ajku: user's manual

Continuous-wave I/Q and noise measurements on the Red Pitaya

With a continuous-variable quantum-optics tutorial and laboratory characterization procedures

Rok Žitko
Jozef Stefan Institute, Ljubljana, Slovenia

Release 1.0: FPGA/driver 1.0, build `0xfb4e4824`; host processing revision 5
Board-qualified 2026-11-03; November 3, 2026

Abstract

`ajku` records two simultaneous voltage channels, computes continuous raw statistics and histograms, and analyses finite I/Q records on a PC. It is intended for balanced optical receivers and other continuous-wave noise experiments. The software retains native ADC codes, calibration and measurement context, distinguishes raw broadband quantities from explicitly filtered quantities, and records the gaps between retained waveforms.

This manual has three parts. Part I explains the instrument and its six implemented first-priority workflows. Part II connects the harmonic-oscillator description of light—coherent states, displacement, squeezing, non-classical states and phase-space representations—to photodetection and the quantities displayed by the analyser. Part III gives procedures for characterizing *the user's own board and optical chain*, followed by the measured digital operating envelope and an interface reference. No optical sensitivity, shot-noise clearance or state-reconstruction performance is inferred from the digital qualification runs. The theory figures are explicitly computed illustrations, not optical measurements.

Which calibration belongs to you?

The procedures transfer between instruments; numerical calibration constants do not. Board calibration is shared by designs on *one physical board*. Receiver gain, load impedance, optical loss, LO power, quadrature angle and cable delay belong to the particular experiment. Use your own measurements and preserve their provenance.

Contents

I Using the analyser **4**

1 What ajku measures **4**
1.1 Three different records of the experiment . . . 4
1.2 Signal path and responsibilities . . . 5
1.3 Windows, timing, slow references, higher orders and spectra (0.3–0.5) . . . 5

2 First measurement **6**
2.1 Equipment and installation . . . 6
2.2 Command-line check . . . 6
2.3 Notebook workflow . . . 7
2.4 Stopping and ownership . . . 7

3 Reading spectra, moments and distributions **8**
3.1 From native codes to volts . . . 8
3.2 Welch spectra and the I/Q matrix . . . 8
3.3 Coherence and phase uncertainty . . . 9
3.4 Explicit FIR filtering and decimation . . . 9
3.5 Time moments, cumulants and circularity . . . 10
3.6 Histograms, tails and uncertainty . . . 10
3.7 Amplitude, squared-amplitude and complex correlations . . . 11

4 References, recordings and reproducibility **12**
4.1 Reference comparisons . . . 12
4.2 Run contents and coverage . . . 12
4.3 Reanalysis and full-waveform export . . . 13

5 LO-power-dependent spectral measurements **13**
5.1 Model, errors and interpretation . . . 14

6 Drift monitoring and event-selected captures **14**

II Continuous-variable quantum optics and its measurements **16**

7 Optical modes, quadratures and quantum states **16**
7.1 One mode is one harmonic oscillator . . . 16
7.2 Pure states, mixtures and moments . . . 16

8 Coherent states and displacement **17**
8.1 The displacement operator . . . 17

9 Gaussian states, squeezing and loss **18**
9.1 Thermal light . . . 18
9.2 Squeezing as a quadrature transformation . . . 18
9.3 Efficiency and phase noise . . . 19

10 Non-classical states and phase-space representations **19**
10.1 Several meanings of “non-classical” . . . 19
10.2 Wigner representation . . . 20
10.3 Husimi Q and the ordering of moments . . . 21
10.4 Two-mode correlations and entanglement . . . 22

11 From optical operators to electrical I/Q data **22**
11.1 Balanced homodyne detection . . . 22
11.2 Dual homodyne and heterodyne: the added vacuum . . . 23
11.3 From volts to shot-noise units . . . 24
11.4 Why a joint ADC histogram is not a Wigner reconstruction . . . 25
11.5 Continuous waves, bandwidth and temporal modes . . . 25
11.6 Higher moments and photon statistics . . . 26

12 Worked connections and student exercises **26**
12.1 A coherent state in two measurement schemes . . . 26
12.2 Squeezing, loss and phase . . . 27
12.3 Classical mixtures versus cats . . . 27
12.4 A negative fourth cumulant from a classical state . . . 27
12.5 ADC resolution and a quiet signal . . . 27
12.6 Bandwidth and the apparent noise reduction . . . 27
12.7 Coverage and events . . . 27
12.8 No observed tail events . . . 28
12.9 Design a defensible observation . . . 28

III Characterization, qualification and reference **29**

13 Characterizing the user's board and optical chain **29**
13.1 Separate three calibration layers . . . 29
13.2 Board offset and DC gain . . . 29
13.3 Frequency response and channel delay . . . 30
13.4 Anti-alias rejection . . . 31
13.5 Electronic floor and LO-power sweep . . . 31
13.6 Optical quadrature angle, gain and coherent probe . . . 32
13.7 A squeezing measurement with a matched reference . . . 32
13.8 Uncertainty, selection and a complete lab record . . . 33

14 Measured digital operating envelope **33**
14.1 Qualification results . . . 33
14.2 Capture sizes, retention and throughput . . . 34
14.3 Real notebook checks and a GUI correction . . . 35
14.4 Supported use and remaining boundaries . . . 35

A Interface and file-format reference **36**
A.1 Python entry points . . . 36
A.2 Wire commands . . . 37
A.3 Raw and processed formats . . . 37

B Troubleshooting and interpretation checklist **38**

C Glossary **38**

Part I
Using the analyser

1 What ajku measures

The input to this instrument is a pair of electrical voltages. A balanced optical receiver may produce those voltages, but the FPGA does not know whether the optical signal is vacuum, a coherent beam, squeezed light or technical noise. That distinction comes from the experimental arrangement, reference measurements and a measurement model. Calling the two connectors I and Q is a channel convention; it does not by itself turn their voltage distribution into a quantum phase-space distribution.

The six first-priority workflows are available: LO-power-dependent spectral fitting; the full I/Q spectral matrix; joint/marginal and tail diagnostics; amplitude and squared-amplitude correlations; third/fourth-order cumulants; and drift monitoring with raw follow-up captures. The instrument is a CW measurement instrument, release 1.0 (FPGA/driver 1.0), with processing revision 5 on the host. It is not a random-bit extractor, optical phase-lock controller or automatic quantum-state tomograph. Versions 0.3 to 0.5 added the features of Section 1.3; release 1.0 was qualified as a whole (Section 14). The PPS and slow inputs have not been tested with a source.

1.1 Three different records of the experiment

Product	Time coverage	What it describes
FPGA statistics	Consecutive 2^{20}-slot blocks while acquired	Raw digitized exact sums through fourth order (means, powers, covariance, all third/fourth-order monomials), extrema and rail occupancy
FPGA histograms	2^{24}-slot epochs	Raw joint and marginal counts in a configurable window (full range by default), with out-of-window counts
DDR captures	Finite selected intervals	Every native I/Q code in those intervals; the basis for PC spectra, filtering and higher-order analysis

Both ADCs run at a nominal 125 MSa/s, with 8 ns clock slots. The signed native range is −8192 through 8191; all 14 physical bits are retained. A pair occupies four bytes in a raw recording. The default capture contains 2^{21} pairs: 8 MiB and 16.777 216 ms of waveform. A 32 MiB capture contains 2^{23} pairs and 67.108 864 ms. These are acquisition durations, not download times.

The statistics block lasts 8.388 608 ms; a histogram epoch lasts 134.217 728 ms. A 64-block FPGA statistics queue retains about 0.537 s. If full, it drops new completed blocks and increments a counter. Each service subscriber also has a bounded queue. Histograms are latest-wins until claimed by a reader; acquisition continues in other banks. A missed histogram epoch is counted. Neither mechanism is an unconditional lossless waveform stream.

A spectrum limit is not an analogue filter

Host science spectra end at 25 MHz. Continuous FPGA statistics and histograms still include the raw digitized response across the ADC band. Cropping a spectrum does not filter a time trace and cannot remove signals that already aliased at the ADC. The optional host FIR is a separate, explicitly identified analysis branch.

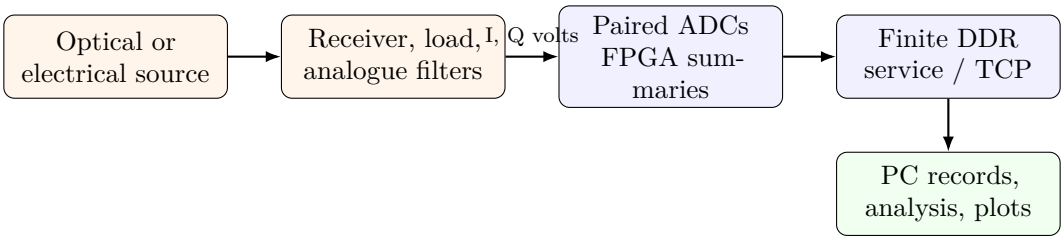


Figure 1: The user supplies and characterizes the source/receiver chain. The FPGA preserves the raw electrical measurement; the PC computes and stores derived products.

1.2 Signal path and responsibilities

Use IN1 for I and IN2 for Q. The two channels share a sampling clock. A sample index is an integer coordinate within an FPGA session, not a UTC timestamp. Session IDs separate reloads, and configuration revisions separate accepted changes. Keep both identifiers with every interval. Nominal sample rate is not a calibration of the board oscillator's absolute frequency.

The current capture engine is finite, without an overwrite ring or pretrigger waveform history. Completion means issued AXI writes have drained. Source/configuration changes abort captures and discard partial statistics blocks. The software never treats an arbitrary timeout as proof that hardware writes have stopped.

1.3 Windows, timing, slow references, higher orders and spectra (0.3–0.5)

Continuous spectra (0.5). The FPGA also transforms every sample: it windows $z = I + \mathrm{i}Q$ in non-overlapping frames of 1,024 samples (Hann), and over each 134 ms epoch of 16,384 frames sums, exactly as integers, the power $|Z_k|^2$ of every bin and the product $Z_k Z_{N-k}$ of each bin with its mirror. From these `spectrum()` returns the one-sided auto spectra of I and Q, their cross spectrum $\overline{X_I^* X_Q}$ (the convention of Section 3.2) and the complex and complementary spectra of z, in V^2/Hz to 25 MHz, in 122 kHz bins (Hann ENBW 183 kHz). They use every frame without a missing sample, not the 1–2 % of time that is captured, so at this resolution their statistical error is 5–8 times smaller in the same time (0.78 % per bin per epoch for Gaussian noise). Finer resolution, low-frequency $1/f$ noise and narrow lines still need host spectra of long captures. Like the other FPGA products they are raw broadband: aliasing is set by the analogue chain, the converter offset leaks into bins 0 and 1, and no converter-skew correction is applied. The fixed-point arithmetic is reproduced bit for bit by `ajku.model`; its rounding adds about 0.3 % to the noise power for an ADC noise of 1.3 codes rms.

Continuous fourth-order statistics (0.4). Every statistics block also holds the exact sums of I^3, I^2Q, IQ^2, Q^3 and $I^4, I^3Q, I^2Q^2, IQ^3, Q^4$ in native codes. From them the PC forms the block's central moments and joint cumulants exactly (integer numerators, one final division), hence skewness, excess kurtosis, $\kappa_{22}/(\kappa_{20}\kappa_{02})$, $\langle r^4\rangle/\langle r^2\rangle^2$ and the kurtosis of any projection X_θ, for every block and for any longer interval by summing blocks. Every sample contributes, so over the same time the statistical error of these non-Gaussianity tests is roughly ten times smaller than from captures alone. They remain raw broadband statistics of the digitized voltages (Section 3). This 0.4 addition passed digital board checks: an actual ADC block matched all fourteen FPGA sums exactly.

Histogram windows. By default a joint bin spans 256 native codes per axis and a marginal bin 16. An LO-off or electronic-noise reference with σ of a few tens of codes then falls into one or two joint bins. A window, given in connector volts as a centre and a span per axis, narrows the FPGA histograms to 2^k codes per bin (joint 1–256, marginal 1–16); the span is rounded up. One code per marginal bin shows missing codes and differential nonlinearity directly. Samples outside the window are counted (joint outside; below and above for each marginal), and every snapshot carries the window it was accumulated with. The panel's *Centre* $\pm 6\sigma$ *on latest block* stages a window around the latest block's means.

Slow inputs and PPS. Each statistics block also holds the mean of the four slow inputs AI0–AI3 (connector E2, 0 V to 7 V nominal; several hundred XADC conversions per block) and the count and sample index of rising edges on DIO0_P (3.3 V, not 5 V tolerant). Connect a receiver's DC monitor outputs or a tap photodiode on the LO to AI0–AI3 to log a relative LO-power reference beside the noise statistics; it becomes a power only after its own calibration against a power meter. A disciplined PPS measures the actual sample rate (index differences between edges) and, if UTC-aligned, ties sample indices to UTC within ±1 clock plus the input delay.

Timed captures. A capture can start at a chosen sample index instead of when the request arrives; `record.py --timed` uses this for a seeded, signal-independent schedule without host jitter.

Stability and skew. `tools/iqstability.py` computes Allan and modified deviations of the block statistics (means, variances, covariance, relative noise power, slow inputs) from 8.4 ms to the length of a gap-free recording. When the board calibration holds measured converter delays for both inputs, processing revision 5 adds skew-corrected cross-spectral products (`skew_*`); a 1 ns skew is 9° of cross-phase at 25 MHz. Receivers and cables are not included.

2 First measurement

2.1 Equipment and installation

You need a Red Pitaya STEMlab 125-14 with the ajku overlay/plugin and common service installed, an Ethernet connection, and a PC with Python. The board runs PYNQ; host analysis uses NumPy. Plots require Matplotlib, ipywidgets and Jupyter; HDF5 export additionally requires h5py. These analysis packages belong on the PC. Normal use does not require Vivado. The repository's `common/SETUP.md` explains board installation. From release 1.0 the SD image builder includes ajku; an image built before 1.0 does not contain it.

On a development machine the uniform `make install-service` target installs the overlay and restarts the common service with ajku selected. It is an installation action: a service restart can load the FPGA and restore settings. It is not needed for each measurement. Use the repository-root `config.mk` for the board address; a design-local file can override it and command-line Make assignments take precedence. Python programs use an explicit address.

Before connecting receivers, check their permitted load, output amplitude and DC offset, the board's LV/HV jumpers, and the required analogue anti-alias filters. Red Pitaya inputs are high impedance. A receiver specified for a 50 Ω load generally needs the appropriate external termination; check the voltage at the actual connection. The software range declaration does not move a jumper. Leave headroom for stochastic tails and transients as well as the RMS signal.

The outputs are not LO actuators in this design. Applying the acquisition preset or `SAFE` selects ADC input and calibrated-zero DAC words. Loading starts at digital zero; connector offsets, startup and unconfigured behaviour are not a precision zero-voltage guarantee. Board calibration can intentionally drive reference levels on both outputs. Disconnect equipment that must not see those levels, and read the root README's equipment section before calibration or changing designs.

