spin: theory manual

Pulsed magnetic resonance and qubit readout, with exercises on the emulator

Rok Žitko
Jozef Stefan Institute, Ljubljana, Slovenia
rok.zitko@ijs.si

Red Pitaya `spin` design, version 1.1
November 8, 2026

Abstract

This manual explains the physics that the `spin` transceiver teaches and the signal processing it does: the Bloch equations and the rotating frame; pulses and their spectra; the free induction decay and its spectrum; spin echoes, CPMG and spin-locked trains; the measurement of T_1 and T_2; phase cycling; nuclear quadrupole resonance; pulsed electron spin resonance through a microwave bridge; dispersive readout of a qubit, single-shot decisions and active reset; the receiver's mixing, filtering and averaging; and the emulator's model of a spin ensemble and of a qubit, with its limits. Each section ends with exercises that run on the emulator, so a class needs no magnet; their expected results are computed here. Figures are computed with the float Bloch equations and `spin`'s own Python model; the board's recordings are in the user's manual (`user_manual.pdf`).

Contents

1 Introduction **3**

2 The Bloch equations and the rotating frame **3**
2.1 Precession . . . 3
2.2 The rotating frame . . . 3
2.3 Relaxation . . . 4
2.4 Why $T_2 \leq 2T_1$ . . . 5
2.5 The qubit's Bloch sphere . . . 5

3 Pulses **6**
3.1 Hard pulses and their band . . . 6
3.2 Calibration and B_1 inhomogeneity . . . 7
3.3 Chirps and adiabatic inversion . . . 7

4 The free induction decay and its spectrum **7**
4.1 Lines and decays . . . 7
4.2 Sampling, quadrature and the spectrum . . . 8
4.3 The receiver's filter . . . 9
4.4 One isochromat per shot . . . 9

5 Spin echoes and echo trains **9**
5.1 The Hahn echo . . . 9
5.2 CPMG . . . 10
5.3 Spin-locked trains . . . 11
5.4 One isochromat, two signs of offset . . . 11

6 Relaxation measurements **11**
6.1 T_1 . . . 11
6.2 The repetition time . . . 12
6.3 T_2 and $T_{1\rho}$ . . . 12

7 Phase cycling **12**
7.1 What a cycle selects . . . 12
7.2 Receiver errors and CYCLOPS . . . 13
7.3 Echoes: EXORCYCLE . . . 13

8 Nuclear quadrupole resonance **14**
8.1 The quadrupole interaction . . . 14
8.2 Powders and the flip angle . . . 14
8.3 Finding a line . . . 14

9 Pulsed ESR through a microwave bridge **15**
9.1 Electron spins at microwave frequencies . . . 15
9.2 Single-sideband up- and down-conversion . . . 15
9.3 Echoes and field sweeps . . . 15

10 Qubit readout and decisions **16**
10.1 Dispersive readout . . . 16
10.2 The matched filter . . . 16
10.3 Thresholds and fidelity . . . 16
10.4 Active reset and repeat-until-success . . . 16

11 The receiver: mixing, filtering, averaging **18**
11.1 Quadrature demodulation with the transmitter's oscillator . . . 18
11.2 The order-2 CIC filter . . . 18
11.3 Averaging . . . 19
11.4 Weighted integration . . . 19

12 The emulator's model **19**
12.1 One vector, one rotation per slot . . . 19
12.2 Limits that follow from the model . . . 20
12.3 The qubit . . . 20
12.4 How it is verified . . . 20

A Expected results of the exercises **21**

1 Introduction

Pulsed magnetic resonance and qubit readout share one picture: a two-level system, or an ensemble of them, driven by short bursts of a radio-frequency field and observed through the field it radiates back (or, for a qubit, through a resonator whose response it shifts). This manual develops that picture as far as `spin` uses it, and explains what the instrument does to the signals on the way. It is written for a lab course: each section ends with exercises that run on `spin`'s built-in emulator (no magnet, probe or qubit needed), with the results to expect.

Notation. Frequencies are in hertz: f_0 a spin's resonance (the Larmor frequency, or a qubit's transition), f_{TX} the transmitter's, $\Delta = f_0 - f_{\mathrm{TX}}$ the offset, f_{R} the Rabi frequency (the nutation rate under a pulse). $\boldsymbol{M} = (M_x, M_y, M_z)$ is the magnetization in units of its equilibrium value; for a qubit the Bloch vector, with $M_z = +1$ the ground state (the convention of `spin`'s emulator). T_1, T_2 and T_2^* are the longitudinal, the transverse and the inhomogeneous decay times. Signals are volts at the instrument's connectors.

The emulator. The exercises use the templates of `spin.experiments` on the emulator (user's manual, sections on the emulator and the experiments). A typical start:

```
from spin import SpinClient, experiments as X, analysis as A
dev = SpinClient("192.168.0.102")
dev.apply({"emulator": {"mode": "ensemble", "frequency": 10.02e6, "rabi": 1.25e6, "t1": 1e-3,
                        "t2": 300e-6, "signal": 0.1, "detuning": {"shape": "gaussian", "width": 20e3
    ↪ }},
           "rx": {"1": {"input": "emulator"}}})
p = X.pulses(f=10e6, amplitude=0.8, t90=0.25e-6, dwell=0.5e-6, points=256, tr=5e-3, scans=64)
r = X.acquire(dev, X.fid(p))
```

The emulator is a spin 1/2 (or a qubit) updated every 128 ns; section 12 says what that implies. `X.simulate(driver, exp)` runs the same experiment on a PC in the bit-exact model, which is slower but needs no board.

Further reading. Levitt's *Spin Dynamics* [1] and Keeler's *Understanding NMR Spectroscopy* [2] cover pulsed NMR; Fukushima and Roeder [3] the experimenter's side; Ernst, Bodenhausen and Wokaun [4] phase cycling and coherence pathways; Das and Hahn [5] NQR; Schweiger and Jeschke [6] pulsed ESR; Krantz et al. [7] and Blais et al. [8] qubit readout.

2 The Bloch equations and the rotating frame

2.1 Precession

A nucleus or an electron with spin carries a magnetic moment $\boldsymbol{\mu} = \gamma\hbar\boldsymbol{I}$; in a field $\boldsymbol{B}_0 = B_0\hat{z}$ an ensemble's magnetization $\boldsymbol{M}$ precesses about $\hat{z}$ at the Larmor frequency $f_0 = \gamma B_0/2\pi$ (42.58 MHz/T for protons, 28.0 GHz/T for a free electron):

$$\frac{\mathrm{d}\boldsymbol{M}}{\mathrm{d}t} = \gamma\,\boldsymbol{M} \times \boldsymbol{B}. \tag{1}$$

In equilibrium $\boldsymbol{M}$ points along $\hat{z}$ with a size set by the Boltzmann populations. A coil around the sample picks up the precessing transverse part $M_\perp = M_x + \mathrm{i}M_y$ as an EMF at f_0: the NMR signal [9].

2.2 The rotating frame

An RF field $2B_1\cos(2\pi f_{\mathrm{TX}}t + \phi)$ perpendicular to $\boldsymbol{B}_0$ is two counter-rotating fields; the one rotating with the spins acts, the other averages out. In a frame rotating at f_{TX} about $\hat{z}$ the RF field stands

still, along the axis set by its phase ϕ ($\phi = 0$: x, 90°: y), and $\boldsymbol{B}_0$ shrinks to the offset's field. The magnetization then turns about an *effective field* $\boldsymbol{\Omega}$ (in angular frequency):

$$\frac{\mathrm{d}\boldsymbol{M}}{\mathrm{d}t} = \boldsymbol{\Omega} \times \boldsymbol{M}, \qquad \boldsymbol{\Omega} = 2\pi\,(f_\mathrm{R}\cos\phi,\ f_\mathrm{R}\sin\phi,\ \Delta), \qquad \Delta = f_0 - f_\mathrm{TX}, \tag{2}$$

with the Rabi frequency $f_\mathrm{R} = \gamma B_1/2\pi$, proportional to the transmitter's amplitude (`spin`'s emulator: `rabi` hertz per volt). `spin`'s receivers demodulate with the transmitter's own oscillator, so they look at $\boldsymbol{M}$ in exactly this frame.

On resonance ($\Delta = 0$) a pulse of length t_p turns $\boldsymbol{M}$ about its axis by the flip angle $\theta = 2\pi f_\mathrm{R} t_p$: a 90° pulse along x turns $+\hat{z}$ into $-\hat{y}$ (`spin`'s sign convention), a 180° pulse inverts it. Off resonance the axis tilts out of the transverse plane by $\arctan(\Delta/f_\mathrm{R})$ and the rotation is faster, at $f'_\mathrm{R} = \sqrt{f_\mathrm{R}^2 + \Delta^2}$: for $\Delta = f_\mathrm{R}$ the tilt is 45° and $\boldsymbol{M}$ never gets further than $M_z = 0.00$ from the pole (figure 1). Starting from $+\hat{z}$, a pulse that turns by $\beta = 2\pi f'_\mathrm{R} t_p$ about the tilted axis leaves

$$|M_\perp| = \frac{f_\mathrm{R}}{f'_\mathrm{R}}\sqrt{\sin^2\beta + \frac{\Delta^2}{f'^2_\mathrm{R}}\,(1-\cos\beta)^2}, \qquad M_z = 1 - \frac{f_\mathrm{R}^2}{f'^2_\mathrm{R}}\,(1-\cos\beta), \tag{3}$$

and the population of the upper state is Rabi's formula [10]

$$P_1(t) = \frac{f_\mathrm{R}^2}{f_\mathrm{R}^2 + \Delta^2}\,\sin^2\!\left(\pi\sqrt{f_\mathrm{R}^2 + \Delta^2}\,t\right). \tag{4}$$

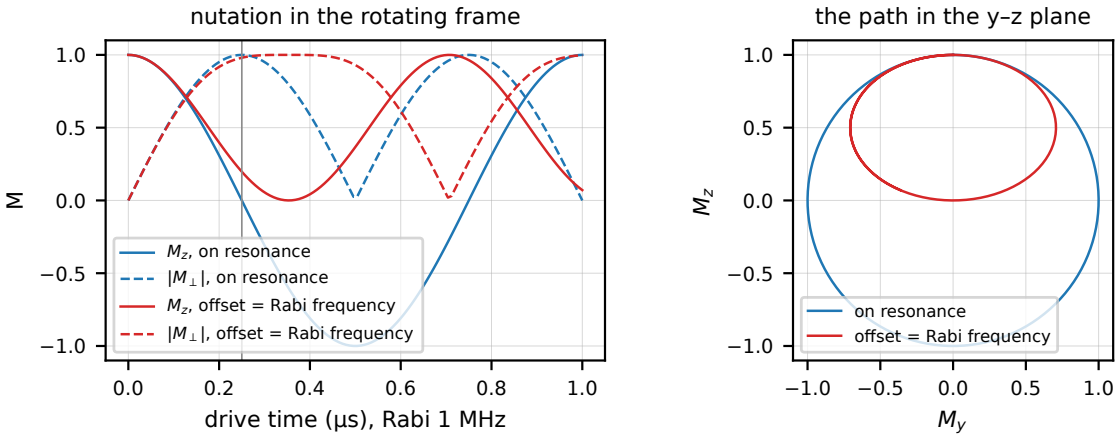


Figure 1: A drive of 1 MHz Rabi frequency along x, on resonance and 1 MHz off resonance (computed): M_z and $|M_\perp|$ against time, and the path in the y–z plane. On resonance the vector passes the transverse plane at 0.25 µs (the 90° pulse); off resonance it turns faster about a tilted axis and never reaches the plane.

