boxcar: theory manual

Pump–probe detection with a boxcar averager, with exercises through a cable loopback

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

Red Pitaya `boxcar` design, version 1.0
November 10, 2026

Abstract

This manual explains the measurement that the `boxcar` instrument makes: what a pump–probe signal is and how small it is; gated integration of detector pulses, the best gate, and what sampling at 8 ns does to a gate's sum; normalization by a reference detector; modulation of the pump (chopping, shot-by-shot categories, pairs) against a laser's slow noise; the statistics of the sums and the errors of $\Delta T/T$; delay scans, their grids and drifts; the boxcar output's running average; when a boxcar beats a lock-in; and the test source, the emulated laser and sample that the exercises use. Each section ends with exercises that run through two cables (OUT1 → IN1, OUT2 → IN2), so a class needs no laser; their expected results are given at the end. Figures are computed (numpy); the board's recordings are in the user's manual (`user_manual.pdf`).

Contents

1 Introduction **3**

2 Pump–probe signals **3**
2.1 What is measured . . . 3
2.2 Time scales . . . 3
2.3 Why so many shots . . . 3

3 Gated integration and sampling **4**
3.1 The gate and its width . . . 4
3.2 Sampling: why the gate's edges matter . . . 5

4 Normalization by a reference detector **6**
4.1 The ratio . . . 6
4.2 The mean of ratios and the ratio of means . . . 6
4.3 What the ratio does not remove . . . 6

5 Modulation: chopping, categories and pairs **7**
5.1 Why modulate . . . 7
5.2 Chopping frequency . . . 7
5.3 Categories from a chopper . . . 7
5.4 Pairs . . . 8

6 Statistics: the sums and the errors **8**
6.1 What the sums give . . . 8
6.2 The error of $\Delta T/T$ . . . 8
6.3 How many shots . . . 8
6.4 The ratio of means and the covariance . . . 9

7 Delay scans **9**
7.1 Mechanical and electronic delays . . . 9
7.2 Grids . . . 9
7.3 Gates that follow the probe . . . 10

8 The boxcar output **10**

9 Boxcar and lock-in **11**

10 The test source: an emulated laser and sample **11**
10.1 Its shots and pulses . . . 11
10.2 What its sample really has . . . 12
10.3 Its limits as a laser . . . 12

11 Answers to the exercises **12**

1 Introduction

A pump–probe measurement asks how a sample changes after a short excitation, as a function of the time since it. The answer is usually a small number per delay, buried in the noise of the laser and the detectors; extracting it is a problem of signal processing as much as of optics. This manual treats that problem with the tools that `boxcar` provides: gates on each detector pulse, categories of shots, exact sums per delay point, and a reference detector. It does not replace an optics text [1, 2]; it explains the electronics' side, with numbers.

Throughout, a *shot* is one laser pulse (or one probe pulse), a *gate* is a time window of a detector's signal summed per shot, *category* 1 means the pump was on for that shot and 0 that it was off, and a *point* is one delay of a scan. D_1 is the signal detector's gate minus its baseline, D_2 the reference detector's, and $Q = D_1/D_2$.

The exercises. Each section ends with exercises for the board with two cables, OUT1 → IN1 and OUT2 → IN2, and `boxcar`'s test source as the laser and the sample. The code assumes a client:

```
import sys, time; sys.path[:0] = ["common/python", "boxcar/control"]
from boxcar import BoxcarClient, driver, model
dev = BoxcarClient("192.168.0.102")
```

The answers, with the values to expect, are in section 11.

2 Pump–probe signals

2.1 What is measured

A probe pulse of energy E crosses a sample of transmission T and reaches a detector: the signal is $S = k\,T\,E$, with k the detector's response. A pump pulse, a time τ earlier, changes the transmission by $\Delta T(\tau)$. Comparing shots with and without the pump,

$$\frac{\Delta T}{T}(\tau) = \frac{S_{\text{on}} - S_{\text{off}}}{S_{\text{off}}}, \tag{1}$$

a quantity independent of k and E if both stay the same between the two kinds of shots. For a thin sample with an absorption change $\Delta\alpha$ over a length L, $\Delta T/T \approx -\Delta\alpha L$; in reflection the same is written for $\Delta R/R$. Typical values range from 10^{-2} (strong excitation of a semiconductor) to 10^{-6} and below (weak excitation, thin films, single layers) [1].

2.2 Time scales

The time resolution is set by the pulses (their durations and their jitter) and by the delay: an optical delay line gives femtoseconds to nanoseconds (0.3 mm of path is 1 ps), an electronic delay between two lasers nanoseconds to milliseconds. `boxcar` times its outputs in 2 ns steps from the trigger's nanosecond: electronic delays are exact to that grid, and their jitter is the lasers'. The detector and the electronics need not be fast: a photodiode's pulse of a microsecond is fine, as long as it has decayed before the next shot; what matters is that each shot's pulse is integrated into one number.

2.3 Why so many shots

The laser's pulse energy fluctuates from shot to shot (about 1 % for an amplified laser), and drifts slowly (power, pointing, temperature). A signal of 10^{-4} is thus a hundred times smaller than one shot's noise; with N shots per category white noise falls as $\sqrt{2/N}$ (section 6), so 10^{-4} needs some 10^4 shots for a signal-to-noise ratio of 1 even after the reference has removed most of the laser's fluctuation, and 10^{-6} needs the reference, fast modulation and millions of shots. The rest of this manual is about these three levers: the reference (section 4), the modulation (section 5) and the number of shots, without losing any (sections 3 and 6).