2.2 Command-line check

From the `ajku/` directory in your chosen host Python environment:

```
export PYTHONPATH="$PWD/../common/python:$PWD/control:$PWD/tools"
export BOARD_HOST=192.168.0.102  # replace with your board's address
python3 -m rpfpga.cli --host "$BOARD_HOST" --preset acquire
python3 -m rpfpga.cli --host "$BOARD_HOST" --statistics
```

An initial statistics or histogram request can precede the first completed block; wait for the declared block/epoch duration and request again. Record the reported hardware/driver version, build ID, common-software version, source and calibration. The internal counter source is useful for digital

integrity checks, but its values are diagnostic native codes and must never be interpreted as voltages or light.

The smallest useful Python example retains the raw data:

```
import os
from ajku import CLIENT
from ajku.recording import save
from iqanalysis import samples, spectra

with CLIENT(os.environ["BOARD_HOST"]) as dev:
    context = dev.snapshot()
    metadata, raw = dev.capture(2**21,
                                expected_revision=context["revision"])
save("first.iqbin", metadata, raw)  # a new path; no overwrite
volts = samples(metadata, raw)
result = spectra(volts, nperseg=65536)
print(result["band_power"], result["band_covariance"])
```

With measured board calibration, `samples` returns connector volts. Nominal calibration is labelled as such; counter data are not a voltage measurement. The guarded request refuses a changed revision rather than acquiring under a stale assumption. There is no automatic replay after a failed network request.

2.3 Notebook workflow

Launch `jupyter notebook tools/live.ipynb` from the same environment. Absolute `PYTHONPATH` entries remain valid if the notebook kernel starts in `tools/`. Opening the panel does not start capture.

1. Read back source, jumper ranges and calibration. Source edits are staged until **Apply source**. If another controller has changed the revision, refresh, discard/reconsider the draft and apply deliberately.
2. Set pairs, FFT length, window and maximum frequency. Defaults are 2^{21} pairs, 65,536-point Hann spectra and 25 MHz. Enter the actual condition and **Chain JSON**: receiver, coupling, attenuation, termination, filters and IDs.
3. Select **Apply analysis** or **Acquire once**. **Live bursts** repeats complete jobs; its delay is after a job, not a guaranteed sampling cadence.
4. Enable FIR, temporal lags or distribution errors when those products are needed. Choose **Raw ADC** or **FIR / decimated**, and the **Spectra / lags**, **Matrix / phase** or **Distributions** panel. Switching a view neither acquires a new record nor changes its processing.
5. Use a new run directory and **Record run** to retain the statistics stream and summaries. Acquire/live bursts add waveform records; monitored recording can instead request event-selected follow-ups (Section 6).

The FPGA histogram, latest-block display and captured waveform need not cover the same interval. Their session/revision must agree before the dashboard overlays them. Ordinary displayed variance history samples the latest block at refresh time; only the recording's statistics stream provides the complete retained sequence. Plotting and widget actions share the notebook event loop. A redraw can briefly delay an input; the shared poller permits one outstanding UI publication so slow rendering cannot build an unbounded callback backlog. The separate statistics recorder continues to drain its stream. After changing views, check the plot's raw/FIR title as the new image appears.

2.4 Stopping and ownership

Turn off live bursts to stop future requests. Turn off recording to stop admitting recording jobs. Already admitted work may finish; wait for the manifest to become `complete` or `complete_with_`

`errors`. Closing the panel signals its workers to stop. A Python owner can join the returned recorder. A completed file/job lifecycle does not by itself mean there were no stream losses.

Use a normal configuration change or `SAFE` for an intentional capture abort; the driver still drains accepted writes. Do not kill the board service to implement ordinary Stop. A second acquisition is `BUSY`, not queued behind the first. An accepted edit can cancel an acquisition, and a reload makes old clients stale. Reconnect explicitly and inspect the state after a transport failure. Active-DMA process-crash recovery is outside the qualified operating envelope. Closing a browser tab is not the same as running cleanup in its Python kernel. Stop live/recorded work and use the notebook's explicit cleanup cell when ending the interactive session; the cell is commented to make ordinary **Run All** safe for opening the dashboard.

3 Reading spectra, moments and distributions

3.1 From native codes to volts

The board calibration uses signed 16-bit stream codes, whereas the recordings use right-aligned native 14-bit codes c. For channel j, let the shared calibration have ADC slope K_j (stream codes per volt) and offset A_j (stream codes). Then

$$v_j = \frac{4c_j - A_j}{K_j}. \tag{1}$$

The factor four changes the code convention, not the physical resolution. It does not create extra measured bits. Calibration and declared jumper range are captured with the data. A later change to the board calibration does not alter an old raw record's calibration context.

This conversion applies gain and offset. It does not automatically de-embed an analogue frequency response, equalize the receiver channels or rotate optical quadratures. Common calibration can store response/delay fields for other designs; their presence is not evidence that ajku has applied those corrections. User-derived corrections need their own explicit model and provenance (Section 13).

3.2 Welch spectra and the I/Q matrix

For each contiguous capture, the host divides the data into overlapping segments, optionally removes each segment's mean, applies a window and averages powers. For a segment of L points with sampling rate f_s, define

$$X_j[k] = \sum_{n=0}^{L-1} w[n](v_j[n] - \bar{v}_j)\mathrm{e}^{-2\pi \mathrm{i} kn/L}, \qquad U = \sum_n w[n]^2. \tag{2}$$

For positive interior frequency bins the one-sided auto/cross estimates are

$$S_{jj}[k] = \frac{2\,\overline{|X_j[k]|^2}}{f_s U}, \qquad S_{IQ}[k] = \frac{2\,\overline{X_I[k]^* X_Q[k]}}{f_s U}. \tag{3}$$

The factor is one at DC and, for even L, Nyquist. The bar means averaging over segments, not taking a coherent average before squaring. Without detrending, the unsubtracted segment is used instead. Densities have units V^2/Hz; square-root auto-densities have units $\mathrm{V}/\sqrt{\mathrm{Hz}}$.

The stored matrix convention is

$$\boldsymbol{S}(f) = \begin{pmatrix} S_{II} & S_{IQ} \\ S_{IQ}^* & S_{QQ} \end{pmatrix}, \qquad S_{IQ} = \overline{X_I^* X_Q}/(\text{normalization}). \tag{4}$$

It is Hermitian. Its eigenvalues describe electrical spectral combinations, with phase information in the off-diagonal elements. If $Q(t) = I(t - \tau)$, the cross-phase is $-2\pi f\tau$. This sign matters when checking a delay against another program. In the lag-domain display, a positive I,Q correlation peak means Q follows I.

The Fourier-bin width is $\Delta f = f_s/L$. The equivalent noise bandwidth of the window is $B_{\mathrm{ENBW}} = f_s U/(\sum_n w[n])^2$, which is $1.5\Delta f$ for the periodic Hann window. At 125 MSa/s with $L = 65536$, these are about 1.907 35 kHz and 2.861 02 kHz. ENBW is not the full measurement bandwidth and not the FIR noise bandwidth. A narrow coherent tone spreads over the window's spectral response; integrate its admitted bins rather than treating its peak as a broadband noise density.

The continuous FPGA spectra (Section 1.3) are Eq. (3) with $L = 1024$, no overlap and no detrending, over every frame of an epoch; their Hann window is sampled at half-integer points, which has the same three-bin transform and the same $U = 3L/8$. `spectra(v, nperseg=1024, overlap=0, detrend=False)` of a capture estimates the same quantities and serves as a cross-check. Non-overlapping Hann frames of white noise are independent, so the number of frames is the number of averages.

Band powers use the actual admitted grid:

$$P_I = \sum_{f_k \le f_{\max}} S_{II}[k]\Delta f, \quad P_Q = \sum_{f_k \le f_{\max}} S_{QQ}[k]\Delta f, \quad C_{IQ} = \sum_{f_k \le f_{\max}} \Re S_{IQ}[k]\Delta f. \tag{5}$$

The integrated real covariance matrix has P_I, P_Q on its diagonal and C_{IQ} off diagonal. Full-Nyquist Parseval products check the normalization against window-weighted segment powers; a science-band integral need not equal the raw time-domain variance. Detrending also removes segment DC information. Report the FFT/window/detrending and frequency interval along with a quoted noise power.

For $z = I + \mathrm{i}Q$, signed-frequency complex spectra retain positive and negative frequencies separately. The complementary spectrum involves $\overline{Z(f)Z(-f)}$, without conjugating the second factor. These products expose image/asymmetry and non-circular structure in the electrical signal. They use a two-sided convention: do not apply an additional one-sided factor of two.

3.3 Coherence and phase uncertainty

The magnitude-squared coherence is

$$\gamma^2(f) = \frac{|S_{IQ}(f)|^2}{S_{II}(f)S_{QQ}(f)}. \tag{6}$$

A shared tone can have high coherence even when noise in the rest of the band is uncorrelated. A single segment trivially gives unity coherence where both transforms are nonzero; this is not evidence of a physical connection. More averages are needed.

The host computes an overlap-corrected effective number of averages K_{eff} under white-input assumptions. With per-bin level α, the displayed approximate null threshold is

$$\gamma^2_{\mathrm{thr}} = 1 - \alpha^{1/(K_{\mathrm{eff}}-1)}, \qquad K_{\mathrm{eff}} > 1. \tag{7}$$

Where exceeded, the phase is reported with asymptotic standard error

$$u_\phi \simeq \sqrt{\frac{1-\gamma^2}{2K_{\mathrm{eff}}\gamma^2}}. \tag{8}$$

DC/Nyquist and one-segment significance are excluded. The assumptions are Gaussian, stationary channels under the null and a white-input overlap approximation; coloured or nonstationary records can violate them. The threshold is pointwise, not a family-wise guarantee across thousands of bins. The phase error excludes systematic channel-delay and receiver calibration uncertainty.

3.4 Explicit FIR filtering and decimation

An optional real, symmetric, unity-DC Kaiser FIR acts identically on I and Q. The default passband ends at 20 MHz, the stopband begins at 25 MHz, and the decimation factor is $M = 2$. The current

default design has 137 taps, numerically checked stopband attenuation 85.36 dB and peak-to-peak passband ripple 0.000 861 dB. Its one-sided white-noise bandwidth is 22.03 MHz. These describe the digital coefficients, not an optical/analogue measurement.

Filtering retains both the raw and derived branches. Raw codes, raw rail counts and continuous FPGA summaries remain available. Host decimation does not reduce board-to-PC wire bytes. A "25 MHz stop edge" is not a brick-wall 25 MHz passband: the transition and residual response contribute to filtered time-domain moments.

For L_h taps and an N-sample record beginning at raw index n_0, the valid convolution is

$$y[j] = \sum_{k=0}^{L_h-1} h[k]x[j + L_h - 1 - k], \qquad 0 \le j \le N - L_h. \tag{9}$$

Let $D = (L_h - 1)/2$ and let p be the requested raw-counter residue. The first admitted offset is $j_0 = (p - n_0 - D) \bmod M$. Retained output centres are

$$n_r = n_0 + D + j_0 + rM. \tag{10}$$

These are exact integer coordinates, including for counters beyond the exact integer range of float64. Each centre requires a complete input window; there is no reflected/zero padding or borrowed history across captures. Centre-referenced time is an offline delay convention, not a zero-latency filter.

An 8 MiB default record yields 1,048,508 derived pairs. The same 65,536-point FFT then has half the raw branch's bin width and a different segment count. PSD units remain V^2/Hz; no extra factor M belongs in the density. Comparing isolated raw/FIR peak heights with different frequency resolutions can misstate a filter gain. Use its saved response or appropriately integrated powers.

Choose a stop edge no higher than the new Nyquist frequency. Decimation without appropriate filtering would introduce new aliases. The profile records exact coefficients, sampled response, coefficient identity and edge policy. A loaded reference can stage its exact FIR; this does not automatically copy its other analysis choices. See `docs/FILTERING.md` for all design limits and replay APIs.

3.5 Time moments, cumulants and circularity

For centred channel values $i = I - \bar{I}$ and $q = Q - \bar{Q}$, write $\mu_{ab} = \overline{i^a q^b}$. Second moments form the covariance matrix. The four third-order cumulants are $\mu_{30}, \mu_{21}, \mu_{12}, \mu_{03}$. Fourth order gives

$$\kappa_{40} = \mu_{40} - 3\mu_{20}^2, \qquad \kappa_{31} = \mu_{31} - 3\mu_{20}\mu_{11}, \tag{11}$$

$$\kappa_{22} = \mu_{22} - \mu_{20}\mu_{02} - 2\mu_{11}^2, \qquad \kappa_{13} = \mu_{13} - 3\mu_{02}\mu_{11}, \tag{12}$$

$$\kappa_{04} = \mu_{04} - 3\mu_{02}^2. \tag{13}$$

Standardized cumulants divide by $\sigma_I^a \sigma_Q^b$ where those scales exist. These are finite-record population-moment estimators; higher-order cumulants are not silently replaced with unbiased k-statistics. Gaussian distributions have vanishing population cumulants above second order, but their finite-record estimates do not vanish exactly. Nonzero cumulants can also arise from drift, clipping, intermittent interference or mixtures of classical states.

For centred $z = i + \mathrm{i}q$, the ordinary variance is $C = \langle |z|^2 \rangle = \mu_{20} + \mu_{02}$ and the pseudo-variance is $P = \langle z^2 \rangle = \mu_{20} - \mu_{02} + 2\mathrm{i}\mu_{11}$. The displayed circularity magnitude is $|P|/C$ when $C > 0$. Unequal channel gain or a non-quadrature receiver can make it large. Whitening the measured ellipse would remove that evidence; it would not establish optical quadrature calibration or eliminate quantum squeezing.

3.6 Histograms, tails and uncertainty

Since version 0.4 the FPGA statistics blocks carry exact third- and fourth-order sums, so the cumulants below are also available continuously for the raw broadband signal (`statistics()["higher_`

`order"]`, `model.higher_order`); captured records add band-limited (FIR) versions, tails and batch-jackknife errors.

The continuous FPGA joint histogram has 64×64 bins, and each marginal has 1,024 bins. Over the full range (the default), joint bins span 256 native codes per axis, marginal bins 16; a window narrows them down to one code per bin (Section 1.3). A captured full-code marginal accounts for each of the 16,384 native codes. Counts are not density-normalized.

Optional captured diagnostics additionally standardize each centred axis by its own fitted standard deviation, with default range $\pm 6\sigma$ and 64 bins per axis. Out-of-view counts remain explicit. This plotting scale is neither optical calibration nor covariance whitening. Quantiles, empirical-CDF points with ties, Gaussian Q–Q coordinates, and positive/negative tail counts are saved. A Q–Q tail deviation is a reason to inspect the records and measurement conditions, not an automatic non-classicality witness.

A fitted correlated Gaussian supplies bin-integrated expected counts. Singular covariance or an unresolved numerical integral gives an unavailable fit, not an arbitrary regularization. The residual display $(N - \lambda)/\sqrt{\lambda}$ masks expected counts below five. Its colour scale is descriptive, not an independent-bin chi-square test: parameters were fitted, samples can be correlated, and ADC codes are discrete.