2.3 Relaxation

Bloch added two phenomenological decays [9]:

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = (\boldsymbol{\Omega}\times\boldsymbol{M})_x - \frac{M_x}{T_2}, \quad \frac{\mathrm{d}M_y}{\mathrm{d}t} = (\boldsymbol{\Omega}\times\boldsymbol{M})_y - \frac{M_y}{T_2}, \quad \frac{\mathrm{d}M_z}{\mathrm{d}t} = (\boldsymbol{\Omega}\times\boldsymbol{M})_z - \frac{M_z - M_\mathrm{eq}}{T_1}. \tag{5}$$

T_1 returns the populations to equilibrium (exchanging energy with the lattice); T_2 destroys the phase coherence between the spins (flip-flops, fluctuating local fields). Every physical system has

$$T_2 \le 2\,T_1, \tag{6}$$

because the processes that relax M_z also destroy coherence: in the quantum description the density matrix stays positive only then, and the Bloch vector stays inside the unit sphere. With $T_2 > 2T_1$ the equations still integrate, but $|\boldsymbol{M}|$ grows beyond 1, up to 1.41 for $T_2 = \infty$ (figure 2). `spin`'s emulator refuses such settings (its vector's components are clamped to ±1, which is exact inside the ball only); its `t2` null means $2T_1$, relaxation without extra dephasing.

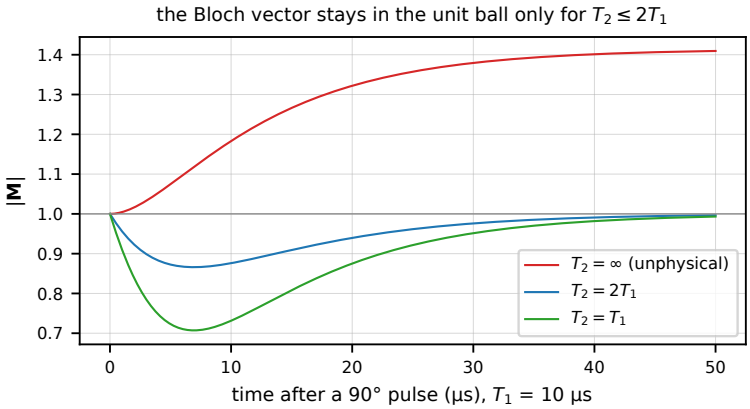


Figure 2: $|\boldsymbol{M}|$ after a 90° pulse with $T_1 = 10\,\mu\text{s}$ (computed): it grows beyond 1 only for the unphysical $T_2 > 2T_1$.

2.4 Why $T_2 \leq 2T_1$

For a spin 1/2 (or a qubit) the state is a density matrix $\rho = \frac{1}{2}(\mathbb{1} + \boldsymbol{M} \cdot \boldsymbol{\sigma})$, with the Pauli matrices $\boldsymbol{\sigma}$; its eigenvalues $\frac{1}{2}(1 \pm |\boldsymbol{M}|)$ are probabilities, so $|\boldsymbol{M}| \leq 1$. Relaxation is the sum of two processes. Energy relaxation (amplitude damping) at the rate $1/T_1$ moves population to the ground state and, because a coherence lives between the two levels, also damps the coherences at $1/(2T_1)$. Pure dephasing (random phase kicks that leave the populations alone) adds a rate $1/T_\varphi \geq 0$. Hence

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\varphi} \;\geq\; \frac{1}{2T_1}, \tag{7}$$

with equality when energy relaxation alone limits the coherence: the case `spin`'s emulator takes for `t2` null.

2.5 The qubit's Bloch sphere

A qubit is a two-level system $\alpha|0\rangle + \beta|1\rangle$; its Bloch vector $(2\,\mathrm{Re}\,\alpha^*\beta,\ 2\,\mathrm{Im}\,\alpha^*\beta,\ |\alpha|^2 - |\beta|^2)$ obeys equation (2) under a resonant drive, with T_1 the energy relaxation and T_2 the coherence time. A measurement projects it: outcome 1 with probability $(1 - M_z)/2$. Rabi oscillations (eq. 4) and Ramsey fringes [11] (two $\pi/2$ pulses separated by a free precession at Δ: the population oscillates at Δ) are the qubit's versions of nutation and the FID. Figure 3 maps Rabi's formula over the drive's time and the offset: the oscillation speeds up and shrinks away from resonance, the chevron that calibrates a qubit's frequency.

Exercises

1. *Nutation.* With the ensemble above, run `X.nutation` over 0.05 … 2 µs at 0.8 V; fit `A.fit_oscillation`. Expect $f_\mathrm{R} = 1.00\,\mathrm{MHz}$ (`rabi` 1.25 MHz/V × 0.8 V). Halve the amplitude: the frequency halves.
2. *Off resonance.* Move the emulator's `frequency` 1 MHz above TX1 and repeat: the oscillation is faster ($\sqrt{2}\,\mathrm{MHz}$) and its amplitude smaller, as in figure 1.
3. *The unit ball.* Try `"t1": 10e-6, "t2": 50e-6`: the driver refuses it. Why must $T_2 \leq 2T_1$?
4. *Rabi and Ramsey on the qubit.* With the qubit (`mode qubit`, the preset `rabi-qubit`), verify equation (4) on resonance; with the qubit 500 kHz below TX1 (`ramsey-qubit`) the Ramsey fringes oscillate at 500 kHz.

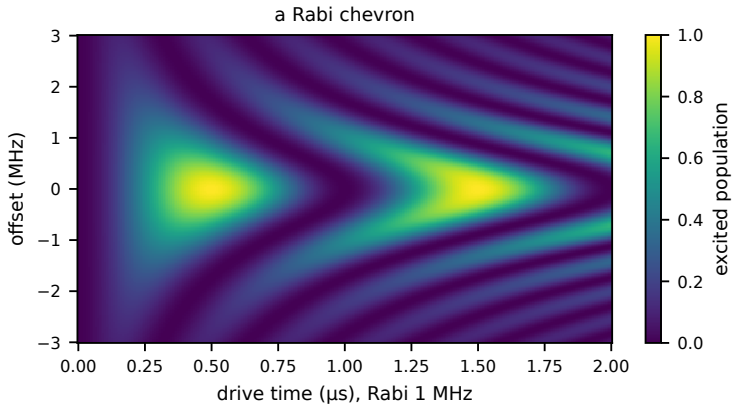


Figure 3: The excited population (eq. 4) against the drive's time and the offset, at 1 MHz Rabi frequency (computed).

3 Pulses

3.1 Hard pulses and their band

A rectangular pulse is short and strong: f_R large against the offsets it must cover. How large must it be? A pulse of length t_p has a spectrum of width about $1/t_p$, but the spins' response is not linear in the drive; figure 4 computes it with the Bloch equations: a nominal 90° pulse of 0.25 µs at 1 MHz Rabi frequency tips at least half the magnetization into the plane over 5.50 MHz (full width at half height), with side lobes far out; a Gaussian envelope of the same area over 1 µs covers only 2.25 MHz but with nearly no side lobes. A spectrum wider than the pulse's band is excited unevenly (and with an offset-dependent phase: the first-order phase of section 4); a narrow pulse selects a slice: ESR's field sweeps and frequency-selective NMR use that.

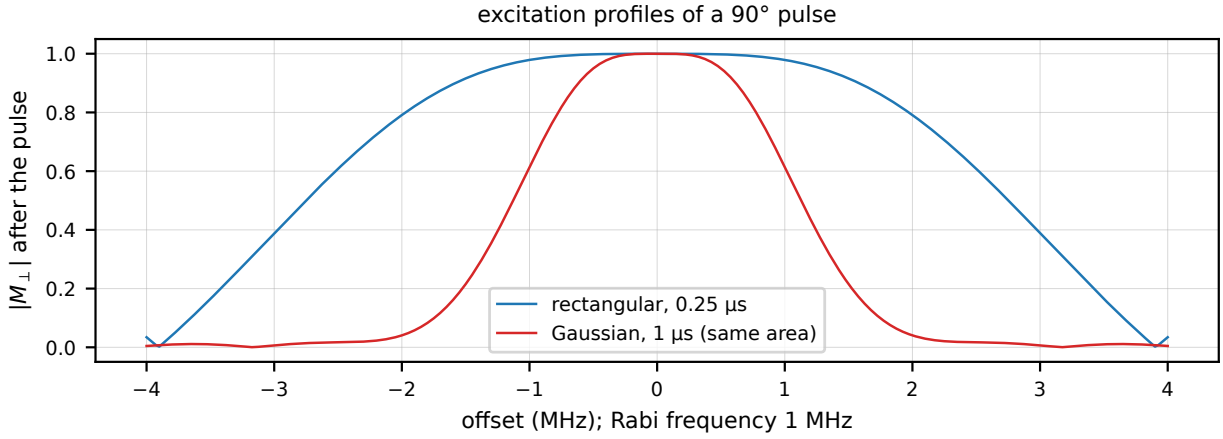


Figure 4: The transverse magnetization after a nominal 90° pulse against the offset (computed, Bloch equations): a rectangular pulse of 0.25 µs at 1 MHz Rabi frequency, and a Gaussian pulse of the same area over 1 µs.

`spin`'s shapes give any envelope in amplitude and phase (31 shapes on a shared 4096-point table, interpolated at every sample); the emulator drives its spin with the envelope's amplitude and phase exactly as a real coil would.