Exercises

1. Load `loopback-demo` and read the programmed $\Delta T/T$ of the test source (`dev.config()["test"]["dtt"]`) and the value its sample really has (`driver.emulated_dtt(dev.config(), dev.calibration())`). Why do they differ? (Section 10.)
2. Run it for 2 s and read `dev.points()[0]["dtt"]`. How many shots per category were summed, and how large is the error of $\Delta T/T$ compared with $\Delta T/T$ itself?

3 Gated integration and sampling

3.1 The gate and its width

A detector turns a probe pulse into a voltage pulse $x(t) = Af(t)$, with f its shape and A its size, plus noise $n(t)$. A gate sums the samples x_k of a window of W samples: $g = \sum_k x_k$. With white noise of density e_n (V/$\sqrt{\text{Hz}}$) the gate's mean level $\bar{x} = g/W$ has a noise of $e_n/\sqrt{2T_\text{g}}$, $T_\text{g} = W \cdot 8\,\text{ns}$ the gate's duration, while the pulse contributes $A\int_\text{gate} f\,\text{d}t/T_\text{g}$. The signal-to-noise ratio

$$\text{SNR}(T_\text{g}) = \frac{A\sqrt{2}}{e_n}\,\frac{\int_0^{T_\text{g}} f(t)\,\text{d}t}{\sqrt{T_\text{g}}} \tag{2}$$

grows while the window still gains pulse area and falls once it only gains noise. The best any linear filter can do is the *matched filter*, which weights each sample by the pulse's shape: its SNR is $\frac{A\sqrt{2}}{e_n}\left(\int f^2\,\text{d}t\right)^{1/2}$. Figure 1 compares the two: for an exponential decay of time τ_d the best rectangular gate is $1.25\,\tau_d$ long and reaches 0.903 of the matched filter's SNR; a gate of $3\tau_d$ still gives 0.776. A rectangular gate placed on the pulse thus loses only about 10 %, which is why boxcar averagers use them (weighted gates are not in `boxcar`).

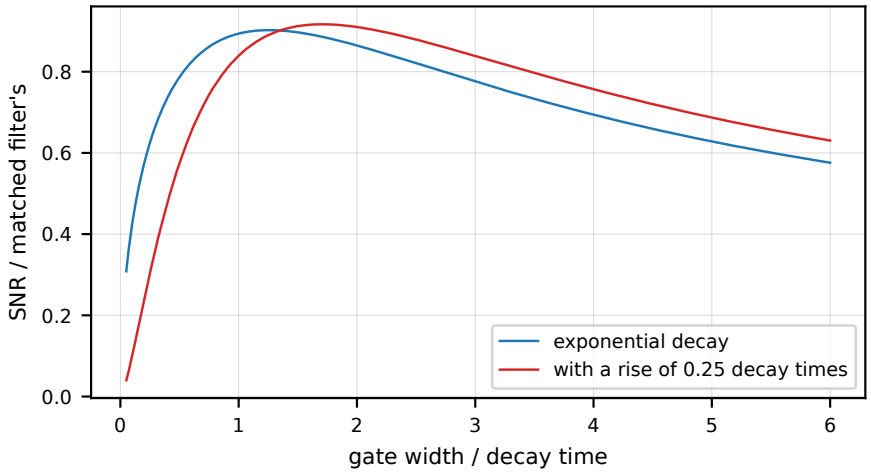


Figure 1: The SNR of a rectangular gate that starts at the pulse, against its width, relative to the matched filter, for an exponential pulse and for one with a rise of a quarter of its decay time.

The baseline gate. A detector's offset drifts, and the pump may leave a slow background. $D = G_a - (W_a/W_b)G_b$ subtracts the mean level of a baseline window G_b before the pulse: any level constant over the two windows cancels, and so does noise slower than their distance. The price is the baseline's own noise, added with the weight $(W_a/W_b)^2$: a baseline window at least as wide as the signal's keeps it small.

3.2 Sampling: why the gate's edges matter

The inputs are sampled every $T = 8\,\text{ns}$, and an asynchronous laser's pulse falls anywhere within a sample. A gate sums samples, not the continuous pulse: does the sum depend on where the pulse fell? The answer is the Poisson summation formula: for a pulse f with Fourier transform F,

$$\sum_k f(kT - \delta) = \frac{1}{T} \sum_m F\left(\frac{m}{T}\right) \mathrm{e}^{-2\pi \mathrm{i} m \delta / T}. \tag{3}$$

If the pulse has no spectrum at multiples of $1/T = 125\,\text{MHz}$, only the $m = 0$ term remains: the sum of all its samples is its area over T, at every position δ. A window holding the whole pulse therefore sums the same for every shot; a window that cuts the pulse does not, because cutting is a sharp edge in time. Figure 2 shows it for Gaussian pulses: a window holding a pulse of 4 ns rms varies by $\pm 1.4\,\%$ with the position, one of 2 ns rms (whose spectrum reaches 125 MHz) by $\pm 57.2\,\%$, and a tight window of two samples by several percent for any pulse ($\pm 3.3\,\%$ at 4 ns).