Cumulant errors use a delete-one-contiguous-batch jackknife with leave-sample re-centring. For B equal complete batches and leave-batch estimates $\theta_{(-b)}$,

$$\widehat{\mathrm{SE}}(\theta) = \sqrt{\frac{B-1}{B} \sum_b (\theta_{(-b)} - \bar{\theta}_{(-)})^2}. \tag{14}$$

The full-record estimate, retained-subset estimates, bias estimates and SEs are separate products. Base, twice-base and four-times-base batch lengths are attempted, each needing at least eight complete batches. Derive the omitted remainder at each scale from its saved length/count; no incomplete batch is invented.

Tail errors use variation of batch indicator means while holding the full-record fitted thresholds fixed. They exclude uncertainty of the thresholds themselves. They are withheld below ten events in four occupied batches. Zero observed events therefore do not produce a zero error or an IID upper bound. Increase observation time and inspect dependence rather than treating millions of correlated samples as millions of independent trials.

3.7 Amplitude, squared-amplitude and complex correlations

The lag convention is $R_{ab}[k] = \overline{a[n]b[n+k]}$. Amplitudes are centred with the whole-record channel means; squared-amplitude fluctuations are $i^2 - \overline{i^2}$ and $q^2 - \overline{q^2}$. The radial quantity is $i^2 + q^2 - \overline{i^2 + q^2}$. The module stores their covariance/correlation functions and ordinary/complementary complex correlations. These electrical squared amplitudes must not automatically be called photon counts or optical $g^{(2)}$.

FFT computation is divided into contiguous blocks. Pairs crossing a block or capture boundary are omitted, and each lag carries its actual pair count. Longer lags thus have fewer pairs; gaps are never concatenated. A filtered branch's lag samples and batch lengths use its own sampling rate, with raw-clock stride also recorded.

Temporal-analysis batch SEs apply to the saved means, variances and covariance; they are not a universal confidence band for every plotted lag. Amplitude and power can have different correlation times, so one "effective sample number" cannot be reused for all statistics. Inspect error estimates as the batch length increases. All batch methods here require approximate stationarity and weak dependence between sufficiently long batches; none includes a complete systematic calibration budget.

4 References, recordings and reproducibility

4.1 Reference comparisons

Establish the reference condition physically, then acquire and save it. For an electronic-background reference, document which optical ports were blocked, whether the receiver remained powered, and whether its gain, load and coupling changed. Typing "vacuum" or "LO off" does not verify that condition.

A new reference directory contains the original raw capture, processed arrays, metadata and hashes. Comparisons require matching source/configuration, input ranges, board-calibration identity, recorded DAC words, declared measurement-chain identity and the full processing contract. The latter includes revision, frequency grid, filter coefficients/phase, temporal settings and diagnostic settings. Condition labels may differ. A change to calibration, filter or analysis requires a compatible new reference, or explicit reanalysis of the old raw data with matching settings.

The displayed clearance is $10\log_{10}(S/S_{\mathrm{ref}})$ and the excess is the *signed* difference $S - S_{\mathrm{ref}}$. Negative differences are retained; undefined ratios are unavailable. Subtraction does not prove quantum origin. A reference captured yesterday can be a poor physical baseline today even if its software contract matches perfectly.

```
python3 tools/iqmeasure.py --host "$BOARD_HOST" --output /data/run-off \
  --filter --lags 128 --diagnostics --condition "actual LO-off condition" \
  --save-reference /data/reference-off
python3 tools/iqmeasure.py --host "$BOARD_HOST" --output /data/run-on \
  --filter-profile /data/reference-off/filter.json --lags 128 --diagnostics \
  --reference /data/reference-off --condition "actual LO-on condition"
```

These are examples of new paths, not supplied calibration data. Include the same `--chain-json` metadata in both commands when declaring your optical chain. CLI measurement commands use the existing source; select ADC first for physical data.

4.2 Run contents and coverage

```
run.json                    lifecycle, capture index, hashes, losses, coverage
captures/000000.iqbin        original paired native codes and captured context
captures/000000.npz          processed arrays and JSON provenance
statistics.jsonl            complete retained statistics packets
histograms.jsonl            retained histogram snapshots
events.jsonl                headers, comparisons, failures, trigger outcomes
drift.jsonl                 aggregate windows for monitored runs
reference-<id>.npz          processed reference used by a comparison
```

NPZ loading disables pickle. A capture is published in the manifest after its raw and processed files are complete. Failed writes and unfinished admitted jobs remain errors rather than successful captures. Preserve the original reference directory if later reanalysis of its raw data will be needed. Hashes establish byte identity, not truth of an operator-supplied optical-condition label.

Raw coverage is retained ADC time divided by host elapsed recording time. The numerator is the union of retained sample intervals within each session, preventing double counting when a record is saved again. A waveform gap contains no recoverable raw samples. Continuous summaries may describe the gap, but cannot reconstruct its waveform or its higher-order temporal structure. Host timestamps are not synchronized ADC UTC time. A `complete` run can still contain explicitly counted FIFO losses, subscriber losses or missed histogram epochs.

```
python3 tools/iqmeasure.py --host "$BOARD_HOST" --output /data/run-many \
  --captures 100 --pipeline --filter --lags 128 --diagnostics \
  --condition "document the actual experiment"
```

The pipeline admits at most two jobs, including downloading, queued and processing work. One downloader completes each transfer before requesting another capture. Backpressure limits memory and slows admission. Settings, conditions, references and recording ownership are pinned per job.

This is throughput-driven selection, not a fixed-cadence unbiased sample of a fluctuating experiment. The separate `tools/record.py` supports a seeded schedule and skips missed deadlines; check its selection metadata when comparing runs.

4.3 Reanalysis and full-waveform export

```
python3 tools/iqmeasure.py --input /data/run-many/captures/000000.iqbin \
  --output /data/reanalysis --nfft 32768 --filter --diagnostics \
  --export-hdf5 /data/reanalysed.h5
```

Reanalysis preserves original capture context and does not claim new live coverage. Supply the intended condition/chain metadata when making a new processed reference; do not silently relabel an old physical experiment. NPZ traces are previews of up to 4,096 pairs. The optional HDF5 export contains native codes, all products and, when enabled, the complete derived FIR waveform with integer centre/stride metadata. The CLI exports its final measurement, not an entire multi-capture run into one file. Check the HDF5 root `complete` marker. It is set after write/flush completion, not a guarantee against power loss.

5 LO-power-dependent spectral measurements

The LO-level recorder associates repeated ADC captures with an independent optical power reading. A level is first established externally; ajku does not infer watts from a DAC setting, attenuator command, ADC RMS or detector full-scale utilization. The user supplies the actual meter/actuator drivers for automated sweeps.

An illustrative reading file has the following form; replace every physical value and provenance field with your own measurement:

```
{"origin":"measured","power_w":0.001,"standard_uncertainty_w":0.00001,
 "instrument":"your meter ID","method":"your calibrated measurement",
 "lo_state":"on"}
```

Unknown uncertainty is `null`, not silently zero. For unknown power, set `origin` to `unavailable`, `power_w` to `null`, and declare the LO state. It can be recorded but cannot be fitted. A declared off state requires zero power; an on state requires positive power. Synthetic inputs require an explicit fitting opt-in and remain labelled synthetic in outputs.

```
python3 tools/iqsweep.py record --host "$BOARD_HOST" \
  --directory /data/lo-level-001 --power-json /data/power-001.json \
  --captures 4 --interval 0.25 --filter --diagnostics \
  --condition "signal port condition, receiver settings, optical arrangement"
python3 tools/iqsweep.py fit /data/lo-level-*/level.json \
  --output /data/lo-fit.npz --power-scale-relative-std 0.01 \
  --plot /data/lo-fit.png
```

The meter-scale value above is illustrative too. Omitting it means that part of the uncertainty is unknown. Use `--view filtered` for the filtered fit plot. Saved-level fitting is host-only and does not open a board connection, including from the GUI. The **LO-power calibration** tab exposes level recording and fitting. Stop ordinary acquisitions/recording before starting a level. **Stop level** prevents future repetitions and lets admitted work finish; a cancelled level is retained as incomplete rather than silently promoted to complete data.

For automation, `iqsweep.record_sweep` takes a sequence of requested controls, `set_control` and independent `read_power` callbacks, settling time and an optional cleanup callback. It records actuator readback separately from watts, and checks a second power reading after each level. With manual levels and no after-reading, stability is explicitly operator-assumed. The software does not randomize the control order; interleave/repeat levels and references to reveal drift.

5.1 Model, errors and interpretation

The fitted response is

$$S(f;P) = A(f) + B(f)P + C(f)P^2. \tag{15}$$

There are separate fits for $S_{II}, S_{QQ}, \Re S_{IQ}, \Im S_{IQ}$ and for integrated I/Q powers, covariance and total power. Raw and filtered branches remain distinct. Integrated-band errors are fitted from each capture's band quantities; they are not made by pretending frequency bins have independent errors.

At least four distinct powers including a declared LO-off zero, and at least two complete captures per level, are required. Processing/calibration/chain contracts must agree. Clipped, trigger-selected, repeated or overlapping capture intervals are refused, and unstable levels are ineligible. This avoids silently dropping inconvenient observations, but absence of ADC clipping does not prove receiver linearity or independence of repeated captures.

The fit retains signed coefficients. Their units are response, response/W and response/W^2. OLS HC3 covariance describes statistical errors conditional on the polynomial model and independent stationary capture repeats. Meter-reading errors are linearized, shared by repeats within a level; a common multiplicative meter-scale error is propagated as a correlated variable. Unknown errors stay unavailable. The report includes residuals and a lack-of-fit/pure-error diagnostic, without an automatic p-value or assertion that the model is adequate.

If explicitly requested with `--quantum-assumptions`, the conditional model assigns BP to quantum noise and $A + CP^2$ to classical noise:

$$\mathrm{Q/C} = \frac{BP}{A + CP^2}. \tag{16}$$

It is reported only where the auto-component signs and positive power permit that interpretation. Negative fitted components remain visible. Linear power scaling alone is not proof of quantum origin; the report always retains `quantum_origin_verified=false`. This ratio is not an entropy or security bound.

6 Drift monitoring and event-selected captures

Drift monitoring drains the retained raw-statistics stream, accumulates exact integer sums into trend windows and evaluates configurable rules. Available metrics are means, central variances, covariance, RMS including DC, rail/near-rail fractions and circularity. Volts and native codes are distinct choices; voltage rules refuse the counter source. Means/RMS use the chosen amplitude unit, variances/covariance its square, and occupancy/circularity are dimensionless.

An illustrative rules file is:

```
{"units":"volts","baseline_blocks":32,"trend_blocks":128,"pre_blocks":64,
 "max_trigger_age_s":2.0,
 "rules":[{"name":"I-offset","metric":"mean_i","mode":"deviation",
           "threshold":0.001,"hysteresis":0.0002,
           "persistence_blocks":3,"refractory_s":2.0}]}
```

Choose thresholds from your experiment, not from these example numbers. Above/below rules compare absolute values; deviation rules compare with a frozen baseline from consecutive valid blocks. Persistence counts consecutive violations. Hysteresis sets the clearing/rearming threshold, and refractory time uses the sample clock. Missing/invalid blocks, changed loss counters and context changes split trends and reset baselines. Empty rules provide trend-only recording.

```
python3 tools/iqdrift.py --host "$BOARD_HOST" --rules /data/rules.json \
  --output /data/drift-run --seconds 120 --filter --diagnostics --lags 128
python3 tools/iqdrift.py --input /data/drift-run/statistics.jsonl \
  --rules /data/rules.json --output /data/replayed-events.jsonl
```

In the GUI, enable **Drift monitoring during Record run**, enter rules and start recording. Periodic live bursts and automatic monitored recording are mutually exclusive in the GUI; another controller can still supersede or contend for capture.

Only one automatic capture job is admitted at a time. Extra detections receive explicit busy skips. Current statistics and revision-guarded capture admission reject stale context or an old backlog; requests are not automatically retried as though the event were still current. Stop prevents new jobs, while admitted work drains and is saved or recorded as failed before sealing the run.

The event contains its first/end sample indices, rule, value, baseline and preceding *statistics*. The capture stores its actual delay, overlap/late flags and `pretrigger_raw_pairs=0`, with selection `drift_followup`. Measured host delays in short qualification checks were about 0.25 s; a short event may already have ended. Replay reproduces deterministic event IDs/windows without requesting new captures. Never pool event-selected records into an unbiased noise distribution or ordinary LO-power sweep. Hardware pretrigger history is a separate feature, not an implicit property of these finite host buffers.

Part II
Continuous-variable quantum optics and its measurements

7 Optical modes, quadratures and quantum states

The voltage records become quantum-optical measurements only after specifying *which optical mode is measured, by which detector, with which calibration.* This section fixes conventions and develops that connection. Standard treatments are given by Gerry and Knight, Walls and Milburn, and the Gaussian-state review of Weedbrook *et al.* [1–3]; references may use different factors of two in their quadratures.

7.1 One mode is one harmonic oscillator

For one normalized optical mode, the annihilation and creation operators obey

$$[\hat{a}, \hat{a}^\dagger] = 1, \qquad \hat{n} = \hat{a}^\dagger \hat{a}, \qquad \hat{H} = \hbar\omega_0(\hat{n} + \tfrac{1}{2}). \tag{17}$$

The positive-frequency electric field of that mode is proportional to $\hat{a}\mathrm{e}^{-\mathrm{i}\omega_0 t}$; the dimensional proportionality factor depends on mode normalization and geometry. Photon number is dimensionless. A 125 MSa/s ADC does not resolve an optical carrier; the photodetector/LO arrangement converts optical information to an electrical band that the ADC can sample.

We use dimensionless quadratures

$$\hat{X} = \frac{\hat{a} + \hat{a}^\dagger}{\sqrt{2}}, \qquad \hat{P} = \frac{\hat{a} - \hat{a}^\dagger}{\mathrm{i}\sqrt{2}}, \qquad [\hat{X}, \hat{P}] = \mathrm{i}, \tag{18}$$

so $\hat{a} = (\hat{X} + \mathrm{i}\hat{P})/\sqrt{2}$. A rotated quadrature is

$$\hat{X}_\theta = \hat{X}\cos\theta + \hat{P}\sin\theta = \frac{\hat{a}\mathrm{e}^{-\mathrm{i}\theta} + \hat{a}^\dagger \mathrm{e}^{\mathrm{i}\theta}}{\sqrt{2}}. \tag{19}$$

The angle is relative to a phase reference, normally the LO. A free optical phase without a specified reference is not a measurable absolute angle.

For any state, $\Delta X\, \Delta P \geq 1/2$. The vacuum $|0\rangle$, defined by $\hat{a}|0\rangle = 0$, has zero means and

$$\mathrm{Var}(X) = \mathrm{Var}(P) = \tfrac{1}{2}, \qquad \Delta X = \Delta P = 1/\sqrt{2}. \tag{20}$$

Some courses use $x = (a + a^\dagger)/2$, giving vacuum variance 1/4; some use $x = a + a^\dagger$, giving variance 1. Always translate the convention before comparing an experiment to a theoretical uncertainty bound. "One shot-noise unit" often sets the *variance ratio* of a chosen detector reference to one and does not alone specify the canonical commutator.