3.2 Calibration and B_1 inhomogeneity

The flip angle is calibrated by nutation: the signal after a pulse goes as $\sin\theta$; the first maximum is the 90° pulse, the first zero the 180° (which is sharper, so it is the better calibration). In a real coil B_1 varies across the sample: the nutation damps, the 180° pulse never nulls the signal completely, and composite pulses or adiabatic pulses become useful.

3.3 Chirps and adiabatic inversion

A pulse whose frequency sweeps through the resonance inverts the magnetization *adiabatically* when the effective field turns slowly against its own size: in the frame of the sweep, $\boldsymbol{M}$ follows the effective field from $+\hat{z}$ to $-\hat{z}$. For a linear sweep at $\dot{f}$ (Hz/s) the Landau–Zener formula gives the probability of *not* following, $\mathrm{e}^{-\pi Q/2}$, with the adiabaticity $Q = 2\pi f_\mathrm{R}^2/\dot{f}$ (figure 5) [12]: an inversion to 99 % needs $Q \approx 3.37$. Above that the inversion is insensitive to B_1 (any f_R large enough works) and to the offset (anywhere within the sweep): the reason adiabatic pulses dominate where B_1 is inhomogeneous. `spin`'s chirp registers give linear sweeps from 3.5 mHz/s to 4.9×10^{14} Hz/s, with the envelope from a shape.

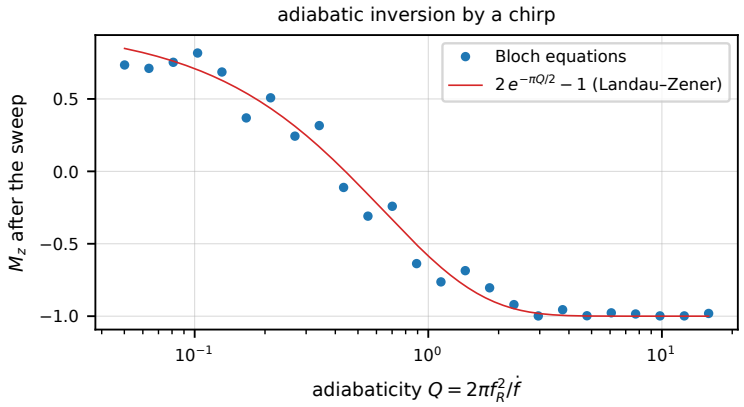


Figure 5: M_z after a linear chirp through resonance against the adiabaticity (computed: Bloch equations, a 10 MHz sweep at 0.5 MHz Rabi frequency), with the Landau–Zener formula.

Exercises

1. *Excitation profile.* Keep TX1 at 10 MHz and step the emulator's `frequency` from 6 to 14 MHz (it must stay within 3 MHz of TX1: use 7 … 13 MHz) with `detuning none`; record the FID's first point after a 0.25 µs pulse at 0.8 V. Compare with figure 4.
2. *A soft pulse.* Repeat with a Gaussian shape of 1 µs and the amplitude raised so that its area gives 90°: the band narrows to about 2.25 MHz.
3. *Adiabatic inversion.* With a chirp register at 10^{12} Hz/s and a 4 µs pulse that starts 2 MHz below the spin (a 4 MHz sweep, inside the emulator's ±3 MHz), measure M_z by a following 90° pulse; vary the amplitude: inversion becomes complete once $Q = 2\pi f_\mathrm{R}^2/\dot{f}$ exceeds about 3 ($f_\mathrm{R} = 1$ MHz gives $Q = 6.3$, 0.5 MHz only 1.6).

4 The free induction decay and its spectrum

4.1 Lines and decays

After a 90° pulse each spin precesses at its own offset; the signal is their sum, the free induction decay (FID). Its envelope is the Fourier transform of the distribution of offsets (the line), times the irreversible e^{-t/T_2}. A Lorentzian line of half width Γ (Hz) decays as e^{-t/T_2^*} with $T_2^* = 1/(2\pi\Gamma)$ (5 kHz: 31.8 µs); a Gaussian line of rms width σ as $\exp(-(2\pi\sigma t)^2/2)$ (full width at half height $2\sqrt{2\ln 2}\,\sigma$:

5 kHz rms is 11.77 kHz). Field inhomogeneity across the sample, chemical shifts and unresolved couplings make up the line; $T_2^* < T_2$ always.

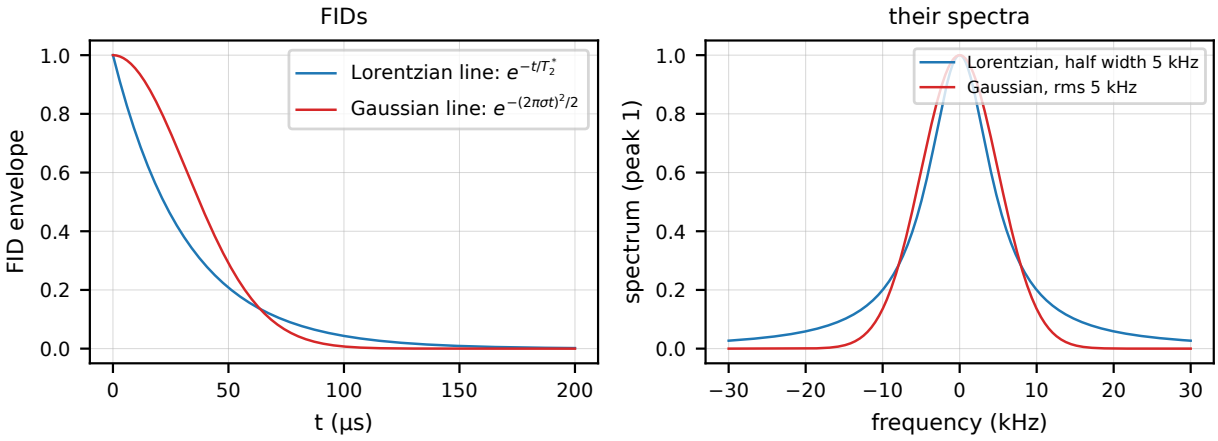


Figure 6: FIDs and spectra of a Lorentzian and a Gaussian line of the same width parameter (computed).

4.2 Sampling, quadrature and the spectrum

`spin`'s receiver mixes the input with the transmitter's oscillator (both cos and sin: a complex, *quadrature* signal), so positive and negative offsets are told apart, and samples it every dwell time t_d. The spectrum (the discrete Fourier transform of the complex FID [13]) spans $\pm 1/(2t_d)$ around the transmitter's frequency; an offset beyond the band folds back (aliases). Three habits make clean spectra: weight the first point by ½ (the discrete transform counts it twice otherwise, raising the baseline), zero-fill to at least twice the length (the spectrum is then interpolated), and phase it.

Phase. The FID's start has a phase from the receiver's and the pulse's phases (zero order: one angle for the whole spectrum, `A.phase`) and, because acquisition starts a dead time t_{dead} after the pulse, a phase $2\pi\Delta\, t_{\mathrm{dead}}$ that grows with the offset (first order). The real part of a phased spectrum is the absorption line. For a dead time of 10 µs a line 20 kHz off resonance is turned by 72° (modulo 360°); the first-order correction $\mathrm{e}^{-2\pi \mathrm{i} f t_{\mathrm{dead}}}$ restores every line at once (figure 7).

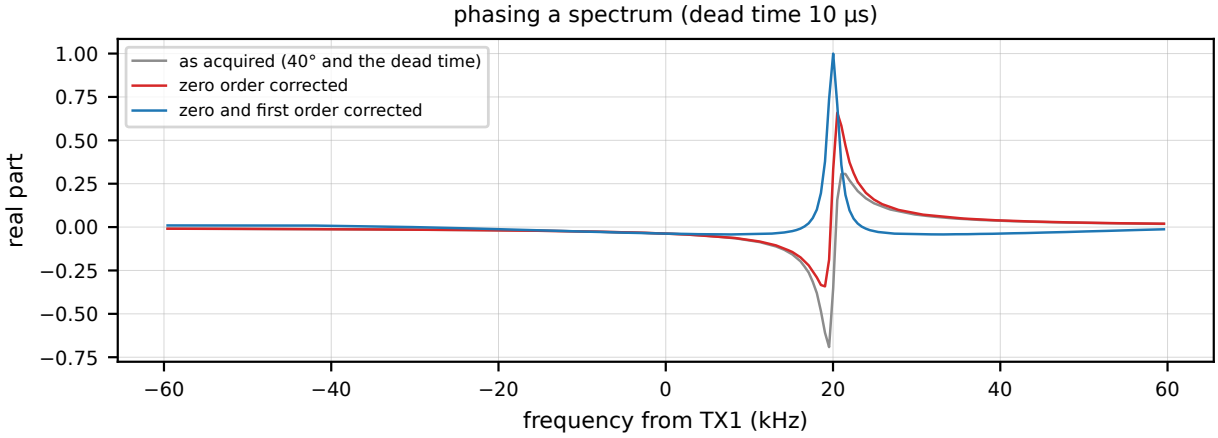


Figure 7: The real part of a spectrum as acquired (a 40° receiver phase and a 10 µs dead time), after the zero-order correction, and after both (computed).

4.3 The receiver's filter

Each point of `spin`'s receiver is a triangle-weighted mean of $2R-1$ samples (an order-2 CIC filter [14], R the decimation, $t_d = 8\,\text{ns} \times R$). Its response,

$$H(f) = \left(\frac{\sin(\pi f R\, t_s)}{R \sin(\pi f t_s)}\right)^2, \qquad t_s = 8\,\text{ns}, \tag{8}$$

is 1 at the centre and falls towards the band's edge, 0.405 ($-7.8\,\text{dB}$) at $\pm 1/(2t_d)$ for $R = 62$ (figure 8); `A.spectrum` divides it out. Its zeros sit at multiples of the output rate, so it rejects the frequencies that would alias onto the centre; at 1.5 times the output rate it still passes $-26.9\,\text{dB}$, which folds onto the band's edge. Keep the line well inside $\pm 1/(2t_d)$.