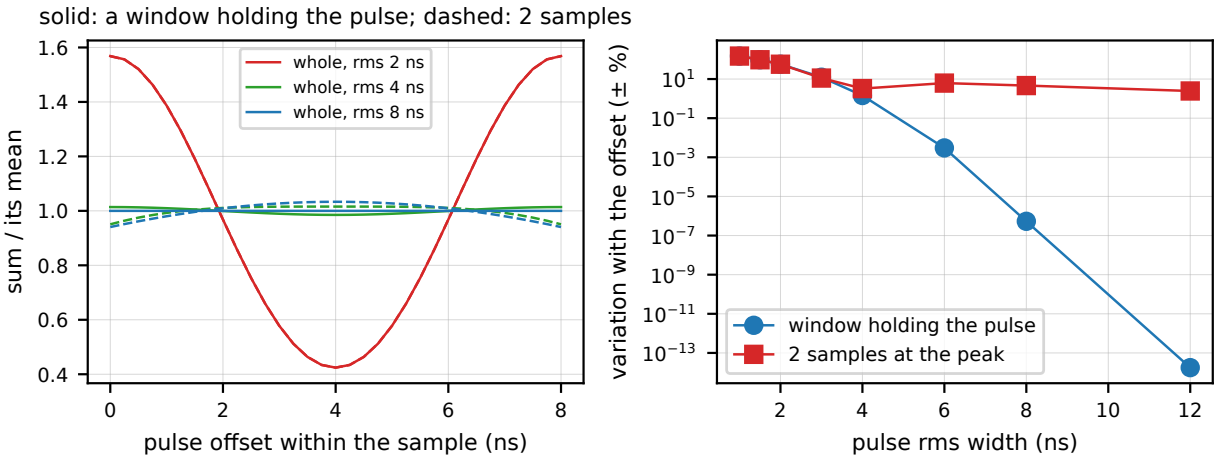


Figure 2: Sums of 8 ns samples of Gaussian pulses against the pulse's position within a sample. Left: three widths, a window holding the pulse (solid) and two samples at the peak (dashed). Right: the variation against the pulse's width.

The board's inputs pass about 60 MHz, so a real detector's pulse reaches the converters smoothed to a few nanoseconds rms at least: a gate that holds the whole pulse with margin, edges on the flat baseline, is insensitive to the laser's phase. The user's manual (figure "phase") measures the same on the board with the test source; there a very short pulse varies even in the widest gate, because the test source itself is sampled at 8 ns (section 10). With a second board running `awg` as a truly asynchronous laser (its band-limited pulses walking through `boxcar`'s samples), a gate holding the whole pulse varied by $\pm 0.05\,\%$ and a tight one by $\pm 35\,\%$ (the user's manual, figure "awg_session"). `boxcar` records every shot's trigger time to 1 ns, so a residual phase dependence can be corrected offline.

Exercises

1. With `loopback-demo`, make G1 two samples wide around the pulse's peak (read it from the averaged traces) and compare the spread of D_1 per shot with the 25-sample gate's. Which gate gives the smaller relative spread, and why? (The test source's pulses are synchronous with the board's clock here: the phase does not change.)
2. Repeat the user's manual's phase measurement (`tools/record_manual_data.py -only phase`) and fit the tight gate's variation against the pulse's position.

4 Normalization by a reference detector

4.1 The ratio

A beam splitter before the sample sends part of the probe to a reference detector. If the laser's pulse energy is $E(1+\epsilon)$ with a fluctuation ϵ (about 1 % per shot for an amplified laser), the signal and the reference are

$$D_1 = k_1 T E(1+\epsilon_1) + n_1, \qquad D_2 = k_2 E(1+\epsilon_2) + n_2, \tag{4}$$

where $\epsilon_{1,2}$ are the fluctuations as each detector sees them and $n_{1,2}$ the detectors' own noise. If both see the same fluctuation ($\epsilon_1 = \epsilon_2$), the ratio $Q = D_1/D_2 = (k_1/k_2)\,T$ is free of it. In general, to first order, the relative variance of Q is

$$\sigma_Q^2 = \sigma_{\epsilon_1}^2 + \sigma_{\epsilon_2}^2 - 2\rho\,\sigma_{\epsilon_1}\sigma_{\epsilon_2} + \sigma_{n_1}^2 + \sigma_{n_2}^2, \tag{5}$$

with ρ the correlation of the fluctuation at the two detectors. Figure 3 shows how strongly the result depends on ρ: for a 1 % laser and detectors with 10^{-4} of their own noise, $\rho = 0.999$ leaves 0.047 % and $\rho = 0.99$ already 0.142 %. The reference removes the laser's noise only as far as both detectors see the same laser: the same part of the beam (pointing, mode), the same spectrum, the same linear range [3, 4].