7.2 Pure states, mixtures and moments

A pure state is represented by $\rho = |\psi\rangle\langle\psi|$; a classical uncertainty about preparation gives a mixture $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, with $p_j \geq 0$ and $\sum_j p_j = 1$. Observable averages are $\langle O\rangle = \mathrm{Tr}(\rho O)$. A mixture is not the same as the coherent superposition $\sum_j c_j|\psi_j\rangle$: the latter retains off-diagonal interference terms.

Define $\hat{\boldsymbol{R}} = (\hat{X}, \hat{P})^T$, displacement vector $\boldsymbol{d} = \langle\hat{\boldsymbol{R}}\rangle$, and symmetrized covariance

$$V_{jk} = \tfrac{1}{2}\langle \Delta\hat{R}_j \Delta\hat{R}_k + \Delta\hat{R}_k \Delta\hat{R}_j\rangle. \tag{21}$$

The uncertainty principle can be written

$$\boldsymbol{V} + \frac{\mathrm{i}}{2}\boldsymbol{\Omega} \succeq 0, \qquad \boldsymbol{\Omega} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{22}$$

For one mode this implies $\det \boldsymbol{V} \geq 1/4$ as well as nonnegative variances. The energy is related to both mean and fluctuations:

$$\langle n \rangle = \frac{\operatorname{Tr} \boldsymbol{V} + \boldsymbol{d}^T \boldsymbol{d} - 1}{2}. \tag{23}$$

The subtraction is the vacuum contribution in this convention. Applying this formula to uncalibrated ADC variances would be meaningless.

From the course to the instrument

Ajku computes an *electrical* mean vector and covariance from two voltage channels. They have units V and V^2, not the dimensionless units of Eqs. (18)–(23). An optical-mode definition, detector model, gain/phase calibration, and treatment of detector noise are needed before identifying those numbers with $\boldsymbol{d}$ or $\boldsymbol{V}$.

8 Coherent states and displacement

A coherent state is an eigenstate of the annihilation operator, $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$, with complex amplitude α. In the number-state basis,

$$|\alpha\rangle = \mathrm{e}^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \tag{24}$$

Its photon-number distribution is Poissonian: $p_n = \mathrm{e}^{-|\alpha|^2}|\alpha|^{2n}/n!$, with $\langle n \rangle = \mathrm{Var}(n) = |\alpha|^2$. Its quadrature means are

$$\langle X \rangle = \sqrt{2}\Re\alpha, \qquad \langle P \rangle = \sqrt{2}\Im\alpha, \qquad \boldsymbol{V} = \tfrac{1}{2}\mathbf{1}. \tag{25}$$

Increasing the coherent amplitude moves the centre without changing the ideal quadrature-noise width. Relative intensity fluctuations decrease as $1/\sqrt{\langle n \rangle}$, while the absolute quadrature vacuum variance remains 1/2. An experimental laser may additionally have classical amplitude/phase noise.

8.1 The displacement operator

The unitary displacement operator is

$$\hat{D}(\alpha) = \exp(\alpha\hat{a}^\dagger - \alpha^*\hat{a}), \qquad \hat{D}^\dagger(\alpha)\hat{a}\hat{D}(\alpha) = \hat{a} + \alpha. \tag{26}$$

It creates a coherent state from vacuum, $|\alpha\rangle = \hat{D}(\alpha)|0\rangle$. More generally it shifts the quadrature means by $(\sqrt{2}\Re\alpha, \sqrt{2}\Im\alpha)$ without changing centred covariance or higher centred cumulants. Displacements compose according to

$$\hat{D}(\alpha)\hat{D}(\beta) = \exp\left[\frac{\alpha\beta^* - \alpha^*\beta}{2}\right] \hat{D}(\alpha + \beta). \tag{27}$$

The extra phase reflects noncommuting conjugate translations. It cancels in the density matrix of an isolated displaced state, but relative phases matter in interference experiments.

Optically, mixing a signal with a strong coherent field at a highly transmitting beam splitter approximates displacement. Finite reflectivity also attenuates the signal, so a real displacement apparatus has a loss model. Adding a DC number to an acquired voltage trace is only a change of data coordinates, not an optical state-preparation operation.

Displacement and a measured trace

For $\alpha = 1 + \mathrm{i}/2$, $\boldsymbol{d} = (\sqrt{2}, 1/\sqrt{2})$ and $\langle n \rangle = 1.25$. Each ideal quadrature still has variance 1/2. A homodyne phase scan has mean $\langle X_\theta \rangle = \sqrt{2}\cos\theta + (1/\sqrt{2})\sin\theta$. An unknown electronic offset shifts the same measured mean. Separate offset and phase calibration is therefore essential before interpreting a displaced cloud.

Phase-averaging a coherent state is another useful distinction. Randomizing its phase gives a mixture distributed around a ring, with zero average field but nonzero energy. A zero I/Q mean is therefore not evidence that the optical input was vacuum. It can also result from drift over many phases or from AC coupling.

9 Gaussian states, squeezing and loss

A Gaussian state has a Gaussian phase-space characteristic function; its first and second moments completely determine the state. Vacuum, coherent, thermal and squeezed thermal states are examples. This quantum definition does not say that any Gaussian voltage histogram is a calibrated Gaussian optical state.

9.1 Thermal light

A single-mode thermal state with mean population $\bar{n}$ is

$$\rho_{\mathrm{th}} = \sum_{n=0}^{\infty} \frac{\bar{n}^n}{(1+\bar{n})^{n+1}} |n\rangle\langle n|, \qquad \boldsymbol{V} = (\bar{n} + \tfrac{1}{2})\mathbf{1}. \tag{28}$$

Its mean field vanishes and its photon-number variance is $\bar{n}(1+\bar{n})$. It is a classical mixture of coherent states with Gaussian-distributed amplitudes. Classical receiver noise can also broaden a measured Gaussian cloud; width alone does not distinguish thermal photons from detector noise.

9.2 Squeezing as a quadrature transformation

For $\zeta = r\mathrm{e}^{\mathrm{i}\phi}$, use the convention

$$\hat{S}(\zeta) = \exp\left[\frac{\zeta^* \hat{a}^2 - \zeta \hat{a}^{\dagger 2}}{2}\right], \qquad \hat{S}^\dagger \hat{a} \hat{S} = \hat{a}\cosh r - \mathrm{e}^{\mathrm{i}\phi}\hat{a}^\dagger \sinh r. \tag{29}$$

For squeezed vacuum $\hat{S}(\zeta)|0\rangle$,

$$\mathrm{Var}(X_\theta) = \tfrac{1}{2}\big[\cosh(2r) - \sinh(2r)\cos(2\theta - \phi)\big]. \tag{30}$$

The squeezed axis is $\theta = \phi/2$, with variance $\mathrm{e}^{-2r}/2$; the conjugate axis has variance $\mathrm{e}^{2r}/2$. Their product is $1/4$. The mean field is zero, but $\langle n\rangle = \sinh^2 r$: photons can be present without a coherent displacement.

Writing $\boldsymbol{R}(\theta)$ for the ordinary real plane-rotation matrix,

$$\boldsymbol{V} = \tfrac{1}{2}\boldsymbol{R}(\phi/2) \begin{pmatrix} \mathrm{e}^{-2r} & 0 \\ 0 & \mathrm{e}^{2r} \end{pmatrix} \boldsymbol{R}(\phi/2)^T. \tag{31}$$

For a squeezed thermal state, replace the factor $1/2$ with $\bar{n} + 1/2$. It is quadrature-squeezed below vacuum only if $(2\bar{n}+1)\mathrm{e}^{-2r} < 1$. An ellipse by itself is insufficient: a classical thermal ellipse may have *both* axes above the vacuum variance.

The convention for squeezing in decibels is a variance ratio,

$$R_\theta = \frac{\mathrm{Var}(X_\theta)}{1/2}, \qquad s_\theta = 10\log_{10} R_\theta. \tag{32}$$

Negative s_θ denotes reduced variance; "6 dB squeezing" usually means $s_{\min} \simeq -6\,\mathrm{dB}$. A statement using amplitude standard deviations must use $20\log_{10}$ instead, giving the same numerical result when conventions agree.

Displacement and squeezing do not generally commute. The state $\hat{D}(\alpha)\hat{S}(\zeta)|0\rangle$ has displacement α and squeezed covariance. For $\hat{S}(\zeta)\hat{D}(\alpha)|0\rangle$, the mean amplitude is instead $\alpha\cosh r - \alpha^* \mathrm{e}^{\mathrm{i}\phi}\sinh r$. State-preparation order therefore changes the mean even when the final covariance is the same.

9.3 Efficiency and phase noise

Optical loss/detection inefficiency is modelled as a beam splitter with vacuum:

$$\hat{a}_{\text{out}} = \sqrt{\eta}\,\hat{a}_{\text{in}} + \sqrt{1-\eta}\,\hat{v}, \quad \boldsymbol{d}_{\text{out}} = \sqrt{\eta}\,\boldsymbol{d}_{\text{in}}, \quad \boldsymbol{V}_{\text{out}} = \eta \boldsymbol{V}_{\text{in}} + (1-\eta)\mathbf{1}/2. \tag{33}$$

For pure squeezed input the variance ratios become $R_{\mp} = 1-\eta+\eta \mathrm{e}^{\mp 2r}$. Even arbitrarily strong input squeezing approaches the floor $1-\eta$. Electronic noise adds another contribution in the detector's units; it is not automatically included in the optical efficiency.

With phase error δ about the squeezed axis,

$$R(\delta) = R_- \cos^2\delta + R_+ \sin^2\delta. \tag{34}$$

Large anti-squeezing makes small angle errors costly. For zero-mean Gaussian phase jitter with variance σ_θ^2, its average is

$$\overline{R} = \tfrac{1}{2}(R_- + R_+) + \tfrac{1}{2}(R_- - R_+)\mathrm{e}^{-2\sigma_\theta^2}. \tag{35}$$

At small jitter, $\overline{R} \simeq R_- + (R_+ - R_-)\sigma_\theta^2$.

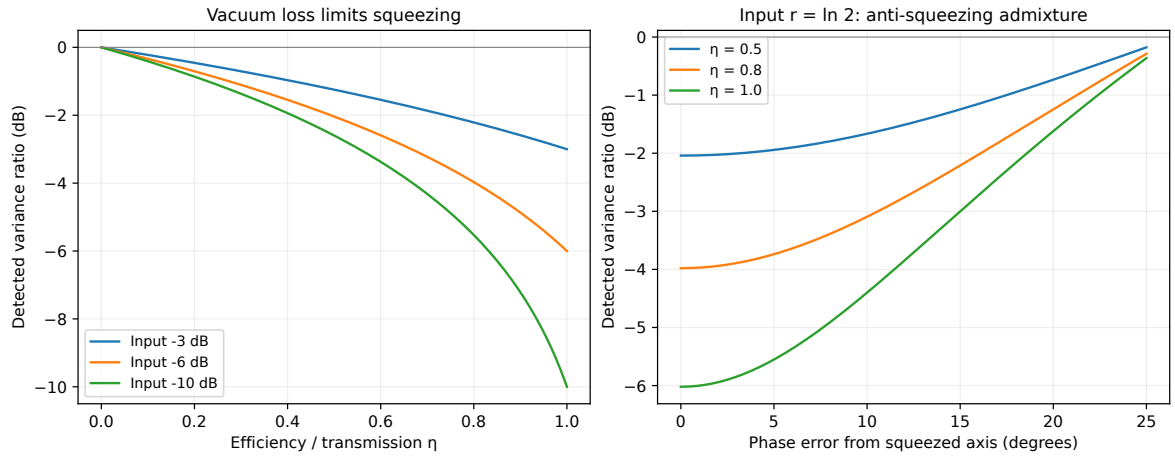


Figure 2: Computed.Ideal squeezing after vacuum loss, and the mixing of anti-squeezing into the measured quadrature by phase error. All curves are models, not a characterization of this board or a receiver.

A quantitative loss example

Take $r = \ln 2$ and $\eta = 0.8$. The input ratios are $1/4$ and 4, or $\mp 6.02\,\text{dB}$. The detected ratios are 0.4 and 3.4, giving $-3.98\,\text{dB}$ and $+5.31\,\text{dB}$. The output covariance determinant is $(0.4/2)(3.4/2) = 0.34$, above $1/4$. For this Gaussian state the purity is $1/(2\sqrt{\det V}) \simeq 0.857$. The loss has made it mixed.

From the course to the instrument

Ajku can display covariance ellipses, spectral matrices and phase dependence, but it does not establish η, an optical vacuum variance or the LO angle for you. Those quantities come from the procedures in Part III. Equalizing I/Q variances by an arbitrary numerical scaling would erase the very squeezing/gain distinction that the experiment must resolve.

10 Non-classical states and phase-space representations

10.1 Several meanings of "non-classical"

In quantum optics a state is conventionally classical if it can be expressed as a positive probability mixture of coherent states,

$$\rho = \int \mathrm{d}^2\alpha\, P(\alpha)|\alpha\rangle\langle\alpha|, \qquad P(\alpha) \geq 0, \tag{36}$$

allowing positive delta measures. This is the Glauber–Sudarshan representation [4, 5]. Some quantum states require a nonpositive or more singular distribution and are non-classical in this sense.

This terminology does not mean that coherent light lacks quantum fluctuations: its quadrature vacuum noise and photon shot noise are still present. A "quantum contribution to the detector noise" and a "non-classical optical state" are different claims. Neither is established merely by seeing random ADC samples.

Several experimental witnesses are useful, but are not equivalent:

- A calibrated quadrature variance below vacuum witnesses squeezing. Squeezed vacuum is Gaussian and non-classical.
- Sub-Poissonian photon statistics or antibunching can witness non-classicality in suitable photon-counting measurements. These are not raw ADC variance tests.
- Negative Wigner regions witness non-classicality, but Wigner positivity does not imply classicality: squeezed vacuum has a positive Wigner function.
- Non-Gaussianity alone is insufficient. A classical mixture of two coherent states is non-Gaussian; a single coherent superposition of those states is different.

For a number state $|n\rangle$, both quadrature variances are $n + 1/2$. Thus the non-classical single-photon state has variances $3/2$, above vacuum, and is not quadrature-squeezed. Its ideal X probability is

$$p_1(x) = \frac{2x^2}{\sqrt{\pi}} \mathrm{e}^{-x^2}. \tag{37}$$

This distribution is non-Gaussian and has a node at zero. Loss, finite resolution, mode mismatch and added detector noise change that observed distribution.

An even/odd coherent superposition (a "cat state") is

$$|\mathcal{C}_\pm\rangle = \frac{|\alpha\rangle \pm |-\alpha\rangle}{\sqrt{2(1 \pm \mathrm{e}^{-2|\alpha|^2})}}. \tag{38}$$

Its interference structure distinguishes it from the incoherent mixture $\frac{1}{2}(|\alpha\rangle\langle\alpha| + |-\alpha\rangle\langle-\alpha|)$. Merely seeing two clouds in an electrical histogram does not establish that coherence.