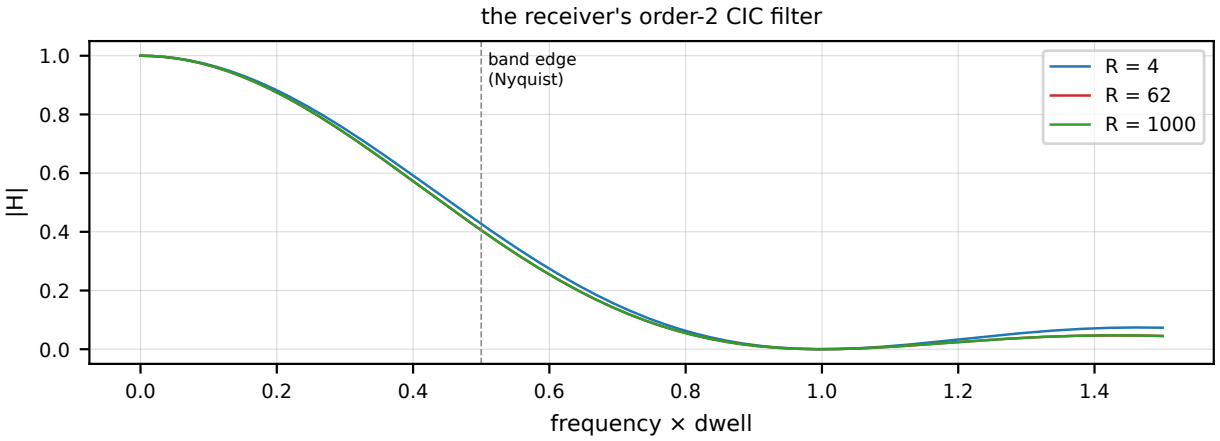


Figure 8: The receiver's order-2 CIC filter against the frequency in units of the output rate, for three decimations (computed).

4.4 One isochromat per shot

`spin`'s emulator draws one offset from the line at every shot; a single shot's FID does not decay (one spin precesses forever, except for T_2), and the average of N shots approaches the ensemble's FID with a residue: N random phasors of length 1 sum to about $\sqrt{N}$, so a fraction near $\sqrt{\pi/(4N)}$ remains after the decay. With 256 shots that is about 5 %. The spectrum is the histogram of the drawn offsets, broadened by T_2. A real sample's 10^{20} spins leave no such residue.

Exercises

1. *Lineshapes.* Record FIDs of a Gaussian and a Lorentzian line of 5 kHz (the emulator's `detuning`), with $T_2 = 1\,\text{ms}$ and 256 scans; fit the envelope's decay and compare with $T_2^* = 31.8\,\mu\text{s}$ for the Lorentzian and the Gaussian's $\exp(-(2\pi\sigma t)^2/2)$.
2. *Aliasing.* With dwell 2 µs (band ±250 kHz) move the emulator 300 kHz above TX1: the line appears at −200 kHz, attenuated by the CIC's response.
3. *The residue.* Compare FIDs of 16, 64 and 256 shots after the decay: the residue falls as $1/\sqrt{N}$.
4. *First-order phase.* Lengthen the dead time from 2 to 10 µs with a line 20 kHz off resonance: the spectrum's phase changes by $2\pi \times 20\,\text{kHz} \times 8\,\mu\text{s} \approx 58°$.

5 Spin echoes and echo trains

5.1 The Hahn echo

The FID's decay with T_2^* is mostly reversible: each spin keeps its own offset. A 180° pulse at time τ turns every phase φ into $-\varphi$; the spins keep precessing at their offsets and come back into phase

at 2τ: the spin echo [15]. Only the irreversible part is lost, so the echo's height is the signal times $e^{-2\tau/T_2}$ (figure 9: the computed top 0.607 against $e^{-2\tau/T_2} = 0.607$). Plotting echo heights against 2τ measures T_2 through any inhomogeneity. With the refocusing pulse along y (90x–τ–180y–τ) the echo forms along the same axis as the FID would have.

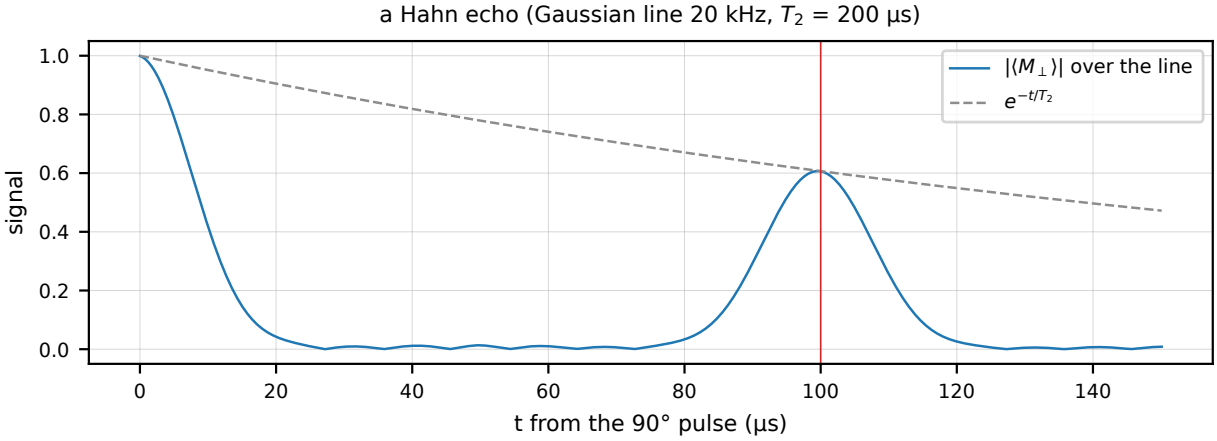


Figure 9: A Hahn echo of a Gaussian line (rms 20 kHz) with $T_2 = 200\,\mu s$, $\tau = 50\,\mu s$ (computed: 64 isochromats over the line's quantiles).

Diffusion. Spins that move in a field gradient G change their offset between dephasing and rephasing: the echo decays faster, as $\exp(-\frac{2}{3}\gamma^2 G^2 D\tau^3)$ on top of T_2 (Torrey [16]). Short echo spacings suppress it; that is one reason for trains.

5.2 CPMG

Carr and Purcell [17] repeated the 180° pulse every 2τ and recorded every echo in one scan. Their pulses along x accumulate the flip-angle error from echo to echo. Meiboom and Gill [18] put the refocusing pulses along y, the axis of the refocused magnetization: an error then only tilts the vector about its own axis, and alternate echoes undo it. Figure 10 shows 170° pulses: after 24 echoes CP keeps 0.34 of the signal, CPMG 1.00. `spin`'s `X.cpmg` uses 180°y pulses and takes each echo into its own averager region.

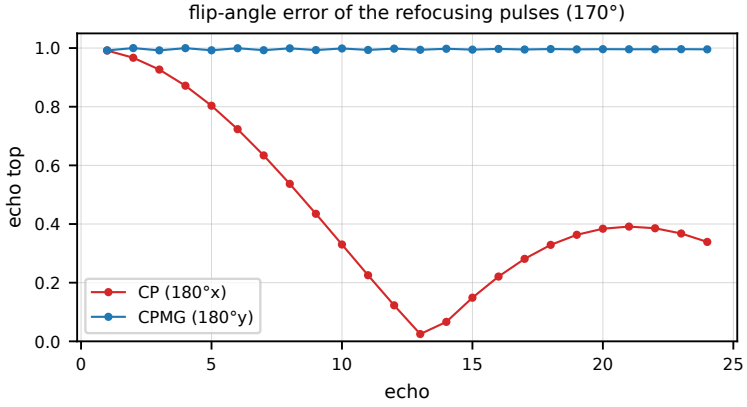


Figure 10: Echo trains with 170° refocusing pulses, 20 µs apart, on a Gaussian line (computed): CP (along x) against CPMG (along y).

5.3 Spin-locked trains

With refocusing pulses smaller than 180° the echoes mix spin echoes with stimulated echoes (from magnetization stored along z), and after a few echoes a pseudo-steady state forms whose height depends on each spin's offset. In NQR the multiple-pulse sequence with 90° pulses, the spin-locked spin echo (SLSE) [19], holds a quadrupolar powder's signal for an effective time $T_{2\text{eff}}$ much longer than T_2, which is how weak NQR lines are found. That effect belongs to spins $I > 1/2$ with several transitions in a powder. For a spin 1/2, like `spin`'s emulator, the 90° train's echoes are offset-dependent and its average over a line depends on where the line sits: the user's manual shows a measured train against the Bloch equations.

5.4 One isochromat, two signs of offset

With strong longitudinal relaxation during an echo even a single isochromat's echo depends on the sign of its offset (figure 11): M_z regrows during τ (toward $+\hat{z}$), and the finite 180° pulse carries it through the transverse plane while the spin precesses at its offset, leaving a transverse residue whose sign follows the offset. With $T_1 = 10\,\mu\text{s}$, $T_2 = 20\,\mu\text{s}$ and $\tau = 10\,\mu\text{s}$ an isochromat 23.5 kHz below the transmitter gives 0.395, one 23.5 kHz above 0.333. A line's echo averages over both, but a few shots of a single isochromat each (the emulator's ensemble) scatter accordingly.

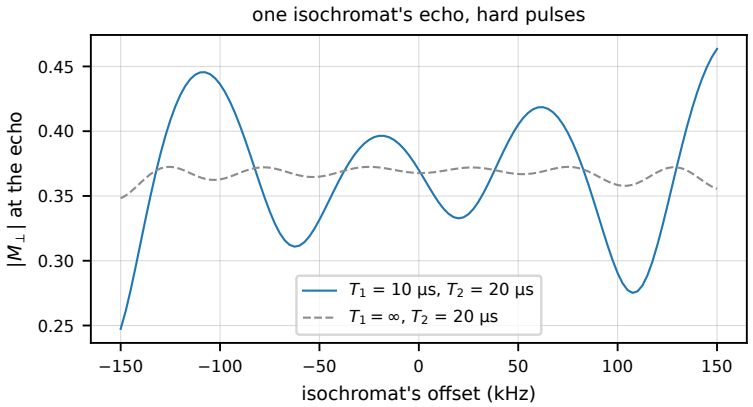


Figure 11: One isochromat's echo height against its offset (computed: hard 0.25/0.5 µs pulses, $\tau = 10\,\mu\text{s}$), with and without T_1.

Exercises

1. T_2 *through* T_2^*. With a 20 kHz Gaussian line ($T_2^* \approx 8\,\mu\text{s}$) and $T_2 = 300\,\mu\text{s}$, run `X.t2` for $\tau = 50 \ldots 600\,\mu\text{s}$; `A.fit_decay` of the echo tops against 2τ gives T_2.
2. *CP against CPMG.* Mis-set the refocusing pulse (a t_{90} 6 % short) and compare `X.cpmg` with a CP train built by hand (refocusing phase register 0): the CP train decays, CPMG holds (figure 10).
3. *A 90° train.* Run `X.slse` and `X.cpmg` with $\tau = 10\,\mu\text{s}$ on the line above, then move the line onto resonance (`frequency` 10 MHz): the 90° train changes much more than CPMG.