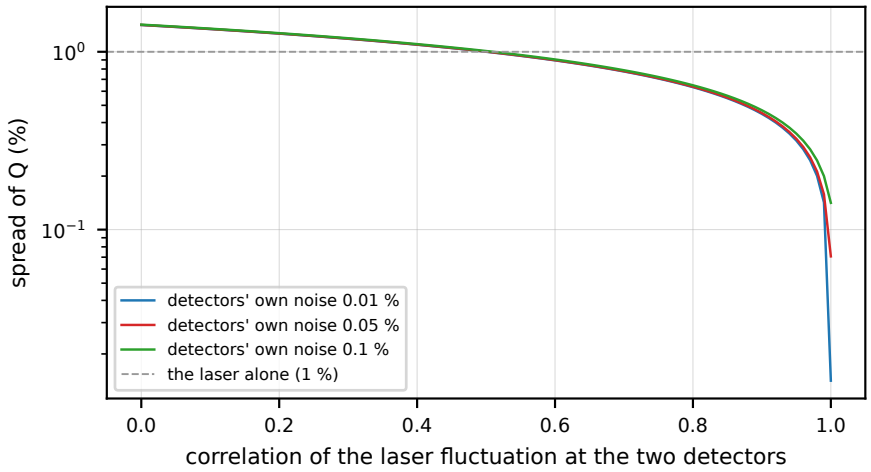


Figure 3: The relative spread of $Q = D_1/D_2$ against the correlation of the laser's fluctuation at the two detectors, for a 1 % laser and three levels of the detectors' own noise.

4.2 The mean of ratios and the ratio of means

`boxcar` forms Q for every shot and sums it: $\langle Q \rangle$ is a mean of ratios. With a very noisy reference this is biased ($\langle 1/D_2 \rangle \neq 1/\langle D_2 \rangle$, by about the reference's relative variance) and noisier than necessary; the ratio of the means $\langle D_1 \rangle / \langle D_2 \rangle$ is then better, and its error needs the covariance of D_1 and D_2, which is why the board also sums $D_1 D_2$ per point. For a reference with a relative noise below about 1 % both agree to 10^{-4}, and the bias cancels in $\Delta T/T$ (it is the same factor for both categories).

4.3 What the ratio does not remove

The detectors' own noise adds in quadrature (both of them); a nonlinearity that differs between the detectors turns the laser's fluctuation into a residual; fluctuations that reach one detector only (pointing changes on a small detector, an aperture, scattered pump light) stay. Matching the detectors (the same type, the same signal level, well within their linear range) matters more than any electronics.

Exercises

1. With `loopback-demo` (20 % shared noise), record SHOT records and compute the relative spread of D_1, D_2 and Q for category 0. Then set `test.noise` to 0 and repeat: what is left in Q, and where does it come from?
2. Halve the reference's amplitude (`test.out2.amplitude`) and repeat with the noise on: how does the spread of Q change? Predict it from the equation above, with the converters' noise from the noise-free run.

5 Modulation: chopping, categories and pairs

5.1 Why modulate

$\Delta T/T$ compares two measurements, with the pump and without. Anything that changes between them that is not the pump enters the result: the laser's power and pointing, the detectors' gain, the sample's temperature. A laser's noise is not white: below some frequency it grows, roughly as $1/f$ (flicker, air currents, thermal drifts) [3]. Alternating the two measurements fast moves the comparison to a frequency where the noise is lower; alternating them shot by shot moves it to half the repetition rate, as high as it can go.

5.2 Chopping frequency

Figure 4 simulates a laser with 1 % of noise per shot, half of it white and half $1/f$, and a pump that is on for m shots and off for m. With the categories alternating shot by shot ($m = 1$) the error of $\Delta T/T$ from 4096 shots per category is 0.74 times the white noise's alone: the $1/f$ half has cancelled. With $m = 100$ (a chopper at a hundredth of the laser's rate) the error is 1.9 times larger: the slow noise is back. A chopper at half the laser's rate, synchronized to it, is the ideal for an amplified laser; at high repetition rates a modulator driven by the board does the same (an output in mode `next`).

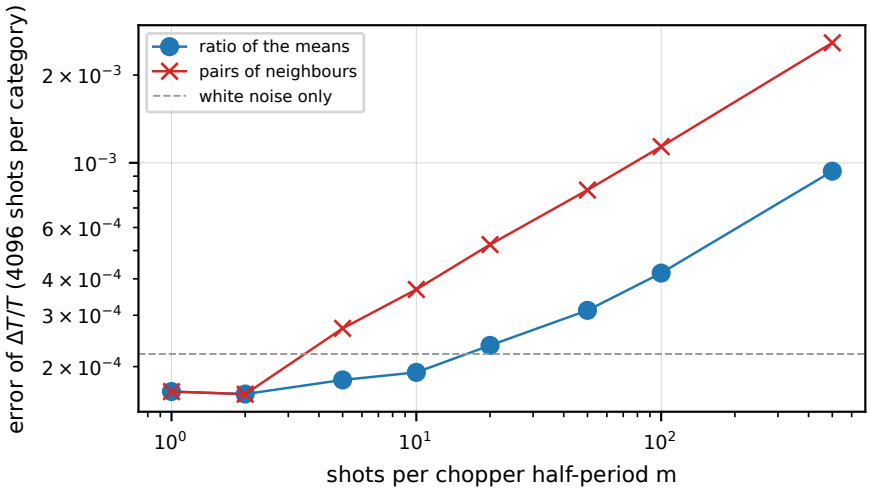


Figure 4: Simulated: the error of $\Delta T/T$ from 4096 shots per category against the number of shots per chopper half-period, for a laser with half white and half $1/f$ noise; the ratio of the means and the pairs of neighbouring shots.

5.3 Categories from a chopper

A mechanical chopper's blade takes time to cross the beam: shots near an edge see a partly blocked pump. `boxcar` reads CHOP at each shot and marks shots within a guard time of an edge as *invalid*; they are counted but kept out of both categories. A chopper locked to the laser puts its edges between

shots, and the guard then rejects nothing; a free-running chopper loses the fraction $2t_{\text{guard}}f_{\text{chop}}$ of the shots.

