzurak and E. Barnes, Geometric formalism for constructing arbitrary single-qubit dynamically corrected gates, Phys. Rev. A 99, 052321 (2019); we used arXiv:1811.04864v2, Eqs. (1)–(8).