10.2 Wigner representation

With the convention of Eq. (18), define

$$W(x,p) = \frac{1}{2\pi} \int \mathrm{d}y\, \mathrm{e}^{\mathrm{i}py} \langle x - y/2|\rho|x + y/2\rangle. \tag{39}$$

It is real and normalized, $\int W\, \mathrm{d}x\mathrm{d}p = 1$, but need not be nonnegative. Its marginals are actual quadrature probabilities: $\int W(x,p)\mathrm{d}p = p_X(x)$ and $\int W(x,p)\mathrm{d}x = p_P(p)$. Rotating the coordinates gives the marginal for X_θ. It represents *symmetrically ordered* moments; for example $\int xpW\, \mathrm{d}x\mathrm{d}p = \langle XP + PX\rangle/2$.

For a Gaussian state,

$$W(\boldsymbol{r}) = \frac{1}{2\pi\sqrt{\det \boldsymbol{V}}} \exp\Big[-\tfrac{1}{2}(\boldsymbol{r} - \boldsymbol{d})^T \boldsymbol{V}^{-1}(\boldsymbol{r} - \boldsymbol{d})\Big], \quad \boldsymbol{r} = (x,p)^T. \tag{40}$$

In particular $W_0 = \pi^{-1}\mathrm{e}^{-(x^2+p^2)}$ and a coherent state translates this Gaussian to $\boldsymbol{d}$. Displacement changes $W(\boldsymbol{r})$ to $W(\boldsymbol{r} - \boldsymbol{d})$; squeezing stretches one axis and contracts the other while preserving area in a pure lossless transformation.

For number states, with L_n a Laguerre polynomial,

$$W_n(x,p) = \frac{(-1)^n}{\pi} L_n(2(x^2 + p^2)) \mathrm{e}^{-(x^2+p^2)}. \tag{41}$$

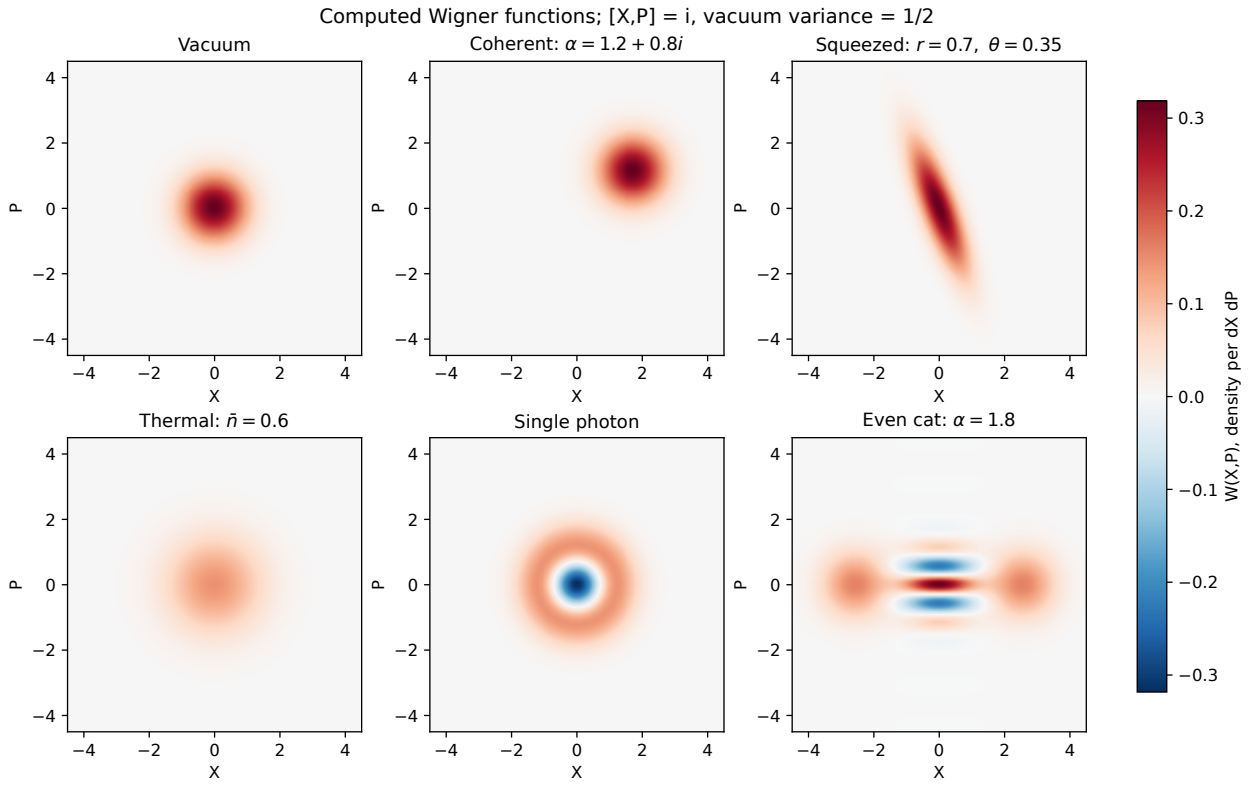


Figure 3: Computed.Wigner functions in the convention $[X, P] = \mathrm{i}$ and vacuum variance $1/2$. Coherent displacement translates the vacuum; squeezing produces an ellipse; thermal noise broadens it. The number state and coherent cat have negative regions. These are analytic quantum-state illustrations, not ajku histograms or reconstructed optical states.

The single-photon value at the origin is $-1/\pi$. More generally $W(0,0) = \langle(-1)^{\hat{n}}\rangle/\pi$: the origin probes photon-number parity. This is why a negative Wigner function cannot be an ordinary simultaneous X,P probability.

For a single-mode *Gaussian* state, purity is $\mathrm{Tr}\,\rho^2 = 1/(2\sqrt{\det V})$. For arbitrary states, purity is $2\pi \int W^2 \mathrm{d}x\mathrm{d}p$, and second moments alone are insufficient. The pure single-photon state has $V = 3\mathbf{1}/2$; incorrectly using the Gaussian formula would give purity $1/3$. This is a useful warning against inferring an entire state from a covariance matrix without checking assumptions.

10.3 Husimi Q and the ordering of moments

The Husimi distribution is

$$Q(\alpha) = \frac{1}{\pi}\langle\alpha|\rho|\alpha\rangle, \qquad \int Q(\alpha)\,\mathrm{d}^2\alpha = 1. \tag{42}$$

It is always nonnegative. For a coherent state $|\beta\rangle$, $Q(\alpha) = \pi^{-1}\mathrm{e}^{-|\alpha-\beta|^2}$; for $|1\rangle$, $Q(\alpha) = \pi^{-1}|\alpha|^2\mathrm{e}^{-|\alpha|^2}$. The Q function is a vacuum-smoothed Wigner representation and describes an ideal heterodyne measurement after the appropriate normalization (Section 11.2).

Since $\alpha = (x+\mathrm{i}p)/\sqrt{2}$, $\mathrm{d}^2\alpha = \mathrm{d}x\mathrm{d}p/2$. Thus the probability density in signal-referred real outcome coordinates is $p_{\mathrm{het}}(x,p) = Q((x+\mathrm{i}p)/\sqrt{2})/2$. Omitting this Jacobian confuses both normalization and vacuum variance.

Representation	Role	Example of ordering
Glauber–Sudarshan P	Coherent-state decomposition; may be nonpositive/singular	$\int \lvert\alpha\rvert^2 P \mathrm{d}^2\alpha = \langle a^\dagger a\rangle$
Wigner W	Quadrature marginals; may be negative	$\int (x^2+p^2) W \mathrm{d}x\mathrm{d}p/2 = \langle n\rangle + 1/2$
Husimi Q	Ideal heterodyne outcome density; nonnegative	$\int \lvert\alpha\rvert^2 Q \mathrm{d}^2\alpha = \langle aa^\dagger\rangle = \langle n\rangle + 1$

The different vacuum terms are physical measurement/ordering differences, not inconsistent algebra. Ajku records voltages and does not choose a quantum ordering for the user simply because a plot is called a joint histogram.

10.4 Two-mode correlations and entanglement

For two modes A and B, operators of different modes commute. A two-mode squeezed state can have $X_A - X_B$ and $P_A + P_B$ simultaneously quiet, while each individual mode has broad thermal-looking marginals. In our convention the equal-weight separability bound is

$$\mathrm{Var}(X_A - X_B) + \mathrm{Var}(P_A + P_B) \geq 2 \quad \text{for separable states.} \tag{43}$$

This is the equal-weight form of a continuous-variable inseparability criterion [6]. Ideal two-mode squeezing gives $2\mathrm{e}^{-2r}$. Unequal loss/gain can motivate a properly normalized weighted criterion; the calibration and the corresponding bound must change together. A positive Gaussian Wigner function can therefore describe an entangled state.

Two electrical channels could measure two homodyne receivers, one for each optical mode. They then provide two selected quadratures at a time, not all four X_A, P_A, X_B, P_B automatically. Both measurement settings and relative LO phases are needed for an entanglement test. Ajku's cross-spectrum/covariance is useful input to that experiment, but electrical correlation alone is not entanglement.

11 From optical operators to electrical I/Q data

11.1 Balanced homodyne detection

Combine a signal mode $\hat{a}$ with a strong LO mode $\hat{b}$ on a balanced beam splitter. Choose output modes $(\hat{a}+\hat{b})/\sqrt{2}$ and $(\hat{b}-\hat{a})/\sqrt{2}$. Their photon-number difference is

$$\hat{n}_1 - \hat{n}_2 = \hat{a}^\dagger\hat{b} + \hat{b}^\dagger\hat{a}. \tag{44}$$

For a strong coherent LO, replace $\hat{b}$ by $\beta = |\beta|\mathrm{e}^{\mathrm{i}\theta}$ in the leading signal term:

$$\hat{n}_1 - \hat{n}_2 \simeq \sqrt{2}|\beta|\hat{X}_\theta. \tag{45}$$

Quantum efficiency, responsivity, transimpedance, bandwidth, loading and attenuation convert this to an electrical voltage. A single calibrated mode readout can be written $v = v_0 + gx_\theta + e$, where e is electronic noise. The gain g grows approximately as $\sqrt{P_{\mathrm{LO}}}$ within the linear detector regime. Thus a fixed quadrature variance produces a voltage variance proportional to LO power.

Balancing cancels the large common LO term, not all technical noise. Detector/gain imbalance converts LO intensity noise into the difference output; spatial mode mismatch reduces signal interference efficiency; unequal paths introduce phase errors. A receiver can compress before the ADC clips. The scale and validity of Eq. (45) must therefore be checked experimentally.

For independent Poisson photoelectrons the one-sided shot-current density of the difference is

$$S_{i,\mathrm{shot}} = 2e(I_1 + I_2) \quad [\mathrm{A}^2/\mathrm{Hz}], \tag{46}$$

where I_1, I_2 are mean diode currents and e is the elementary charge. Multiply by the squared magnitude of the calibrated transimpedance and electrical transfer function for the corresponding voltage density. The factor two is a one-sided spectral convention. DC monitor outputs, if used to determine these currents, need their own calibrated conversion; the balanced RF voltage alone does not determine the individual currents.

A stationary coherent displacement can appear at DC in the chosen LO frame. AC coupling may remove it, and Welch segment detrending removes the segment mean from spectra. Inspect retained means/traces and the receiver coupling rather than equating a missing DC spectral peak with zero optical displacement. RF sideband noise and a single-mode coherent amplitude are related through the chosen mode and detection model, not just the axis labels on an FFT.

From the course to the instrument

The LO polynomial in Eq. (15) is motivated by a constant electronic floor, an approximately linear shot-noise term and possible quadratic technical noise. It remains phenomenological. A classical contribution can also vary linearly over a finite range, and gain drift or compression can distort the fit. Independent power/current measurements and controlled references establish the interpretation.

11.2 Dual homodyne and heterodyne: the added vacuum

Suppose the same signal is split with an independent vacuum mode $\hat{v}$, with $\hat{c} = (\hat{a} + \hat{v})/\sqrt{2}$ and $\hat{d} = (\hat{a} - \hat{v})/\sqrt{2}$. Measure X on c and P on d. After rescaling to the incident signal amplitude, the outcomes correspond to

$$\hat{X}_m = \sqrt{2}\hat{X}_c = \hat{X}_a + \hat{X}_v, \qquad \hat{P}_m = \sqrt{2}\hat{P}_d = \hat{P}_a - \hat{P}_v. \tag{47}$$

They commute: $[\hat{X}_m, \hat{P}_m] = \mathrm{i} - \mathrm{i} = 0$. The two accessible outputs can therefore have an ordinary positive joint distribution without simultaneously measuring noiseless conjugate quadratures of the input. The unused input port's vacuum supplies the extra noise:

$$\boldsymbol{V}_m = \boldsymbol{V}_a + \tfrac{1}{2}\mathbf{1}. \tag{48}$$

For vacuum/coherent input, each signal-referred outcome variance is 1, twice the homodyne variance 1/2. This is the ideal 3 dB simultaneous-measurement penalty in this normalization. Raw splitter-output units would instead have the extra factor $1/\sqrt{2}$ in each amplitude; state which convention you use.

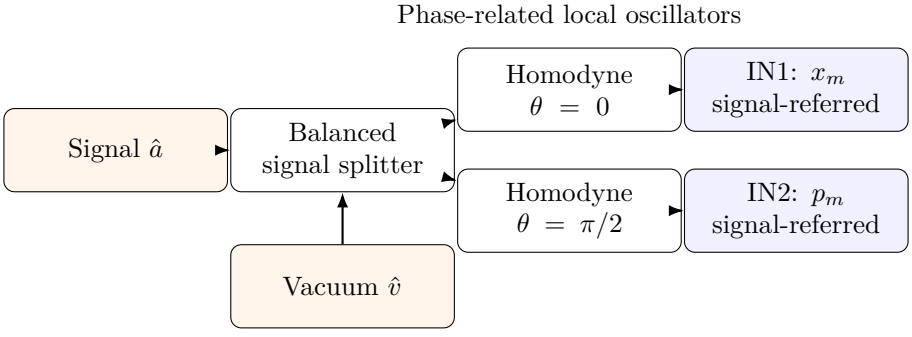


Figure 4: Idealized dual-homodyne topology for Eq. (47). The two LO phases and all optical/electrical gains belong to the external experiment. The vacuum input explains the simultaneous-measurement noise; the ADC labels alone do not establish this topology.

An ideal balanced dual-homodyne/heterodyne measurement realizes

$$\Pi(\alpha)\mathrm{d}^2\alpha = \frac{|\alpha\rangle\langle\alpha|}{\pi}\mathrm{d}^2\alpha, \qquad p(\alpha) = Q(\alpha), \qquad \alpha = (x_m + \mathrm{i}p_m)/\sqrt{2}. \tag{49}$$

The observed Q distribution is a vacuum convolution of the Wigner function. In real signal-referred coordinates,

$$p_{\mathrm{het}}(\boldsymbol{r}) = \int \mathrm{d}^2\boldsymbol{r}'\, W(\boldsymbol{r}')\, \frac{\mathrm{e}^{-|\boldsymbol{r}-\boldsymbol{r}'|^2}}{\pi}. \tag{50}$$

This explains Eq. (48) and why a nonnegative histogram can be entirely consistent with a negative underlying Wigner function.