Chopper edges. A chopper reference or a modulator's drive is a fast logic edge, and its wire runs next to the detectors' cables: an edge shortly before the gates on one category's shots only couples into the inputs and looks like a signal. With `awg` as the laser, an edge about 60 ns before the gate on pump-on shots added 4.3×10^{-4} to $\Delta T/T$ (0.2×10^{-4} with the same timing but no edge on the wire). Put the edges half a period away from the shots, as the presets do.

5.4 Pairs

Besides the means per category, `boxcar` sums the differences of neighbouring shots of different categories, $d = Q_{\text{on}} - Q_{\text{off}}$. Their mean over $\langle Q \rangle_0$ is another estimate of $\Delta T/T$. Its variance is $2\sigma^2(1 - \rho_1)$ per pair, with ρ_1 the correlation of neighbouring shots: slow noise, which is correlated from shot to shot, cancels within each pair. With shot-by-shot categories the two estimates are equivalent; with slow chopping only the shots at the edges form pairs, and the pairs' estimate gets worse (5.1 times the white noise at $m = 100$ in figure 4), while the ratio of the means uses every shot.

Exercises

1. With `loopback-demo`, compare $\Delta T/T$ and its error for the patterns `"01"` and `"0011"` (set both `categories.pattern` and `test.pattern`). The test source's noise is white from shot to shot: what do you expect, and what do the "pairs" give?
2. A chopper at 50 Hz on a 1 kHz laser with a guard of 1 µs: what fraction of the shots is invalid if the chopper runs free? If it is locked to the laser?

6 Statistics: the sums and the errors

6.1 What the sums give

For each point and category the board keeps the number of shots n, $\sum x$ and $\sum x^2$ for every value x (the gates, D_1, D_2, Q), exactly (48 and 80 bits: no rounding and no overflow in any realistic run). From them

$$\langle x \rangle = \frac{\sum x}{n}, \qquad s^2 = \frac{\sum x^2 - n\langle x \rangle^2}{n-1}, \qquad \text{sem} = \frac{s}{\sqrt{n}}. \tag{6}$$

Sums over sweeps are sums of sums: the service adds the POINT records, and the result is the same as one long point. A point longer than 65,536 shots is reported in parts for the same reason.

6.2 The error of $\Delta T/T$

For $r = \langle Q \rangle_1 / \langle Q \rangle_0$, to first order,

$$\left(\frac{\delta r}{r}\right)^2 = \left(\frac{\text{sem}_1}{\langle Q \rangle_1}\right)^2 + \left(\frac{\text{sem}_0}{\langle Q \rangle_0}\right)^2, \tag{7}$$

and $\delta(\Delta T/T) = \delta r$. With a per-shot relative spread σ of Q in both categories and N shots in each, $\delta(\Delta T/T) \approx \sigma\sqrt{2/N}$. The pairs' estimate has $\delta = s_d/(\sqrt{n_p}\,\langle Q \rangle_0)$, s_d the pairs' spread and n_p their number. The user's manual (figure "averaging") measures the $1/\sqrt{N}$ law on the board up to 10,000 shots per category.

6.3 How many shots

Required: $N \approx 2\,(\sigma/\delta)^2$ shots per category for an error δ. With $\sigma = 10^{-3}$ (a good reference) and $\delta = 10^{-6}$, $N = 2 \times 10^6$ per category: 4000 s at 1 kHz, 2 s at 2 MHz. The ultimate floor is the shot noise of the detected photons: a detector that collects N_γ photons per pulse cannot do better than $\sigma = 1/\sqrt{N_\gamma}$ per shot (10^{-4} for 10^8 photons).

6.4 The ratio of means and the covariance

For the ratio of the means $R = \langle D_1 \rangle / \langle D_2 \rangle$ (section 4),

$$\left(\frac{\delta R}{R}\right)^2 = \frac{1}{n}\left(\frac{s_1^2}{\langle D_1\rangle^2} + \frac{s_2^2}{\langle D_2\rangle^2} - \frac{2\,c_{12}}{\langle D_1\rangle\langle D_2\rangle}\right), \quad c_{12} = \frac{\sum D_1 D_2 - n\langle D_1\rangle\langle D_2\rangle}{n-1}, \tag{8}$$

which is why the board keeps $\sum D_1 D_2$: the correlated part of the two detectors' fluctuation cancels in R as it does in Q.

Exercises

1. From `dev.points()` of a `loopback-demo` run, compute $\Delta T/T$ and its error by hand from the means and standard errors of Q, and compare with `["dtt"]`.
2. How long must `loopback-demo` run for an error of 10^{-5} in $\Delta T/T$? Run it and check.

7 Delay scans

7.1 Mechanical and electronic delays

An optical delay line changes the probe's path: 1 ps per 0.15 mm of travel of a retroreflector, from femtoseconds to a few nanoseconds; the PC moves it and the board measures N shots per position (the hold). An electronic delay fires two lasers at different times: nanoseconds to milliseconds, exact to the board's 2 ns steps, limited by the lasers' jitter; the board steps it itself from a table. Both give a curve $\Delta T/T(\tau_p)$ on a grid of points p.