6 Relaxation measurements

6.1 T_1

After an inversion (a 180° pulse) M_z recovers as $M_0(1 - 2\mathrm{e}^{-t/T_1})$; a 90° pulse at the delay t reads it. The signal crosses zero at $t = T_1 \ln 2 = 0.693\,T_1$, a quick estimate. Fitting $A(1 - 2B\mathrm{e}^{-t/T_1})$ allows an imperfect inversion ($B < 1$). Inversion recovery needs the magnetization at equilibrium before every scan: $5\,T_1$ apart. *Saturation recovery* destroys M_z with a comb of 90° pulses instead, and recovers as $M_0(1 - \mathrm{e}^{-t/T_1})$ from zero, whatever came before; scans can follow each other quickly (figure 12).

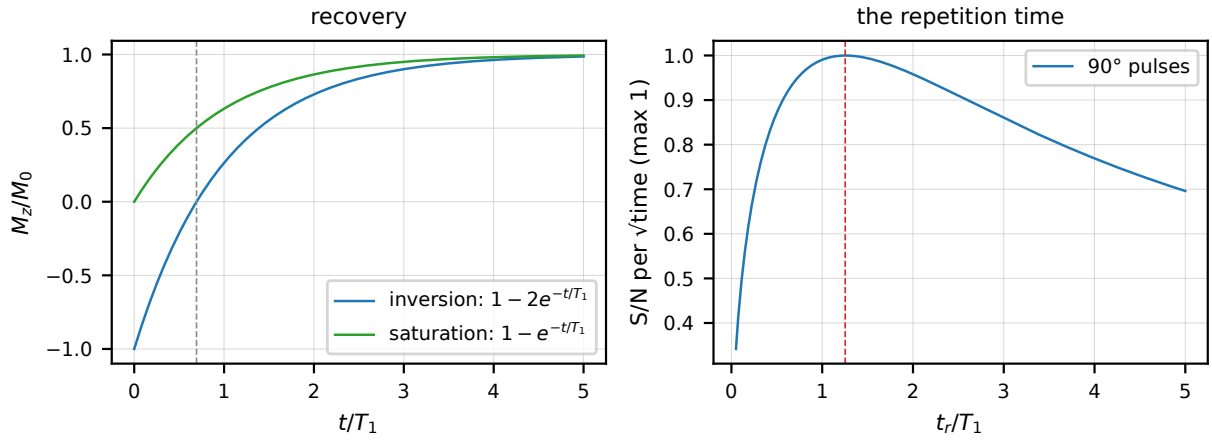


Figure 12: Left: inversion and saturation recovery. Right: the signal per square root of time of a train of 90° pulses against the repetition time (computed).

6.2 The repetition time

Scans repeated every t_r reach a steady state: with 90° pulses each starts from $M_0(1-\mathrm{e}^{-t_r/T_1})$. Averaging N scans in a fixed time $T = Nt_r$ improves the noise as $\sqrt{N}$, so the signal-to-noise per square root of time goes as $(1-\mathrm{e}^{-t_r/T_1})/\sqrt{t_r}$, largest at $t_r \approx 1.25\,T_1$; waiting $5\,T_1$ costs a factor 0.70. With shorter t_r a smaller flip angle does better: the Ernst angle $\cos\alpha = \mathrm{e}^{-t_r/T_1}$ [13] (42° for $t_r = 0.3\,T_1$). In steady-state experiments, dummy scans (scans without acquisition) bring the magnetization to its steady state first: every `spin` template takes `dummies`.

6.3 T_2 and $\mathrm{T}_{1\rho}$

T_2 comes from echoes (section 5). A long pulse along the magnetization (spin-locking) holds it in the rotating frame, where it relaxes with $T_{1\rho}$, sensitive to slow motions; with `spin`'s `keep` a pulse can be continued for milliseconds (watch the duty guard).

Exercises

1. *The null.* With $T_1 = 1$ ms run `X.inversion_recovery` over 0.1 … 5 ms; the signal crosses zero near 0.69 ms; `A.fit_recovery` gives T_1 and B.
2. *Saturation recovery.* Run `X.saturation_recovery` with $t_r = 0.5$ ms: the same T_1 in a quarter of the time.
3. *The optimum t_r.* Average FIDs with $t_r = 0.5$, 1.25, 2.5 and 5 ms for the same total time and compare the signal-to-noise (the emulator's `noise` at 10 mV): the best is near 1.25 ms.

7 Phase cycling

7.1 What a cycle selects

Shifting the phase of a pulse by ϕ multiplies the part of the signal that passed through a coherence change Δp at that pulse by $\mathrm{e}^{-\mathrm{i}\Delta p\,\phi}$ (the coherence order $p = 0$ for z magnetization, ± 1 for transverse; detection sees $p = -1$). Repeating an experiment with a set of phases and adding the scans with matching receiver phases keeps the pathways whose phases cancel the receiver's and removes the others [20]. Two cycles cover most of `spin`'s templates.



Figure 13: Coherence pathways of a Hahn echo: the 90° pulse creates $p = \pm 1$, the 180° pulse turns $+1$ into -1 (the echo, detected) and -1 into $+1$. A cycle of the 180° pulse's phase with the receiver at twice its step keeps the first.

7.2 Receiver errors and CYCLOPS

A real receiver adds a DC offset (a spike at the transmitter's frequency) and, with an I/Q pair, a gain and a phase error between I and Q, which turn a line at $+\Delta$ into an image at $-\Delta$: for a gain ratio g and a phase error ϵ the image is suppressed by

$$\mathrm{IRR} = 10 \log_{10} \frac{1 + 2g\cos\epsilon + g^2}{1 - 2g\cos\epsilon + g^2} \tag{9}$$

(26 dB for 5 % and 5°). CYCLOPS [21] steps all pulses and the receiver through 0°, 90°, 180°, 270° together: the signal adds in every step, while the offset and the image, which do not turn with the transmitter, cancel exactly in the four-step sum (figure 14; what remains at the image's frequency is the line's own wing, −31 dB). The two-step cycle 0°/180° removes the offset only. `spin` applies the receiver's phase exactly, in 90° steps, before averaging, so the cycle is subtracted on the board.

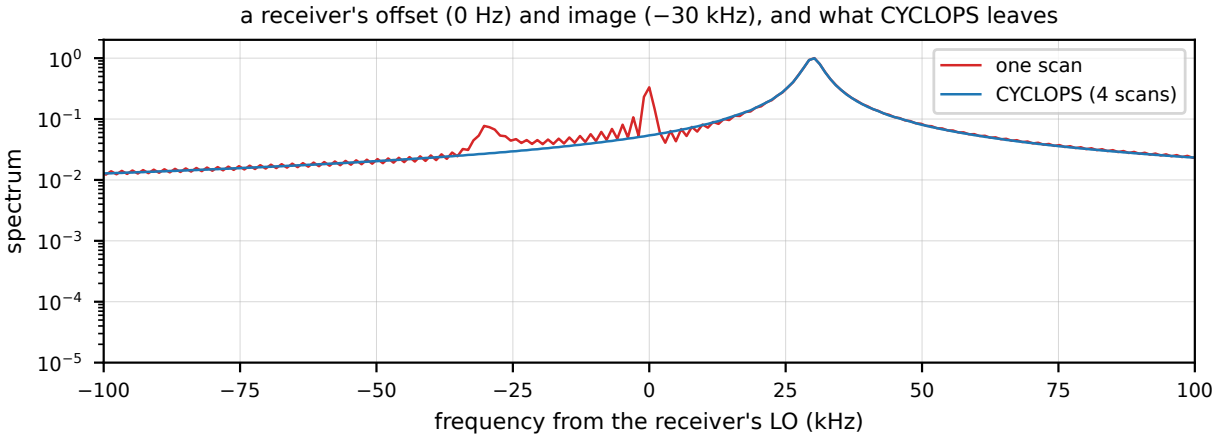


Figure 14: A receiver with an offset and a 5 % / 5° I/Q imbalance (computed): one scan's spectrum shows the offset at 0 Hz and the image at −30 kHz; the CYCLOPS sum removes both.

7.3 Echoes: EXORCYCLE

An imperfect 180° pulse also leaves FID signal and unwanted pathways in an echo. EXORCYCLE [22] steps the 180° pulse through 0°, 90°, 180°, 270° with the receiver at 0°, 180°, 0°, 180°: the echo pathway (p: $+1 \to -1$, $\Delta p = -2$ at the 180°) adds, the others cancel. With `spin` the refocusing pulse takes its phase from a phase register (or a 16-bit phase from a table) and each step's window its receiver phase.

Exercises

1. *CYCLOPS with cables.* In I/Q mode, cable OUT1 → IN1 and OUT2 → IN2 and use input `iq`; give TX1 a CYCLOPS-cycled pulse and TX2 an uncycled tone at the same frequency: the averaged window keeps TX1 and cancels TX2 (the board test does this, below 0.2 %).
2. *An I/Q imbalance.* With the same cables, set TX1's `q_gain` to 1.05 and `q_phase` to 5° and look at

one scan's spectrum: the image at $-\Delta$ is about 26 dB down; CYCLOPS removes it.

3. *EXORCYCLE.* Build a Hahn echo whose four unrolled steps (as the templates unroll their cycles) give the 180° pulse the phase registers 0°, 90°, 180°, 270° and the window the receiver phases `rx_phase` 0, 2, 0, 2; mis-set the 180° pulse and compare with a single step.

8 Nuclear quadrupole resonance

8.1 The quadrupole interaction

A nucleus with spin $I \geq 1$ has an electric quadrupole moment eQ; in a crystal it sits in an electric field gradient (EFG) whose principal values $V_{zz} \geq V_{yy} \geq V_{xx}$ define the coupling e^2qQ/h ($eq = V_{zz}$) and the asymmetry $\eta = (V_{xx} - V_{yy})/V_{zz}$. The interaction splits the nuclear levels without any magnet [5]:

- $I = 3/2$ (^{35}Cl, ^{63}Cu, ^{75}As): one line at $\nu_Q = \frac{e^2qQ}{2h}\sqrt{1 + \eta^2/3}$, the $\pm 1/2 \leftrightarrow \pm 3/2$ transition;
- $I = 1$ (^{14}N): three lines, $\nu_\pm = \frac{3e^2qQ}{4h}(1 \pm \eta/3)$ and $\nu_0 = \nu_+ - \nu_-$.

Typical frequencies: ^{14}N 0.5–5 MHz, ^{35}Cl 30–40 MHz in organic chlorides: `spin`'s range. The EFG depends on the lattice, so the lines move with temperature (several kHz per kelvin for ^{35}Cl) and with pressure: a thermometer, and a fingerprint of a compound.