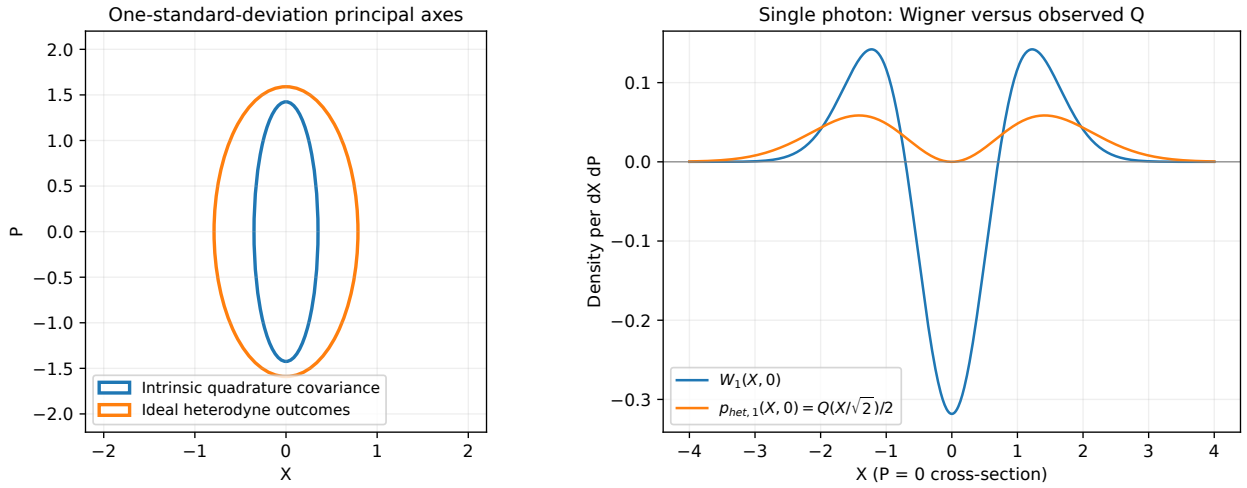


Figure 5: Computed.Left: a squeezed state's intrinsic covariance ellipse and the broader ideal heterodyne covariance after adding $\mathbf{1}/2$. Ellipses show one standard deviation along principal axes, not a 68% joint-confidence region. Right: the single-photon Wigner cross-section is negative at the origin, while its heterodyne outcome density is nonnegative. Their common axes are signal-referred dimensionless quadratures.

For pure squeezed input the ideal heterodyne variance relative to its coherent reference is $(1 + \mathrm{e}^{-2r})/2$ on the squeezed axis. It tends to $1/2$, not zero, as squeezing increases. Removing the known vacuum contribution is a model-based inference with calibration/error propagation, not a way to remove random vacuum noise from individual measured samples.

Optical heterodyne at a frequency-offset LO is closely related, with an image-band vacuum contribution under the usual ideal assumptions. A real CW receiver may involve upper/lower sideband modes, electronic mixers and filters. Establish its mode/phase convention before applying the simple one-mode formulas. Conversely, two independent homodyne receivers can measure two different optical modes; their two ADC voltages are not then a dual-homodyne measurement of one mode.

11.3 From volts to shot-noise units

Consider one fixed homodyne mode/band with $v - v_0 = gx + e$, and assume the electronic noise e is independent and unchanged between measurements. Let V_e be its variance, V_{vac} the LO-on vacuum-reference voltage variance, and V_s the voltage variance in the signal condition. For the convention of this manual,

$$g^2 = 2(V_{\mathrm{vac}} - V_e), \qquad R = \frac{V_s - V_e}{V_{\mathrm{vac}} - V_e}, \qquad \mathrm{Var}(X_{\mathrm{detected}}) = R/2. \tag{51}$$

The denominator must be positive and resolved relative to its uncertainty. This calibrates the detector-output optical state, including prior optical loss. With an independently known efficiency, an input-referred estimate would be $\mathrm{Var}(X_{\mathrm{input}}) = [R - (1 - \eta)]/(2\eta)$ under Eq. (33); it should be reported separately from the detected value.

For ideal signal-referred heterodyne outcomes, the coherent/vacuum variance is 1, so the same voltage model instead gives $g^2 = V_{\mathrm{vac}} - V_e$ and outcome variance R. Inferring the underlying detected signal's single-quadrature variance then subtracts the simultaneous-measurement vacuum

contribution: $\mathrm{Var}(X) = R - 1/2$. The theoretical meaning of the scale is therefore different even if a plotting program labels both references "1 SNU". No subtraction yields an individual noise-free X sample; these are statements about ensemble moments.

Illustrative voltage-variance comparison

Suppose matched homodyne records give $V_e = 0.10$, $V_{\mathrm{vac}} = 1.10$ and $V_s = 0.50$, all in the same arbitrary voltage-squared unit. The electronics-subtracted ratio is $R = 0.40$, or $-3.98\,\mathrm{dB}$, and the detected quadrature variance is 0.20. The unsubtracted ratio is $0.50/1.10 \simeq 0.455$, or $-3.42\,\mathrm{dB}$. Both can be reported, with different definitions and propagated errors. Using the LO-off floor 0.10 as the denominator instead would give a clearance of $10\log_{10}(0.50/0.10) \simeq 6.99\,\mathrm{dB}$, which is not a squeezing ratio.

Use the same temporal/electrical band, LO power, gain and physical chain for these comparisons. Shared references and uncertain electronic subtraction correlate the errors. A changing or correlated electronics contribution violates the simple model. Ajku's ordinary reference display provides the specified voltage ratio and signed excess, not an automatic canonical-quadrature calibration: choose and record the physical reference correctly, and perform any additional inference explicitly.

11.4 Why a joint ADC histogram is not a Wigner reconstruction

Every stored histogram is a nonnegative count distribution of commuting electrical outcomes. To reconstruct a Wigner function from homodyne data, one collects repeated quadrature distributions at known LO angles and uses their relation to rotated Wigner marginals. Mathematically this is a Radon-transform inversion; practical tomography uses finite-sample estimation, often likelihood/Bayesian methods with a physical density-matrix constraint [7–9]. Finite efficiency, phase jitter, digitization and the chosen optical mode enter the likelihood or forward model.

Heterodyne data also permit state inference, but the Q-to-W inversion amplifies noise and is not an innocuous image sharpening. Neither conversion is performed by the current ajku histogram or Gaussian-residual view. The archived literature on dual homodyne and temporal modes is useful background [10]; it is not evidence that these reconstruction algorithms are implemented here.

11.5 Continuous waves, bandwidth and temporal modes

For a travelling field in a rotating frame, write $[\hat{a}(t), \hat{a}^\dagger(t')] = \delta(t - t')$. A normalized temporal mode is

$$\hat{a}_f = \int \mathrm{d}t\, f(t)^* \hat{a}(t), \qquad \int \mathrm{d}t\, |f(t)|^2 = 1, \tag{52}$$

so $[\hat{a}_f, \hat{a}_f^\dagger] = 1$. Orthogonal mode functions give distinct canonical degrees of freedom, commuting between modes; their joint state may still correlate or entangle them. Overlapping functions need not commute in this way. Optical LO shape, detector impulse response and subsequent electrical weighting determine which combination is accessible.

An RF component of a CW homodyne current generally involves *two* optical sidebands. For normalized modes $\hat{a}_+$ and $\hat{a}_-$ around $\omega_0 + \Omega$ and $\omega_0 - \Omega$, its ideal complex amplitude is proportional to

$$\hat{j}_\theta(\Omega) \propto \mathrm{e}^{-\mathrm{i}\theta}\hat{a}_+ + \mathrm{e}^{\mathrm{i}\theta}\hat{a}_-^\dagger. \tag{53}$$

At $\theta = 0$ the cosine and sine electronic components access combinations $X_+ + X_-$ and $P_+ - P_-$. They commute, since the two commutators cancel, just as in Eq. (47). Equivalently they can be viewed as quadratures of two orthogonal sideband-superposition modes. They are not automatically X and P of one optical sideband. An unexcited image sideband supplies vacuum noise; populated or correlated sidebands require their own joint model. This is why "electronic I/Q at an RF frequency" needs a mode definition before a one-mode phase-space interpretation is valid.

An ADC sample is a filtered electrical observation, not automatically a newly prepared independent optical mode. For example, a narrow-band receiver sampled at 125 MSa/s can produce strongly correlated adjacent samples. Squared amplitudes can retain dependence even when an

amplitude correlation is small. This is one reason ajku provides both kinds of correlations and several batch-error scales. Counting raw samples is not a justified count of independent quantum trials.

In a stationary electrical model, filtering a voltage by h gives

$$\mathrm{Var}(v_h) = \int_0^{f_s/2} |H(f)|^2 S_{vv}(f)\mathrm{d}f \tag{54}$$

for a one-sided sampled density, with appropriate discrete-bin approximation in the actual analysis. The raw FPGA variance, a 0–25 MHz spectral integral and a FIR time variance therefore answer different questions. A spectral noise level at one RF frequency does not itself give the photon number of a normalized temporal mode. Mode normalization, detection gain and bandwidth are essential.

The optional FIR supplies a known electrical response and exact sample coordinates; it does not choose an independently calibrated optical temporal mode or implement general mode reconstruction. A CW experiment that requires such reconstruction must specify and validate it separately, retaining the raw records and filter identities needed for reproducibility.

11.6 Higher moments and photon statistics

For ideal, single-mode, calibrated heterodyne outcomes define $M_2 = \mathbb{E}|\alpha|^2$ and $M_4 = \mathbb{E}|\alpha|^4$. Anti-normal ordering gives

$$M_2 = \langle n \rangle + 1, \qquad M_4 = \langle (n+1)(n+2) \rangle. \tag{55}$$

Consequently

$$\langle n(n-1) \rangle = M_4 - 4M_2 + 2, \qquad g^{(2)} = \frac{M_4 - 4M_2 + 2}{(M_2 - 1)^2}, \quad \langle n \rangle > 0. \tag{56}$$

The vacuum corrections are indispensable. Coherent, thermal and number states have $g^{(2)} = 1$, 2, and $1 - 1/n$, respectively, in their ideal single-mode cases. For squeezed vacuum with nonzero mean population it is $3 + 1/\bar{n}$.

Ajku's centred voltage powers/cumulants are not automatically M_2, M_4 in this equation: displacement, normalization, receiver noise, efficiency and optical mode must all be dealt with. Near vacuum, $M_2 - 1$ is small, and subtraction/division can be very ill-conditioned. The software deliberately does not label an electrical power correlation as photon $g^{(2)}$ or a cumulant as an entropy estimate.

A practical interpretation ladder

First establish a valid electrical record and its band. Next establish stable, matched detector references and channel calibration. Then state the optical measurement model and mode normalization. Only then compare inferred quadrature or photon statistics with the quantum-state predictions. The retained provenance and explicit unavailable values let the interpretation improve without changing what the ADC actually measured.

12 Worked connections and student exercises

These exercises use ideal models or the documented digital conventions. They are not new measurements or claims about the installed optical apparatus.

12.1 A coherent state in two measurement schemes

Take $\alpha = 1 + \mathrm{i}/2$. A homodyne measurement at angle zero has mean $\sqrt{2}$ and variance $1/2$; at $\pi/2$ its mean is $1/\sqrt{2}$ with the same variance. An ideal dual-homodyne measurement in signal-referred coordinates has the same two means, but each variance is 1. Explain why this does not contradict uncertainty.

Answer. The commuting measured observables contain an independent vacuum mode, Eq. (47). The joint distribution is Q, not a noiseless joint probability of the original X and P. The extra half-unit variance in each channel is required by the simultaneous-measurement arrangement.

12.2 Squeezing, loss and phase

Use $r = \ln 2$ and $\eta = 0.8$. Compute the observed variance ratios at an angular error of 10° from the squeezed axis.

Answer. The principal-axis ratios are 0.4 and 3.4. Mixing gives $R = 0.4\cos^2(10^\circ) + 3.4\sin^2(10^\circ) \simeq 0.4905$, or about $-3.09\,\mathrm{dB}$. The measurement is nearly 0.9 dB worse than the correctly aligned case, despite identical optical loss. In an experiment, repeat the vacuum reference at matched LO power and electronics settings before attributing the difference.

12.3 Classical mixtures versus cats

For real α, compare $|\mathcal{C}_+\rangle$ with the equal incoherent mixture of $|\alpha\rangle$ and $|-\alpha\rangle$. Both can have two lobes along X. Which evidence distinguishes them?

Answer. The superposition has interference terms between the branches, visible in its phase-space structure and appropriate quadrature distributions; the incoherent mixture does not. Two peaks or nonzero fourth cumulants alone are insufficient. Known measurement phases, detector response and a validated inference procedure are needed. Figure 3 shows a computed superposition, not an automatic interpretation of a two-lobed measured histogram.

12.4 A negative fourth cumulant from a classical state

Randomize the phase of a coherent state of fixed magnitude $|\alpha| = r$. This is a positive mixture of coherent states and hence classical in the optical sense. Its X outcome can be represented as $X = \sqrt{2}r\cos\Theta + \epsilon$, where Θ is uniform and ϵ is independent vacuum Gaussian noise of variance $1/2$. The mean is zero and the variance is $r^2 + 1/2$. Using $\mathbb{E}\cos^2\Theta = 1/2$ and $\mathbb{E}\cos^4\Theta = 3/8$ gives

$$\kappa_4(X) = -\tfrac{3}{2}r^4. \tag{57}$$

The independent Gaussian contribution has zero fourth cumulant. Thus even a large, negative, well-resolved fourth cumulant need not be a non-classicality witness. It distinguishes this phase-randomized coherent mixture from a thermal Gaussian with the same covariance, but does not by itself establish a quantum advantage.

12.5 ADC resolution and a quiet signal

With nominal LV calibration $K = 32768$ stream codes/V, one native code corresponds to $4/K \simeq 122.07\,\mu\mathrm{V}$. An RMS signal occupying only a few codes is not improved by appending two zero bits to the recording. Nor does a large number of samples create a finer instantaneous ADC code. Averaged estimates can become more precise if noise/dither and dependence are correctly modelled; this is a statement about an estimator, not extra converter resolution.

12.6 Bandwidth and the apparent noise reduction

Suppose raw samples contain independent unit-variance white noise. The default FIR has $\sum_k h[k]^2 \simeq 0.35247$, so its time-domain variance is about 0.35247. A raw one-sided spectral integral through 25 MHz covers $25/62.5 = 0.4$ of the Nyquist band, giving about 0.4. Explain the difference.

Answer. They are different response weightings. The default FIR passband ends at 20 MHz and has a transition to 25 MHz; it is not a 25 MHz brick wall. This reduction is digital filtering of white noise, not evidence of optical squeezing. A squeezed-state claim compares matched detection/filtering conditions with a calibrated optical vacuum reference.

12.7 Coverage and events

An 8 MiB capture contains 16.777 216 ms. If one such capture is retained every two seconds, coverage is about 0.839%. Exact streamed statistics can still cover nearly all clock slots when their queues are drained. Why can a transient be visible in statistics but absent from the next raw record?