7.2 Grids

A linear grid suits dynamics on one time scale; a decay over decades needs a logarithmic grid, with the same number of points per decade; a few points before time zero fix the baseline. The order of the points matters when anything drifts: an ordered scan maps a drift onto τ, a shuffled one turns it into scatter, and repeated sweeps average it over the curve. Figure 5 simulates a decay with a response that drops by 10 % over one scan: a fit of the ordered scan gives a decay time of 1.949 instead of 2.000, the shuffled scan 2.000 (a larger scatter instead), and ten sweeps added 1.995.

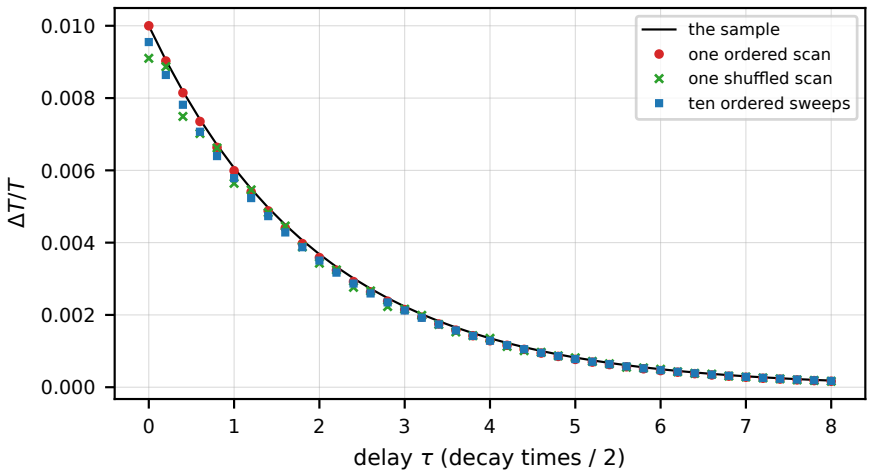


Figure 5: Simulated: a decay measured while the response drifts by −10 % per scan: one ordered scan, one shuffled scan and ten ordered sweeps added.

7.3 Gates that follow the probe

With an electronic delay the probe's detector pulse moves with the delay; the gates follow it with their own table (whole samples), so each point integrates the same part of the pulse. With an optical delay line nothing moves on the electronics' time scale: the gates stay.

Exercises

1. Run `loopback-delay-scan` for 20 sweeps and fit $\Delta T/T = a\,\mathrm{e}^{-\tau/\tau_d}$ to the points after the pump. Compare τ_d with the programmed 2 µs and a with the emulated sample's first value (`driver.emulated_dtt`).
2. Shuffle the delay table (and the gates' table the same way, and the test source's $\Delta T/T$ table with them) and repeat: does the result change?

8 The boxcar output

A boxcar output keeps, per category, the running average

$$A_n = A_{n-1} + \frac{s_n - A_{n-1}}{2^\tau}, \tag{9}$$

updated at each shot of the category (s_n the shot's value), seeded by the first. This is a first-order low-pass filter in units of shots: a step of the input reaches $1-(1-2^{-\tau})^n$ after n shots (99 % after 71 shots for $2^\tau = 16$), and white per-shot noise is reduced to

$$\frac{\sigma_A}{\sigma_s} = \frac{1}{\sqrt{2^{\tau+1}-1}}, \tag{10}$$

0.180 for 16 shots (figure 6). Its time constant is $2^\tau/f_c$, with f_c the shot rate of the category (half the laser's rate with alternating categories). The difference output (on − off) subtracts two such averages: its noise is $\sqrt{2}$ times one's.

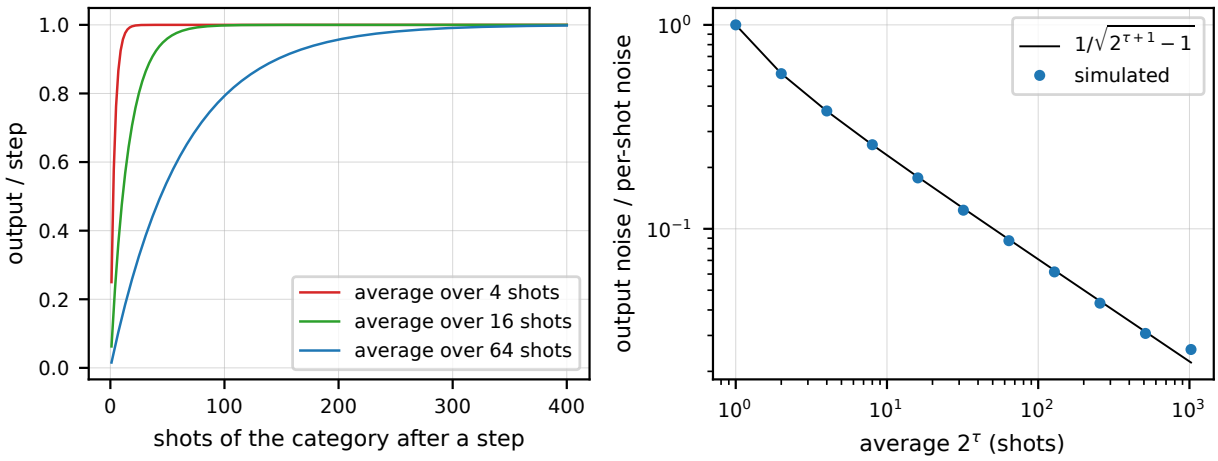


Figure 6: The running average $A \leftarrow A + (s-A)/2^\tau$. Left: the step response for 4, 16 and 64 shots. Right: the output's noise against the number of shots, simulated and computed.