8.2 Powders and the flip angle

In a powder each crystallite's EFG axis lies at its own angle θ to the coil's field. For $I = 3/2$, $\eta = 0$, the transition nutates at $\sqrt{3}\,\omega_1 \sin\theta$ and the coil detects $\propto \sin\theta$; averaged over orientations (solid angle $\sin\theta\,\mathrm{d}\theta$),

$$S(t_p) \propto \int_0^{\pi/2} \sin\!\left(\sqrt{3}\,\omega_1 t_p \sin\theta\right) \sin^2\theta\,\mathrm{d}\theta, \tag{10}$$

which peaks at an effective flip angle $\sqrt{3}\,\omega_1 t_p \approx 102°$ instead of 90° and never nulls (figure 15). Calibrate a powder's pulse by the largest echo.

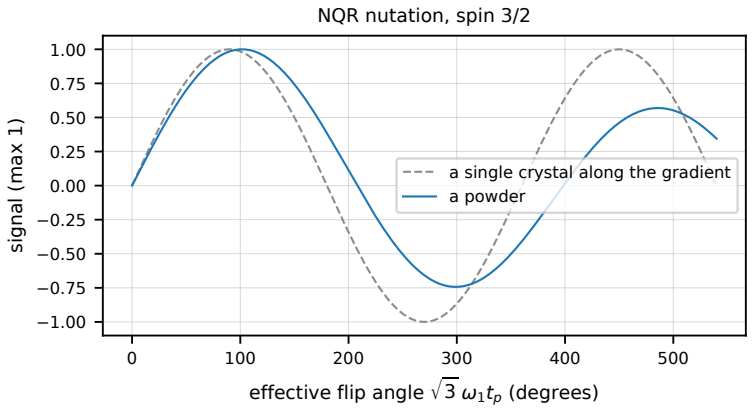


Figure 15: NQR nutation of a spin 3/2 ($\eta = 0$): a single crystal along the gradient against a powder (computed with the model above).

8.3 Finding a line

An unknown line is searched for: step the transmitter by less than a pulse's band (about $1/t_p$: 100 kHz for 10 µs pulses), and at each frequency record a spin-locked echo train (SLSE), which in a powder lasts $T_{2\mathrm{eff}} \gg T_2$ [19]; sum its echoes. A line appears in a few steps' sums; then refine the

frequency, the pulse and the echo spacing. Repetition is limited by T_1, often tens of milliseconds to seconds.

The emulator and NQR. The emulator is a spin 1/2: it cannot show the quadrupolar powder's nutation or its long trains. It does let students practise the search: hide the line by setting the emulator's `frequency` unknown to them, within the transmitter's reach.

Exercises

1. *A search.* A partner sets the emulator's ensemble somewhere in 9.5 … 10.5 MHz (a 2 kHz line, $T_1 = 20$ ms, $T_2 = 10$ ms; scans 0.1 s apart). With 2 µs pulses (band about 500 kHz) step TX1 by 200 kHz (the emulator follows TX1 within 3 MHz), record `X.slse` trains and sum their echoes; then refine.
2. *The powder's optimum.* Integrate $S(t_p)$ above numerically and find its maximum (102°); repeat with a detection weight of 1 instead of $\sin\theta$: the optimum moves.

9 Pulsed ESR through a microwave bridge

9.1 Electron spins at microwave frequencies

An electron's moment is 658 times a proton's: $g \approx 2$ gives 28.0 GHz/T, so ESR works at microwave frequencies (X band: 9.5 GHz at 0.34 T) with relaxation times of nanoseconds to microseconds and lines that are usually broad (anisotropic g, hyperfine couplings) [6]. `spin` cannot make 9.5 GHz; a microwave bridge moves its signals there and back.

9.2 Single-sideband up- and down-conversion

The bridge's IQ modulator multiplies `spin`'s I/Q pair (OUT1, OUT2) by the microwave LO's cos and sin: the output is the LO shifted by each transmitter's frequency, a single sideband (a negative frequency selects the lower). An IQ demodulator brings the echo back to IN1 + i IN2. Two errors of real mixers matter: an imbalance between I and Q leaves an image at $f_{\mathrm{LO}} - f$ (suppressed by the IRR of section 7; figure 16: 1 % and 1° give 40 dB), and LO leakage leaves a line at f_{LO}. `spin` corrects both at the source: each transmitter's Q-path gain and phase (`q_gain`, `q_phase`) null its image at its own frequency, and each output's DC level nulls the leakage; CYCLOPS removes what is left from the averages.

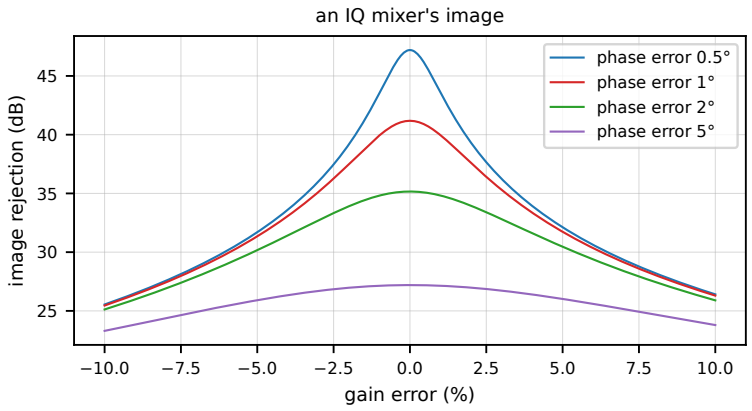


Figure 16: An IQ mixer's image rejection against its gain error, for four phase errors (computed).

9.3 Echoes and field sweeps

After a microwave pulse the resonator rings and the receiver recovers for 100 ns or more, longer than many ESR FIDs: pulsed ESR records echoes. ESR lines are often wider than any pulse's band, so a

spectrum is taken point by point: the *echo-detected field sweep* steps the magnetic field and records the echo's height at each field. A pulse excites a slice of width about $1/t_p$; the swept spectrum is the line convolved with that slice, so softer pulses resolve better at the cost of signal. `spin`'s `field_sweep` steps a slow output (driving the field supply) from a table, waits for it to settle, and records a Hahn echo; the emulator's `ao` moves its spins with AO0 so the experiment runs without a magnet.

Exercises

1. *A field sweep.* With the emulator 200 kHz above TX1, `ao` −1 MHz/V and a 50 kHz line, sweep AO0 over ±0.2 V around 0.9 V: the line's centroid lies where the spins meet TX1 (0.9 V + 0.2 MHz / 1 MHz/V = 1.1 V).
2. *Resolution against signal.* Repeat with t_{90} = 0.5, 2 and 8 µs (the amplitude scaled to keep 90°): the swept line narrows towards the true 50 kHz and the echo weakens.
3. *The image (with cables).* In I/Q mode, loop OUT1/OUT2 to IN1/IN2, mis-set `q_phase` by 1° and `q_gain` by 1 %: the image is about 40 dB down; trim it back.

10 Qubit readout and decisions

10.1 Dispersive readout

A qubit coupled to a resonator (a Jaynes–Cummings system [23]) far from resonance shifts the resonator's frequency by $\pm\chi$ according to its state, without exchanging energy with it: the dispersive regime [8]. A readout pulse near the resonator's frequency returns with a state-dependent phase (and amplitude); after demodulation and integration each shot gives one point in the I/Q plane, near one of two mean values, broadened by the amplifier's noise [7, 24]. `spin`'s emulator returns TX2's readout pulse through a one-pole resonator at $+\chi$ (ground) or $-\chi$ (excited) in TX2's frame.

10.2 The matched filter

For a readout trace $z(t) = s_j(t) + n(t)$ (state j, white noise of variance σ^2 per point), the weighted sum $S = \sum_k w_k z_k$ separates the states best when the weights follow the difference of the two mean traces: $w_k \propto s_1(t_k) - s_0(t_k)$. The separation SNR, $|\langle S\rangle_1 - \langle S\rangle_0|/\sigma_S$, then equals the energy of the difference trace over the noise. A boxcar (equal weights) is optimal only for a constant difference; a resonator's ring-up and decay make the difference trace rise and fall, so matched weights gain. `spin`'s windows take any 1024 complex weights per receiver and integrate at the full rate. Figure 17 computes the gain for `spin`'s emulated readout (a 0.5 µs pulse through a one-pole resonator at ±60°): with a 64 ns ring-down the difference trace is nearly a boxcar and matched weights gain a factor 1.06 in SNR; with 256 ns they gain 1.30, mostly by also integrating the decay after the pulse.

10.3 Thresholds and fidelity

Projecting each shot on the axis between the two means and comparing with a threshold assigns a state; the assignment fidelity is $F = 1 - [P(1|0) + P(0|1)]/2$. For Gaussian blobs at a separation SNR the errors are $\Phi(-\mathrm{SNR}/2)$ each, so $F = \Phi(\mathrm{SNR}/2)$: an SNR of 8 gives 0.99997 (figure 18). Decay during the readout spoils it: a qubit prepared in 1 decays with probability $1 - \mathrm{e}^{-t_{\mathrm{read}}/T_1}$ (2.47 % for 0.5 µs at 20 µs), landing between the blobs, and then about half the time on the wrong side: F falls to about 0.9938 at the same SNR, and the errors become asymmetric (the threshold of fewest errors moves towards the ground blob). Leftover excitation from earlier shots (e^{-t_r/T_1}) and thermal population add errors on the ground side.

10.4 Active reset and repeat-until-success

Waiting $5\,T_1$ resets a qubit passively; measuring and flipping it when found excited resets it faster [25]. The reset's error comes from three places: misassignment, decay between the readout's projection

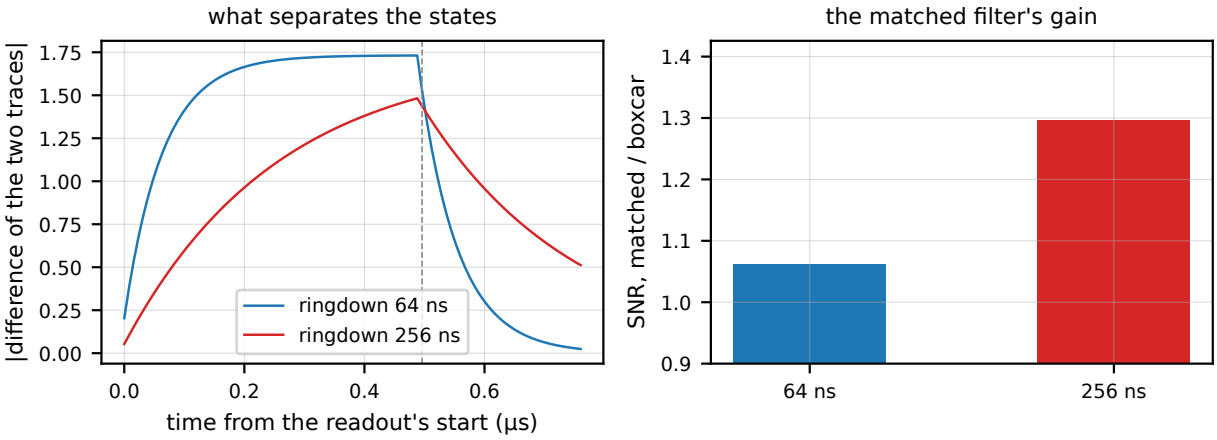


Figure 17: Left: the magnitude of the difference between the two states' readout traces through a one-pole resonator, for two ring-down times. Right: the matched filter's SNR against a boxcar over the pulse (computed).