Answer. Statistics preserve aggregates, not waveform history. A host-triggered capture begins only after the event's block has completed and the command has been processed. Its recorded delay can exceed the transient duration. Inspect event and capture intervals; do not call the follow-up a pretrigger waveform.

12.8 No observed tail events

An 8 MiB record contains 2^{21} paired samples. For an ideal standard Gaussian, the two-sided probability beyond 6σ is about 1.97×10^{-9}. Even if all those samples were independent, the expected count would be only about 0.0041 per channel. Is an observed count of zero evidence that the physical distribution has no such tail?

Answer. No. Zero is overwhelmingly likely for that record length under the Gaussian model itself. Temporal dependence further changes the sampling error. Ajku retains zero counts but withholds an unsupported tail SE; it does not infer a zero probability or an IID confidence bound. Fitted thresholds introduce their own uncertainty, and the Gaussian tail model must be assessed rather than assumed. Likewise, vanishing third/fourth cumulants are compatible with both classical Gaussian states and non-classical squeezed Gaussian states.

12.9 Design a defensible observation

Prepare a short lab plan to distinguish a receiver gain imbalance from a squeezed quadrature. Identify the required electrical gain/delay check, optical vacuum reference, phase reference, loss estimate, matched analysis band and repeatability test. State what remains unknown if only one circular-looking I/Q cloud is available. This exercise has no software-only shortcut: the result depends on the actual detector topology and calibration of the user's apparatus.

Part III
Characterization, qualification and reference

13 Characterizing the user's board and optical chain

This chapter describes work for the end user. The development qualification did not perform these analogue/optical measurements or create transferable calibration constants. Procedures, equations and file conventions transfer; board gains, receiver response, optical efficiency and noise baselines must be measured for the actual apparatus. The existing common board calibration was retained during digital qualification.

13.1 Separate three calibration layers

1. **Board conversion:** native ADC codes to volts at each connector, for the actual board, physical LV/HV jumper and loading. This is the common board calibration, shared by every design on that board.
2. **Electrical chain:** receiver/transimpedance gain, cables, splitters, terminations, attenuation, bandwidth, anti-alias rejection and relative delay.
3. **Optical measurement:** detector topology, LO power/phase, signal/LO mode overlap, efficiency, optical mode definition, vacuum normalization and the state-preparation conditions.

Keep these identities distinct. A fresh board calibration does not establish an optical quadrature angle. A circular vacuum cloud does not establish equal optical gains. A receiver calibration made with one load/cable/optical alignment need not survive a change to those components.

Before each campaign record board identity, software/build IDs, range jumpers, receiver model/serial and gain, wavelength and LO arrangement, optical and electrical filter IDs, cable/termination details, warm-up/temperature conditions, meter IDs, dates and uncertainty statements. Include file hashes/identifiers for external calibration tables in `measurement_chain`; retain the tables themselves. Record observations separately from requested control settings.

13.2 Board offset and DC gain

Use the common procedure in `common/CALIBRATION.md`; ajku reports calibration capability through `status()["calibrates"]`. Board profiles live in `/home/xilinx/rpfpga/calibration.json` and jumper declarations in `board.json`. Preserve a dated backup before replacing them. Do not import the development board's constants into another board.

1. Allow thermal settling and record the warm-up conditions. Set the physical jumpers, then declare the same ranges in software. Keep all test voltages within the range and equipment specifications.
2. For input offsets, disconnect signal sources and terminate/short the inputs appropriately. An open input is not a quiet zero reference. Measure and store offsets with `dev.measure_input_offsets()`.
3. If using the common output-reference/loopback procedure, disconnect sensitive loads from OUT1/OUT2. `dev.hold_reference()` deliberately drives both outputs; measure the actual reference levels with a calibrated DMM, then call `dev.release_reference()`. Record both readings and meter uncertainty.
4. With the prescribed LV loopback cables OUT1→IN1 and OUT2→IN2 and no other loads, run `dev.calibrate(dmm=...)` using those measured readings. The common calibration sweeps output codes; it is not a passive measurement. Check the fit residuals and preserve the returned report.

5. Verify at independent test levels, including both polarities and levels away from the fit points. Repeat the offset check after reconnecting the actual chain. Characterize the HV range separately if it is used; a small LV-sized stimulus occupies little of HV full scale and gives a poorer gain estimate.

For example, after obtaining actual DMM values and wiring the loopback:

```
# measured_out1_v and measured_out2_v are your recorded DMM readings.
report = dev.calibrate(dmm={"out1": measured_out1_v,
                            "out2": measured_out2_v})
print(report)
print(dev.calibration())
```

Without a new DMM reference, the loopback retains stored ADC gains and calibrates the DAC relative to them. If those ADC gains were nominal, loopback agreement alone cannot establish absolute voltage accuracy. Successful offset/loopback steps restore or re-apply configuration; an explicit reference hold persists until release or a superseding edit. Configuration changes cancel conflicting measurements.

The DC model is Eq. (1). For independent uncertainties, a first-order example is

$$u_v^2 \simeq (4/K)^2 u_c^2 + (1/K)^2 u_A^2 + (v/K)^2 u_K^2. \tag{58}$$

Include covariance terms when calibration parameters share a reference or fitted data; the independent formula is not universally appropriate. Gain uncertainty propagates approximately twice fractionally to a variance. The software's statistical SE arrays do not contain this entire systematic budget. A quantization variance $\Delta^2/12$ is only a model under suitable quantization-error/dither assumptions, not an automatically independent noise source in every trace.

13.3 Frequency response and channel delay

Use a characterized low-distortion electrical source, suitable attenuation/load, and an independent amplitude reference at the connector. A generator's front-panel voltage may assume $50\,\Omega$; do not confuse that with the voltage into a high-impedance input. To characterize the full chain, include the intended receiver, cables, filters and loading, using an appropriate optical/electrical stimulus.

1. Sweep frequency over the useful band at a fixed small amplitude. Repeat at several amplitudes to distinguish response from compression. Record source readback and measurement uncertainty at each point.
2. Use enough settling and capture length for the lowest frequency. Integrate the tone's window-broadened bins, or use a separately validated sinusoidal fit. Keep detrending and the admitted frequency interval explicit. The delivered analysis code does not silently fit an analogue correction table.
3. For relative channel response, feed both inputs from a characterized splitter. Correct its imbalance or swap paths and combine the measurements; otherwise the result includes splitter/cable differences. Preserve polarity and phase conventions.
4. With high coherence, unwrap $\arg S_{IQ}$ over valid frequencies. A relative delay $\tau = \tau_Q - \tau_I$ gives $\arg S_{IQ} \simeq \phi_0 - 2\pi f\tau$. Fit the slope only over a region where a delay model is adequate; frequency-dependent distortion requires a richer model.
5. Repeat on another day or after rewiring to determine reproducibility. Store the response, delay, validity band, load, interpolation policy and uncertainty.

For linear channels $Y_I = H_I X_I$ and $Y_Q = H_Q X_Q$, our cross convention gives

$$S_{IQ}^{(Y)} = H_I^* H_Q S_{IQ}^{(X)}, \qquad S_{jj}^{(Y)} = |H_j|^2 S_{jj}^{(X)}. \tag{59}$$

For a general complex matrix $\boldsymbol{Y} = \boldsymbol{H}\boldsymbol{X}$, the same convention gives $\boldsymbol{S}_Y = \boldsymbol{H}^* \boldsymbol{S}_X \boldsymbol{H}^T$. Keep the conjugations consistent if writing an external de-embedding program. Inverting a small response gain amplifies noise and calibration errors; report the valid band rather than producing an uncontrolled inverse. Current ajku processing applies the DC gain/offset and optional FIR, not this frequency-response de-embedding.

13.4 Anti-alias rejection

The sampled raw band ends at the 62.5 MHz Nyquist frequency. A broad optical receiver can deliver noise far above it. The contribution folded into an observed frequency is set by the full analogue transfer function and source spectrum, not by the host's 25 MHz display limit.

1. Identify all analogue poles/filters before the ADC, including receiver and board response. Obtain an independent characterization over the out-of-band range that can contribute appreciably, not only the displayed science band.
2. Inject appropriately attenuated known out-of-band tones and look at their predicted aliases. At nominal $f_s = 125$ MHz, a 110 MHz tone folds to 15 MHz. Confirm actual connector amplitude and avoid overload outside the displayed band.
3. Repeat with the intended external filter and compare rejected power. To bound aliased *noise*, combine the measured response with a justified out-of-band source-noise model or measurement; a few tone checks alone do not specify all noise.
4. Document the resulting useful science band and rejection uncertainty.

Filtering on the PC prevents new aliases from digital decimation. It cannot undo unknown aliases already summed at sampling. A quiet-looking filtered display is therefore not evidence of adequate analogue anti-alias rejection.

13.5 Electronic floor and LO-power sweep

Prepare separate, precisely described conditions:

- terminated ADC input, to study the board/load background;
- powered receiver with optical signal and LO blocked, to study the complete electronic chain under its actual gain/load;
- LO on with the signal port in the intended vacuum/reference condition;
- the actual optical signal or state under study.

These are not interchangeable references. In particular **LO off is an electronic-background condition, not a calibrated homodyne vacuum measurement**. The LO is required to convert the chosen optical quadrature fluctuations into the difference photocurrent. Blocking/terminating changes can also alter receiver noise.

For an LO sweep:

1. Establish the topology, wavelength, matched signal condition and receiver operating range. Use an independent calibrated power meter or calibrated monitor path to obtain the relevant LO power. Document taps, split ratios and diode-current conversion if applicable.
2. Acquire the LO-off reference and at least three distinct positive-power levels, with at least two complete captures per level; more repeats/levels are normally useful. Use `iqsweep.py record` or the supplied-callback API. Keep electrical settings and analysis contracts identical.
3. Allow settling, record power before/after, and interleave repeated reference levels to reveal drift. Account separately for reading uncertainty and common meter calibration/scale uncertainty. A requested attenuator value is not a measured power.

4. Inspect native-code extrema, near-rail occupancy, detector DC monitors and gain-versus-power. Stop increasing power if any receiver stage approaches its specified saturation, even if the ADC still has headroom.

5. Fit Eq. (15). Inspect signed coefficients, residuals, repeated level scatter, lack-of-fit diagnostics and stability across frequency/band choices. Report observed clearance separately from a conditional quantum/classical attribution.

The archived heterodyne/homodyne studies illustrate the importance of calibration and detector-response modelling [11, 12]; their bit rates, security conclusions and optical sensitivities do not transfer to this instrument.

13.6 Optical quadrature angle, gain and coherent probe

Vacuum is rotationally invariant, so it does not identify an arbitrary optical phase axis. Introduce an independently characterized coherent probe and vary its relative phase or beat frequency using the actual experiment's optics. Record the phase-setting/readback uncertainty and the LO-power stability.

In a simple linear baseband model,

$$\begin{pmatrix} v_I \\ v_Q \end{pmatrix} = \boldsymbol{b} + \boldsymbol{G} \begin{pmatrix} x \\ p \end{pmatrix} + \boldsymbol{e}. \tag{60}$$

A phase scan of a fixed-amplitude probe traces the columns of $\boldsymbol{G}$ rather than assuming ideal equal orthogonal channels. Use an absolute probe reference if absolute gain is required. Frequency sweeps reveal delay and frequency-dependent imbalance. Separate actual optical imbalance from splitter/cable/ADC effects. Check linearity at multiple probe amplitudes and repeat the fit after changing the alignment.

An affine fit or whitening of an unknown noise cloud cannot determine $\boldsymbol{G}$ uniquely: genuine state anisotropy and detector imbalance are confounded. Keep raw records and the uncorrected matrix. Apply any independently established correction in a separately versioned analysis, with uncertainty and valid operating conditions. The current GUI does not automatically perform this optical calibration.

13.7 A squeezing measurement with a matched reference

1. Establish a homodyne (or explicitly modelled simultaneous) detection scheme, optical mode and phase reference. Determine losses/efficiency as far as required by the intended claim, keeping their uncertainty distinct from electrical gain.

2. Record a vacuum/coherent reference with matched LO power, receiver gain, loading, optical alignment, analysis band and FIR. Measure the electronic floor separately and repeat both references to assess drift.

3. Scan the relevant LO phase. Retain raw/processed records and phase readbacks at each setting. Look for the expected π-periodic variance dependence and the corresponding anti-squeezing. Check whether phase jitter explains degraded minima.

4. Compare band variances or spectra with matched normalization. If electronic noise is subtracted, preserve both unsubtracted and subtracted values and propagate the uncertainty/covariance of the reference. Do not clip negative estimates to zero.

5. Quote the measured squeezing at the detector. Any inferred source squeezing after loss correction is a separate, model-dependent result; inversion becomes sensitive to efficiency errors when losses are large.

For dual homodyne, account for Eq. (48); a naive homodyne normalization would misidentify the added vacuum. State-tomographic or non-Gaussian claims need the additional measurement settings and inference validation of Section 10, beyond fitting a covariance ellipse.

13.8 Uncertainty, selection and a complete lab record

For each reported quantity state its physical band/mode, estimator, sample interval, selection rule, number of independent repetitions and uncertainty components. Inspect stationarity across capture order and the scale dependence of batch errors. Use independent repetitions or an appropriate dependence-aware resampling model when long-term drift dominates within-record errors. Do not combine event-triggered captures with scheduled/background data without acknowledging their selection.

A useful record accompanies the raw files with a wiring/optical diagram, conditions, meter readings and calibration files, the exact processing profile, excluded/failed records with reasons, and the scripts that produce the reported figure. Evaluate drift, gain/phase uncertainties, power-meter errors, optical efficiency, aliasing, finite-record effects and numerical processing separately before combining them. Shared references cause correlated errors; adding every error term independently in quadrature can be wrong.

The supported software already preserves raw data and the processing context needed for this work. It cannot verify that the experimenter actually blocked the intended port, used the stated load or measured the quoted power. That responsibility belongs to the laboratory record of each user's apparatus.

14 Measured digital operating envelope

Qualification exercises the acquisition, arithmetic, transport, file lifecycle and analysis software. It does not determine an optical sensitivity or supply another board's calibration. This chapter describes release 1.0: FPGA/driver 1.0, build `0xfb4e4824`, host processing revision 5 and common framework 1.3, qualified on 2026-11-03 with every addition of Section 1.3. The tested software was commit `a0a2eca` plus the uncommitted 1.0 changes; the hashes of all 31 board-side driver/framework files were checked against it before and after the qualification sequence. The existing board calibration was retained. A first attempt failed the control maximum below; its correction and the earlier 0.2 baseline (commit `60635cb`) are recorded in the verification record.

The counter source supplies a signed ramp/complement pair with period 16,384 pairs. Every captured sample can be checked against its declared index modulo that period; exact statistics and histogram bins have independent integer predictions. This is strong evidence for digital integrity, but a periodic pattern does not uniquely verify an arbitrary absolute-time displacement of whole periods. Additional checks compare actual ADC samples with a wholly overlapping complete FPGA statistics block. They verify all fourteen exact sums through fourth order, not the ADC's analogue accuracy. Counter spectral epochs are compared sum by sum with the bit-exact model; ADC spectra are checked against Parseval's theorem and host Welch spectra (1–25 MHz band means within 0.4 %).