On a chart recorder this is the classic boxcar trace. In a servo (holding a laser on a resonance by a pump–probe signal, for example) the average's time constant adds to the loop's delay: keep it short against the loop's response time, and protect the actuator's input with the output's limits, locked.

Exercises

1. Record the user's manual's step response (`record_manual_data.py -only output_step`) for an average of 4 and of 64 shots, and compare with the step response above.

2. With a constant input, compare the output's noise for 1 and 16 shots (read it back through OUT2 → IN2 with a wide gate before each shot) with the equation above.

9 Boxcar and lock-in

A lock-in amplifier at the laser's repetition rate is the other classic way to detect a pulsed signal [5, 6]. It multiplies the detector's signal by a sine at the rate and filters: it measures the signal's fundamental Fourier component and rejects everything else, but it also integrates the noise over the whole period. A boxcar integrates only during the gate. For rectangular pulses of duty cycle D (the pulse's width over the period) in white noise, with the same total time, the ratio of the two signal-to-noise ratios is

$$\frac{\mathrm{SNR}_{\mathrm{boxcar}}}{\mathrm{SNR}_{\text{lock-in}}} = \frac{\pi\sqrt{D}}{\sqrt{2}\,\sin(\pi D)} \approx \frac{1}{\sqrt{2D}} \quad (D \ll 1), \tag{11}$$

figure 7: 35.4 for a 400 ns pulse at 1 kHz ($D = 4 \times 10^{-4}$), 7.07 for $D = 0.01$, 2.27 for $D = 0.1$. For the short pulses of pulsed lasers the boxcar wins clearly, and it delivers each shot's value, which a lock-in cannot (no per-shot normalization, no categories, no histograms). A lock-in is the better choice for nearly sinusoidal signals, high duty cycles, or a modulation far below the laser's rate when the electronics cannot see the individual pulses; the repository's `lockin` design is one.

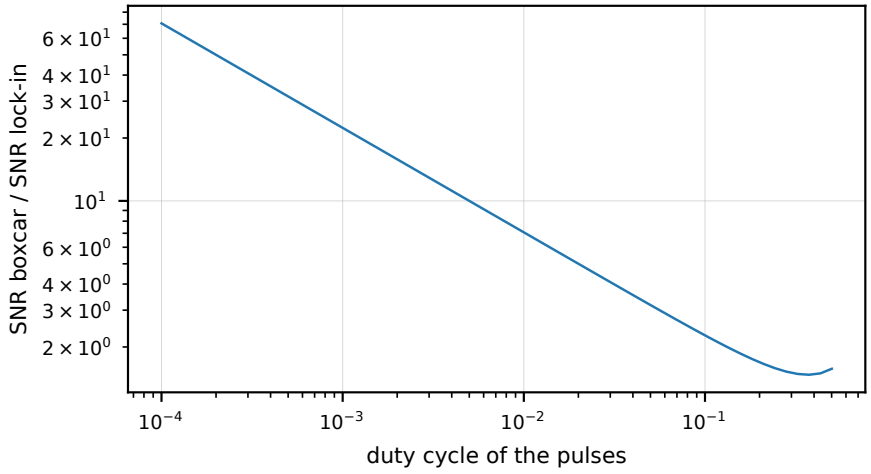


Figure 7: The SNR of a boxcar (a gate on the pulse) over that of a lock-in at the repetition rate, against the pulses' duty cycle, in white noise.

Exercises

1. What is the advantage for the `loopback-demo` pulses (about 200 ns long, every 20 µs)? And for `high-rep-modulator`'s 48 ns gates at 2 MHz?

10 The test source: an emulated laser and sample

The exercises use `boxcar`'s test source, which plays on OUT1 and OUT2 what detectors would see. Knowing its model exactly is what makes the exercises quantitative.

10.1 Its shots and pulses

Its shots come from its own timer (an asynchronous laser) or from a pulse output's pulses (a laser fired by the board). For each shot it draws a noise word n (uniform, 16 bits, signed) and forms the

shared factor $g = 2^{16} + \lfloor \sigma n / 2^{15} \rfloor$ (σ the noise setting times 2^{16}: a uniform distribution of $\pm\sigma/2^{16}$, a standard deviation of $\sigma/(2^{16}\sqrt{3})$, 11.5 % for the 20 % setting). An output in `reference` mode gets the amplitude $a = \lfloor Ag/2^{16} \rfloor$, one in `signal` mode also $\lfloor a\,(2^{16} + \Delta)/2^{16} \rfloor$ on the shots of category 1 of its own pattern, with $\Delta/2^{16}$ the $\Delta T/T$ of the current point. Each output plays a times a shape (1 ns resolution) sampled every 8 ns, after its delay; its sync and chopper levels can go to pins.

10.2 What its sample really has

A is an integer code (2167 for 0.3 V on the test board), and the product is truncated: the sample's $\Delta T/T$ is $\lfloor A(2^{16} + \Delta)/2^{16} \rfloor / A - 1$, smaller than $\Delta/2^{16}$ by up to $1/A$: 3 % of a programmed 0.01 at 0.3 V. With noise the truncation varies from shot to shot and is averaged. `model.test_dtt` computes the exact value (averaged over all noise words, as Q would see it) and `driver.emulated_dtt` the values of a configuration; these, not the table, are what a measurement through the cables should give.