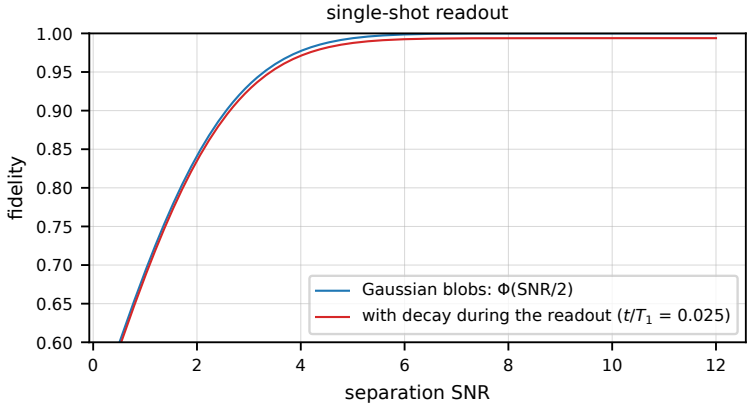


Figure 18: Assignment fidelity against the separation SNR, for Gaussian blobs and with decay during a 0.5 µs readout at $T_1 = 20$ µs (computed).

and the conditional π pulse, and thermal re-excitation. The second dominates when the decision is slow against T_1: a qubit found excited that decays before the π pulse is excited by it. For a half-excited qubit and a projection-to-pulse time t_d the residue is about $p_1(1 - e^{-t_d/T_1})$: with `spin`'s decision (about 1.5 µs from the readout's start to the pulse) and $T_1 = 20$ µs, 0.033 (figure 19). Repeat-until-success prepares and measures until the state is the one wanted: after n attempts with success probability p each, $1 - (1 - p)^n$ succeed (0.9889 for $p = 0.43$ and 8 attempts), less the decay after the last decision.

Exercises

1. *Blobs.* Record `X.readout` with the emulator's `noise` at 10, 20 and 40 mV; `A.readout` gives SNR and F; compare with $\Phi(\mathrm{SNR}/2)$ and the decay correction.
2. *Matched weights.* Average the two states' readout traces (window 0 with 16 points), take their difference as the window's `weights`, and repeat the readout: the SNR improves over the boxcar, most with a long `ringdown`.
3. *The reset's error budget.* Run `X.active_reset` with the emulator's $T_1 = 10$, 20 and 40 µs: the residue scales roughly as $1/T_1$ (about 0.033 at 20 µs).
4. *Repeat-until-success.* Run it with 1 … 8 attempts and compare with $1 - (1 - p)^n$.

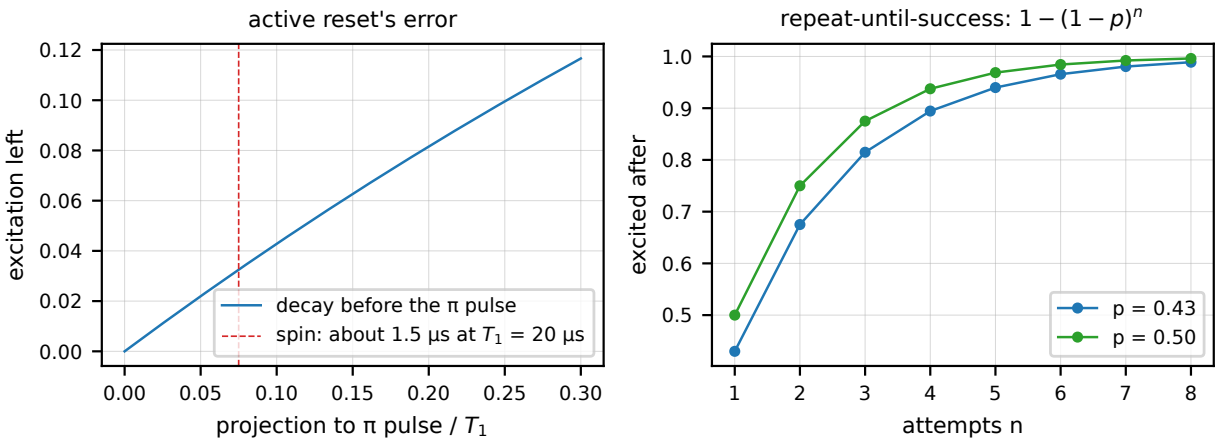


Figure 19: Left: an active reset's residual excitation against the projection-to-pulse time in units of T_1, with `spin`'s decision marked. Right: repeat-until-success against the attempts (computed).

11 The receiver: mixing, filtering, averaging

11.1 Quadrature demodulation with the transmitter's oscillator

The ADC samples the input $x(t)$ every $t_s = 8\,\text{ns}$ (14 bits, $\pm 1\,\text{V}$ on the LV range). The receiver multiplies it by $e^{-i\theta(t)}$, $\theta(t)$ the phase of the transmitter's oscillator at that sample (in `spin`: TX1's or TX2's, per window). A signal $A\cos(2\pi f_{\text{TX}} t + \varphi)$ becomes $\frac{A}{2}e^{i\varphi}$ plus a component at $-2f_{\text{TX}}$ that the filter removes; the driver's gain makes the result read $Ae^{i\varphi}$ in volts. Because the same oscillator drives the transmitter, its phase jumps, frequency switches, chirps and re-referencing to a trigger all appear in the receiver too: transmitter and receiver are coherent by construction, and a phase cycle's receiver phase is an exact 90° step applied after the mixer.

11.2 The order-2 CIC filter

Each output point Z_m is the triangle-weighted sum of $2R-1$ mixed samples (two cascaded boxcars of length R, then decimation by R) [14]:

$$Z_m = \frac{1}{R^2}\sum_{j=0}^{2R-2} h_j\, z(s + mR + 2R - 2 - j), \qquad h = 1, 2, \ldots, R, \ldots, 2, 1, \tag{11}$$

its time label the triangle's centre. In the bit-exact model (the RTL computes the same bits) the two integrators run at every sample and the two combs at every R-th:

spin/control/spin/model.py

```
        z = mix(xr, xi, theta, d["raw"])
        for j in (0, 1):
            st["I1"][j] = wrap(st["I1"][j] + z[j], 48)
            st["I2"][j] = wrap(st["I2"][j] + st["I1"][j], 48)
        if c == st["next"]:
            st["next"] += st["R"]
            c1 = [wrap(st["I2"][j] - st["D1"][j], 48) for j in (0, 1)]
            c2 = [wrap(c1[j] - st["C1"][j], 48) for j in (0, 1)]
            st["D1"], st["C1"] = list(st["I2"]), c1
            if st["out"]:
                self._point(rx, st, c, c2)
            st["out"] = True
```

Its frequency response is section 4's $H(f)$; its noise gain, for white noise of standard deviation σ per sample, is $\sigma\sqrt{\sum h_j^2}/\sum h_j$: 0.4146 for $R = 4$, 0.1037 for $R = 62$ (close to $\sqrt{2/3R} = 0.1037$)

and 0.0258 for $R = 1000$. A longer dwell averages more samples into each point, at the price of bandwidth.

11.3 Averaging

Coherent averaging adds the signal N times and the noise as $\sqrt{N}$: the signal-to-noise ratio grows as $\sqrt{N}$, which the user's manual measures on the board over 1024 scans. `spin` adds the points of every scan into 35-bit sums in the FPGA and the service adds frozen halves into 64-bit sums on the host, so nothing overflows over hours. The ADC's quantization noise, about $2\,\mathrm{V}/2^{14}/\sqrt{12} \approx 35\,\mu\mathrm{V}$ rms per sample, is far below a typical preamplifier's noise; a signal smaller than one LSB still averages out of the noise once the noise dithers it.

11.4 Weighted integration

A weighed window computes $S = \sum_k w_k Z_k$ at the full point rate. Normalized so that a steady tone of amplitude A reads A, S is a volts-scale estimate of the readout's amplitude and phase along the weights' shape; turning the weights by $-\varphi$ puts the decision axis on the real axis, so the state is the sign of $\mathrm{Re}\,S$ − threshold. The noise of S follows from the weights: $\sigma_S^2 = \sigma_Z^2 \sum |w_k|^2 / |\sum w_k|^2$ for uncorrelated points.

Exercises

1. *The noise gain.* With the emulator's signal at 0 and `noise` 10 mV, record one scan at $R = 4$, 62 and 1000 and compute each trace's rms per quadrature: 10 mV times 0.4146, 0.1037, 0.0258.
2. $\sqrt{N}$. Average 1, 4, 16, 64, 256 and 1024 scans of the same noise: the rms falls by half per factor four.
3. *Coherence.* Switch TX1 between two frequency registers inside a scan and back: a window at the first frequency after the switch continues the first phase, as if nothing had happened.