10.3 Its limits as a laser

The test source's noise is white from shot to shot (no $1/f$, so the benefit of fast modulation cannot be seen with it), common to both outputs exactly (a perfect $\rho = 1$, section 4), and its pulses are sampled at 8 ns: a pulse shorter than about two samples changes its own area with its position within a sample, which a real detector's pulse, smoothed by the inputs' 60 MHz, does not. Fired by a pulse output, the shot that ends a point plays the next point's $\Delta T/T$ (it reads the point after the advance). Within these limits it is bit-exact: the model in `control/boxcar/model.py` predicts every output word.

Exercises

1. Compute the emulated $\Delta T/T$ for a programmed 0.01 at amplitudes of 0.1, 0.3 and 1.0 V (`driver.emulated_dtt`) and explain the trend.
2. Measure $\Delta T/T$ through the cables for a programmed 0.01 at 0.3 V and compare with both the table and the emulated value.

11 Answers to the exercises

Values measured on the test board (build 0xe4e10263) are quoted from the user's manual; computed ones from this manual's tools.

Section 2. (1) The table says 0.1, the emulated sample has 0.09979: the test source's amplitudes are integer codes and its product is truncated (section 10). (2) 2 s at 50 kHz are 50,000 shots per category; with Q's per-shot spread of 0.165 % the error is about 1.0e−05, a part in 10^4 of $\Delta T/T$.

Section 3. (1) The tight gate's relative spread is about 2.14 times the wide gate's (white noise, the test pulse's shape: the tight gate collects much less of the pulse but only somewhat less noise); the wide gate's was 0.255 % for D_1 without the shared noise. With the shared noise on, both show its 11.5 %: the laser's noise is in the pulse, not added to it. (2) The tight gate varies by about ±16.9 % for the smooth pulse (the user's manual).

Section 4. (1) D_1 and D_2 spread by 11.5 % and 11.5 %, Q by 0.165 %; without the shared noise Q spreads by 0.162 %, what the converters and the source add (less than D_1's 0.255 %: part of it is a gain noise common to both channels, which the ratio removes too). (2) The reference's own noise doubles relative to it: if the two channels contributed equally, Q's spread grows from 0.162 % to about $\sqrt{1+4}/\sqrt{2} \approx 1.6$ times that.

Section 5. (1) The same $\Delta T/T$ and the same errors for both patterns, by both estimates: the test source's noise is white, there is nothing slow to cancel. (2) Free-running: $2 \times 1\,\mu\text{s} \times 50\,\text{Hz} = 10^{-4}$ of the shots; locked: none.

Section 6. (1) $r = \langle Q \rangle_1 / \langle Q \rangle_0$ and δr from the formula agree with `["dtt"]` exactly (it is computed the same way). (2) $N = 2(\sigma/\delta)^2 \approx 54000$ shots per category for $\sigma = 0.165\,\%$, about 2.2 s at 50 kHz.

Section 7. (1) The board's recording (the user's manual) gives $\tau_d = 1.972\,\mu\text{s}$ and $a = 0.0496$ (the emulated first value 0.04978). The programmed 2 µs is shortened by the test source's truncation, which takes up to $1/A \approx 5 \times 10^{-4}$ off every value and so steepens the small tail: the emulated values themselves fit to 1.942 µs over the same points. (2) The same within the errors: the loopback does not drift on these time scales.

Section 8. (1) The read-back follows $1 - (1 - 2^{-\tau})^n$ with a scatter of about a millivolt (the user's manual: 0.8 mV for 16 shots). (2) The ratio of the noises is about $1/\sqrt{31} = 0.180$ for 16 shots.

Section 9. `loopback-demo`: $D \approx 0.01$, an advantage of about 7.07; `high-rep-modulator`: $D \approx 0.1$, about 2.27.

Section 10. (1) 0.00970, 0.00969 and 0.00997 for amplitude codes of 722, 2167 and 7223: the truncation costs up to $1/(A\Delta)$ of $\Delta T/T$, less with a larger amplitude, but not monotonically (it depends on the fractional part of $A\Delta/2^{16}$). (2) The board measured 0.009718 with 20 % noise for an emulated 0.009755 at 0.3 V (at 2 MHz, the user's manual): −0.38 % from the emulated value, 3 % from the table.

References

[1] Jagdeep Shah. *Ultrafast Spectroscopy of Semiconductors and Semiconductor Nanostructures.* Springer, Berlin, 2nd edition, 1999.

[2] Claude Rullière, editor. *Femtosecond Laser Pulses: Principles and Experiments.* Springer, New York, 2nd edition, 2005.

[3] Christopher A. Werley, Stephanie M. Teo, and Keith A. Nelson. Pulsed laser noise analysis and pump-probe signal detection with a data acquisition card. *Review of Scientific Instruments*, 82: 123108, 2011.

[4] Yuan Feng, Ilia Vinogradov, and Nien-Hui Ge. General noise suppression scheme with reference detection in heterodyne nonlinear spectroscopy. *Optics Express*, 25:26262–26279, 2017.

[5] M. L. Meade. *Lock-in Amplifiers: Principles and Applications.* Peter Peregrinus, London, 1983.

[6] Paul Horowitz and Winfield Hill. *The Art of Electronics.* Cambridge University Press, Cambridge, 3rd edition, 2015.