12 The emulator's model

12.1 One vector, one rotation per slot

The emulator integrates equation (2) for one Bloch vector in the frame of TX1's oscillator. Every 128 ns (a *slot* of 16 samples) it sums the slot's drive (TX1's amplitude after its envelope, at its phase: a rotation vector (w_x, w_y)) and offset (the spin's frequency minus TX1's, sample by sample, plus the shot's drawn offset and AO0's shift: w_z), and turns the vector by the exact rotation about $\boldsymbol{w}$ (Rodrigues' formula, from a table of $\sin\theta/\theta$ and $(1-\cos\theta)/\theta^2$), with the slot's relaxation applied after it:

$$\boldsymbol{M} \leftarrow E\,R(\boldsymbol{w})\,\boldsymbol{M} + (1 - \mathrm{e}^{-t_{\mathrm{slot}}/T_1})\,M_{\mathrm{eq}}\hat{z}, \qquad E = \mathrm{diag}(\mathrm{e}^{-t_{\mathrm{slot}}/T_2},\ \mathrm{e}^{-t_{\mathrm{slot}}/T_2},\ \mathrm{e}^{-t_{\mathrm{slot}}/T_1}). \tag{12}$$

This is exact when the drive and the offset are constant within a slot, and accurate to second order otherwise: within 0.1 % of the float Bloch equations for hard and soft pulses at offsets of ±100 kHz, with physical T_1 and T_2. The offset is clamped to ±0.4 turn per slot (±3.125 MHz), the drive limited to 1.95 MHz of Rabi frequency at full amplitude. The slot's end in the emulator's model (`emulator.py`; the RTL is bit-exact with it):

spin/control/spin/emulator.py

```
            if s.mode == EMU_QUBIT and st.reading:
                if st.out and not st.rd_first and ((r >> 20) & M24) < s.p1:
                    st.out = 0
                st.v[0] = [0, 0, -V_ONE if st.out else V_ONE]
            else:
                wx, wy = self._drive(st.dx, st.dy)
                wz = _clamp(wzf + st.delta + wao, WZ_MAX)
                st.v[0] = self.rotate(st.v[0], wx, wy, wz)
```

12.2 Limits that follow from the model

- $T_2 \leq 2T_1$: the vector's components are clamped to ± 1, exact inside the unit ball, where every physical vector stays (section 2); the driver refuses longer T_2.
- **The first pulse after a RUN**: TX1's oscillator stands still (step 0) until its first coherent start, so the spin precesses at its full frequency in the frame, clamped to 0.4 turn per slot. A pulse that starts inside such a slot turns about a tilted axis in that slot: a 90° pulse comes out at 67° to 89°, depending on where in its slot it starts (figure 20). Starting TX1 without output one slot before the first pulse (the templates' preludes) avoids it.
- **One isochromat per shot**: an ensemble's line is a draw of one offset per SHOT, so averages keep a residue near $\sqrt{\pi/(4N)}$ after the decay, and offset-sensitive sequences (90° trains, soft pulses) scatter from run to run as $1/\sqrt{N}$ of the per-isochromat spread.
- **Spin 1/2**: no quadrupolar or coupled spins, no diffusion, no field drift, no probe ring-down or receiver recovery.

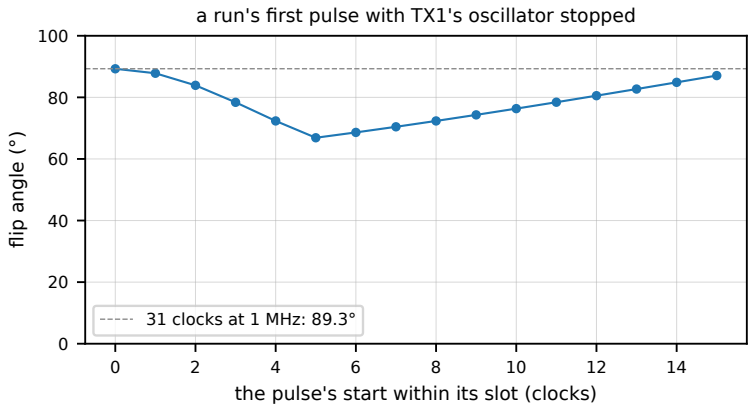


Figure 20: A run's first 90° pulse (31 clocks at 1 MHz) when TX1's oscillator has not yet run, against where in its 128 ns slot the pulse starts (computed: one rotation for the slot's clamped offset and drive, then the rest of the pulse).

12.3 The qubit

A TX2 readout pulse makes the emulator draw an outcome from $(1 - M_z)/2$ at the readout's start (from the state at the end of the last slot before it), collapse the vector to that state, and let it decay with T_1 during the readout (the outcome can flip from 1 to 0 at a random slot). The readout's signal is TX2's amplitude turned by $\pm\chi$, passed through a one-pole resonator of time constant `ringdown`, plus Gaussian noise: a model of dispersive readout without photon-number dependence or measurement-induced transitions.

12.4 How it is verified

The emulator's RTL equals its Python model (`emulator.py`) bit for bit: the output at every clock and the vector after every update, in three test scenarios that every one of 55 deliberate errors in the RTL fails. The model is compared with the float Bloch equations (`bloch.py`) in the host tests, and the board's emulator with both: seeded runs equal the model, free runs reproduce T_2^*, T_2, T_1, Rabi and Ramsey frequencies and echoes within their statistics.

Exercises

1. *The slot.* Compare a 31-clock and a 32-clock 90° pulse at offsets of 0, 0.5 and 1 MHz in `X.simulate` and in `spin.bloch`: the emulator follows the float equations within 0.1 %.

2. *The first pulse.* In a hand-written program, put the first pulse right after the RUN without starting TX1 first; compare its nutation with the template's: the flip angle is smaller.
3. *The residue.* Average FIDs of 16, 64, 256 and 1024 shots of a broad line and measure the signal after the decay: it falls as $1/\sqrt{N}$.

A Expected results of the exercises

The values below follow from the formulas and computations of this manual for the settings the exercises name; on the board each is met within its noise (the user's manual records the board's own runs of most of them).

Exercise	Expected
2: nutation at 0.8 V, `rabi` 1.25 MHz/V	$f_R = 1.00$ MHz; 1 MHz off resonance: $\sqrt{2}$ MHz, smaller amplitude
2: Rabi and Ramsey on the qubit	Rabi 1 MHz; Ramsey fringes at the offset (500 kHz)
3: excitation profile of a 0.25 µs 90° pulse	half height over 5.50 MHz; a 1 µs Gaussian: 2.25 MHz
3: adiabatic inversion	complete above $Q \approx 3.37$; $f_R = 1$ MHz at 10^{12} Hz/s: $Q = 6.3$
4: a 5 kHz Lorentzian	$T_2^* = 31.8$ µs; the Gaussian's full width 11.77 kHz
4: the residue after the decay	about $\sqrt{\pi/(4N)}$: 22 % at 16 shots, 5.5 % at 256
4: a 10 µs dead time, 20 kHz off resonance	a first-order phase of 72°
5: CP against CPMG with 170° pulses	after 24 echoes 0.34 against 1.00
6: inversion recovery, $T_1 = 1$ ms	the null at $0.693\,T_1$; the best t_r $1.25\,T_1$
7: an I/Q imbalance of 5 % and 5°	the image 26 dB down in one scan; CYCLOPS removes it
8: a spin-3/2 powder	the largest signal at 102° of effective flip angle
9: an IQ mixer at 1 % and 1°	image rejection 40 dB
10: blobs at SNR 8	$F = 0.99997$; with decay in a 0.5 µs readout at 20 µs: 0.9938
10: matched weights	SNR gain 1.06 (64 ns ring-down), 1.30 (256 ns)
10: active reset at $T_1 = 20$ µs	about 0.033 left; 8 attempts at $p = 0.43$: 0.9889
11: noise per point, 10 mV per sample	0.4146, 0.1037, 0.0258 of it for $R = 4$, 62, 1000
12: a run's first pulse with TX1 stopped	a 90° pulse turns 67° … 89°

References

[1] M. H. Levitt. *Spin Dynamics: Basics of Nuclear Magnetic Resonance.* Wiley, 2nd edition, 2008.

[2] J. Keeler. *Understanding NMR Spectroscopy.* Wiley, 2nd edition, 2010.

[3] E. Fukushima and S. B. W. Roeder. *Experimental Pulse NMR: A Nuts and Bolts Approach.* Addison-Wesley, 1981.

[4] R. R. Ernst, G. Bodenhausen, and A. Wokaun. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions.* Oxford University Press, 1987.

[5] T. P. Das and E. L. Hahn. *Nuclear Quadrupole Resonance Spectroscopy.* Solid State Physics, Supplement 1. Academic Press, 1958.

[6] A. Schweiger and G. Jeschke. *Principles of Pulse Electron Paramagnetic Resonance.* Oxford University Press, 2001.

[7] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. A quantum engineer's guide to superconducting qubits. *Appl. Phys. Rev.*, 6:021318, 2019.

[8] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff. Circuit quantum electrodynamics. *Rev. Mod. Phys.*, 93:025005, 2021.

[9] F. Bloch. Nuclear induction. *Phys. Rev.*, 70:460–474, 1946.

[10] I. I. Rabi. Space quantization in a gyrating magnetic field. *Phys. Rev.*, 51:652–654, 1937.

[11] N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Phys. Rev.*, 78:695–699, 1950.

[12] M. Garwood and L. DelaBarre. The return of the frequency sweep: Designing adiabatic pulses for contemporary NMR. *J. Magn. Reson.*, 153:155–177, 2001.

[13] R. R. Ernst and W. A. Anderson. Application of fourier transform spectroscopy to magnetic resonance. *Rev. Sci. Instrum.*, 37:93–102, 1966.

[14] E. B. Hogenauer. An economical class of digital filters for decimation and interpolation. *IEEE Trans. Acoust., Speech, Signal Process.*, 29:155–162, 1981.

[15] E. L. Hahn. Spin echoes. *Phys. Rev.*, 80:580–594, 1950.

[16] H. C. Torrey. Bloch equations with diffusion terms. *Phys. Rev.*, 104:563–565, 1956.

[17] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, 94:630–638, 1954.

[18] S. Meiboom and D. Gill. Modified spin-echo method for measuring nuclear relaxation times. *Rev. Sci. Instrum.*, 29:688–691, 1958.

[19] R. A. Marino and S. M. Klainer. Multiple spin echoes in pure quadrupole resonance. *J. Chem. Phys.*, 67:3388–3389, 1977.

[20] G. Bodenhausen, H. Kogler, and R. R. Ernst. Selection of coherence-transfer pathways in NMR pulse experiments. *J. Magn. Reson.*, 58:370–388, 1984.

[21] D. I. Hoult and R. E. Richards. Critical factors in the design of sensitive high resolution nuclear magnetic resonance spectrometers. *Proc. R. Soc. Lond. A*, 344:311–340, 1975.

[22] G. Bodenhausen, R. Freeman, and D. L. Turner. Suppression of artifacts in two-dimensional J spectroscopy. *J. Magn. Reson.*, 27:511–514, 1977.

[23] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proc. IEEE*, 51:89–109, 1963.

[24] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Approaching unit visibility for control of a superconducting qubit with dispersive readout. *Phys. Rev. Lett.*, 95:060501, 2005.

[25] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo. Initialization by measurement of a superconducting quantum bit circuit. *Phys. Rev. Lett.*, 109:050507, 2012.