14.1 Qualification results

Quantity	Complete 8 MiB workload	Complete 32 MiB workload
Elapsed recording time (s)	3603.473	139.902
Retained captures	1657	14
Exact statistics blocks	429,551	16,675
Histogram snapshots	8,820	385
Raw coverage (%)	0.771	0.672
Control p99 / max (ms)	141.2 / 334.0	138.9 / 245.1
Capture/download median (s)	0.702	2.383
Analysis median (s)	1.745	9.327
Subscriber loss / invalid slots	0 / 0	0 / 0
FIFO loss in counter phase	0	0

The source preflight matched 31 board-side files. The digital suite passed captures through 32 MiB, a mid-transfer cancellation (shown by the capture's own abort record) with drain and a subsequent

capture, and exact ADC/statistics agreement. The priority repeat reproduced its recorded trigger decisions and windows, with a measured follow-up delay of 141.8 ms. The stalled subscriber reported 6760 lost blocks while the fast reader and FPGA FIFO remained loss-free; completed-transfer disconnect recovery also passed.

During the one-hour counter phase, 106 valid steady-run OS samples were retained after a five-minute warm-up exclusion. Board RSS ranged 72.90–81.18 MiB; its first/last-quarter medians were 73.52/74.70 MiB. Host RSS ranged 161.6–292.9 MiB, with first/last-quarter medians 193.9/167.1 MiB. These sampled ranges include different acquisition/analysis phases; they are not an instantaneous bound or an extrapolation to arbitrary run duration.

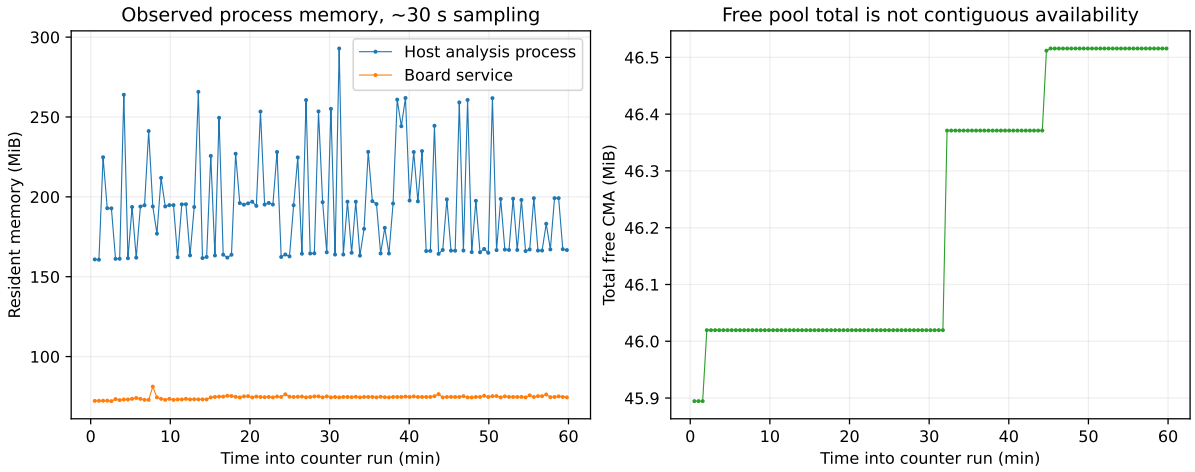


Figure 6: Measured digital qualification.Sampled process RSS and board free contiguous-memory-pool capacity during the qualified one-hour counter workload. Free CMA is a total, not the size of the largest allocatable block. The waveform/statistics files consume host disk space as recording proceeds.

The complete workload includes finite counter captures and raw/FIR processing, 65,536-point Hann spectra, ±128 temporal lags, 32-bin-per-axis captured diagnostics with 16,384-pair batches, statistics and histogram recording, matched references and independent control probing. Default FIR settings are the ones in Section 3.4. The ordinary GUI defaults to 64 diagnostic bins; the qualified workload's 32-bin choice is recorded explicitly rather than presented as universal performance.

Resource observations sample the host process and free disk space, and read-only board OS process/memory information over SSH roughly every 30 seconds. They are finite-rate observations, not a proof of all instantaneous memory states or a worst-case execution-time bound. Control-envelope checks use p99 at most 250 ms and observed maximum at most 500 ms for the stated workload. The common aspirational 100 ms response target is not a measured guarantee. The first 1.0 attempt failed this envelope: one of 18,381 probes took 515.6 ms. Each status merged the board calibration file onto its defaults seven times, about 5 ms each on the board's processor; the common framework now caches the merged result by the file's time stamp, which halved the idle status time (46 ms to 19 ms). The repeat, whose results are tabulated above, passed.

14.2 Capture sizes, retention and throughput

Use 8 MiB as the default operating point. Repeated 32 MiB captures provide a larger qualified finite-record option when allocation succeeds. The 64 MiB API ceiling has repeatedly failed contiguous allocation on the development board; do not budget an experiment on its availability. Total free CMA memory does not establish one sufficiently large contiguous region. A nominal maximum is not an allocator promise.

Throughput is accepted as a measured operating limitation for this baseline. No universal 3–5% raw coverage is claimed, and further optimization is not a prerequisite for using the documented workflows. Inspect actual retained intervals, elapsed time, selection policy and analysis settings.

Default-size raw/FIR/distribution processing has a materially lower duty cycle than a raw-transfer-only test. A streamed summary can cover a waveform gap without retaining the samples in that gap.

Stalled-reader qualification deliberately pauses one subscriber and checks that its loss count grows while a fast subscriber remains correct. Completed-transfer disconnect checks close a binary transfer after the hardware capture has drained and verify that the next acquisition succeeds. Cancellation checks schedule a hardware-timed capture, supersede it through a normal configuration edit once the sample clock has passed its start, confirm from the capture's own record that it stopped mid-transfer, and verify a subsequent capture. These are distinct from terminating the service during active DMA.

14.3 Real notebook checks and a GUI correction

The 1.0 real Notebook/Chromium run recorded two ADC captures with 1,996 statistics blocks and 10 histograms, saved and compared a raw/FIR reference, verified the displayed matrix and distribution PNGs against expected renders and the FPGA spectrum line of the live summary, and ran a monitored recording with 443 blocks, 28 trend windows and one follow-up capture (470.3 ms after its event, zero pretrigger raw pairs), all without loss or invalid slots. The rest of this subsection describes the 0.2 baseline's browser checks, which found the GUI correction.

The first browser run of 0.2 exposed an unbounded UI callback backlog: repeated summary reads could queue plots faster than the notebook rendered them, delaying controls and cleanup. The shared `ManagedPoller` now allows one outstanding publication; a slow renderer backpressures polling, while the independent statistics recorder continues normally. Close releases a worker even if a queued callback is never delivered. Three regression tests cover this bound, cancellation/replacement, disposal and publisher failure. After the correction, 105 ajku host tests and 323 common tests passed (20 common skips in the notebook-enabled environment).

The final 0.2 Notebook/Chromium repeat recorded two ADC captures with 2,006 statistics blocks, saved/compared a raw/FIR reference, and verified the actual displayed matrix and distribution PNGs against expected renders. A separate monitored run retained 494 statistics blocks, 31 trend windows and one follow-up capture. Both runs had zero FIFO/subscriber loss and zero invalid slots. The latter delay was 305.1 ms, with zero pretrigger raw pairs. Explicit synthetic fitting and normal panel cleanup also passed. These short UI runs are functional checks, not throughput benchmarks. A notebook-frontend startup `schema` exception was logged separately; it did not prevent the exercised workflows. Earlier unsuccessful automation attempts and their explanations are retained in the verification record.

The starter notebook's cleanup cell is commented so **Run All** opens the dashboard without immediately closing it. The actual notebook was also executed through a Jupyter kernel. The GUI correction is recorded by source hash in the 0.2 browser receipt.

14.4 Supported use and remaining boundaries

The measured envelope supports controlled finite acquisitions and bounded concurrent statistics/histogram readout under the documented software lifecycle. It does not establish arbitrary-duration unattended service, power-loss durability or active-DMA process-crash recovery. A longer run gives stronger finite evidence; it does not turn finite queues into a lossless stream or prove all failure recovery.

The user's own electrical and optical characterization remains necessary for physical interpretation and is described in Section 13. The manual's computed quantum-state figures are not optical validation data. Exact dated receipts and any failed/incomplete checks are retained in `docs/data/`; `docs/VERIFICATION.md` explains their scope.

A Interface and file-format reference

This manual describes release 1.0: FPGA/driver 1.0, build `0xfb4e4824`, register ABI `0x00040000`, host processing revision 5 and common software 1.3, qualified on the board on 2026-11-03 (`docs/data/board-1.0-*`). The supported interface is frozen in `docs/RELEASE.md`. Processing revisions keep explicit identity rather than silently changing old results: revision-4 references are reanalysed before comparison. Low-level register details are in `docs/REGISTERS.md`; normal users should use the checked driver/client. Reading registers of an unconfigured FPGA can hang the board, which is why every hardware access goes through the common Overlay checks.

A.1 Python entry points

Entry point	Meaning
`CLIENT(host, port=5025)`	PC connection to the existing service; use as a context manager
`dev.snapshot()`	Atomic source, session/revision, status, ranges and calibration readback
`dev.apply(patch)`	Accepted configuration edit: source (ADC or counter) and the histogram window in volts
`dev.apply_snapshot(..., expected_revision=r)`	Revision-guarded edit and its readback
`dev.safe()`	Select ADC/calibrated-zero and abort an active capture; accepted writes still drain
`dev.statistics()`	Latest complete raw block and contextual conversion
`dev.statistics_stream()`	Owned stream; use `read()` to receive complete packets
`dev.histogram()`	Claim/read/release a complete raw histogram snapshot
`dev.spectrum()`	Claim/read/release the latest continuous spectral epoch: exact sums and densities
`dev.capture(pairs, expected_revision=r)`	PC tuple (metadata, immutable native bytes); guard is optional
`dev.capture(pairs, start_index=n)`	Hardware-timed capture from sample index n (`dev.sample_index()` gives the board's clock)
`iqmeasure.analyze(meta, raw, ...)`	Host measurement products for one contiguous record
`LiveSession.configure(...)`	Validated/pinned host settings, condition and chain metadata
`start_recording(path, monitoring=...)`	New run plus complete retained statistics; optional drift tracker
`stop_recording()`	Stop admission; returns run/recorder, whose `join()` waits for owned work
`CapturePipeline(session, ...)`	At most two admitted capture/download/analysis jobs
`load_measurement(path)`	Validate/open NPZ products without pickle
`load_reference(path)`	Validate reference identities/files and return measurement/manifest
`reconstruct_filtered(measurement, raw)`	Recreate complete FIR waveform with saved coefficients and timing

The board-side `DRIVER.capture` returns metadata, with a separate `capture_data` call for bytes. Do not confuse it with the PC convenience method that performs both phases. Clone a client for independent concurrent requests; one TCP connection handles one command at a time.

A.2 Wire commands

The service listens on TCP 5025. Ordinary responses are one line beginning with `OK` or `ERR`. The shared session-aware REQUEST envelope guards Python control calls. The design-specific commands are:

Command	Result
`STATS?`	Latest block, context and optional physical conversion
`STATS DRAIN [maximum]`	Destructive read of up to 32 complete queued blocks
`STATS STREAM`	Dedicated connection: header, then statistics JSON packets
`HIST?`	Complete claimed histogram, geometry and loss/coverage metadata
`SPEC?`	Latest claimed spectral epoch: exact sums, densities, frames and losses
`CAPTURE [pairs]`	Long acquisition; metadata after completion and copy
`CAPTURE DATA <id>`	Separate connection: metadata, exactly the declared bytes, EOF

`CONFIG?`, `SNAPSHOT?`, `STATUS?`, `SAFE`, presets, range and calibration commands follow the common framework. Configuration accepts `{"source":"adc"}` or `{"source":"counter"}`; unsupported keys are rejected before writes. Direct DRAIN must not compete with the service's active statistics subscriber; there is one destructive FPGA drainer.

Only the latest completed capture remains available at the service. An older ID cannot retrieve a newer capture accidentally. A transfer that already owns immutable bytes can finish while another capture completes. The binary client checks identity, session, pair count, completeness and payload length. Stream generators do not access MMIO. No partial/uncertain acquisition is automatically replayed.

A.3 Raw and processed formats

Format	Contract
`.iqbin`, metadata `ajku-iq-v1`	Four-byte `IQ01`, little-endian u32 JSON byte length, UTF-8 JSON, then exact raw payload
Raw payload	Interleaved $I_0, Q_0, I_1, Q_1, \ldots$; little-endian signed int16, right-aligned native codes $-8192 \ldots 8191$; no zero-padding substitution
`ajku-measurement-v1`	NPZ arrays plus JSON document; analysis revision is inside the processing contract, currently 4
`ajku-reference-v1`	Reference directory/manifest with raw and processed hashes and exact comparison context
`ajku-run-v1`	Run index, file hashes, selection/intervals, lifecycle, coverage, errors and stream losses
`ajku-fir-v1`, `ajku-filter-profile-v1`	Validated descriptor, exact coefficient identity and replayable coefficient list
`ajku-measurement-hdf5-v1`	Native codes, processed arrays, full FIR waveform when enabled, JSON context, integer timing and completion marker
`ajku-lo-level-v1`, `ajku-lo-sweep-v1`	Independent power provenance and complete measurement links; failed/incomplete levels are ineligible
`ajku-lo-calibration-v1`	Signed fits, residuals, statistical/meter errors, explicit synthetic/conditional interpretation
`ajku-drift-window-v1`	Exact aggregate raw statistics, context, band, units and sample interval

Native files have a 1 MiB metadata bound and a 64 MiB payload ceiling. The latter is an API limit, not a contiguous-memory availability guarantee. Filenames and output directories must be new; loaders reject truncation/trailing data. Keep NaNs in arrays for undefined normalized quantities distinct from finite JSON metadata.

First indices and FIR centres/strides are integer raw-clock coordinates. Never subtract rounded floating-point absolute timestamps to recover short intervals. Do not reuse a raw capture's sample rate or pair count to describe its FIR branch. Old processed references can be loaded where supported, but their analysis contract must match before comparison; reanalyse retained raw files when migrating revisions.

B Troubleshooting and interpretation checklist

Observation	Interpretation and next action
`BUSY`	Another acquisition owns the slot, or a histogram epoch is not yet ready. Finish the owner/wait for an epoch; this is not an automatic queue.
`CONFLICT`	An expected revision is stale. Refresh, reconsider and apply a new explicit request; do not overwrite another controller's change.
`STALE`	The FPGA session changed. Reconnect and rebuild the panel/context.
`CANCELLED`	A request was superseded or its acquisition deadline expired. Inspect state; accepted DMA writes still have to drain.
Allocation failure	Contiguous DDR allocation could not be satisfied. Use the qualified smaller size and inspect ownership/state; total free memory is not one contiguous region. Do not assume a failed 64 MiB request corrupts smaller records.
Reference rejected	Read the listed processing/source/calibration/range/chain difference. Obtain or explicitly reanalyse a matching reference.
FIFO drop count nonzero	Compare counters within the recorded revision and actual index gaps. Initial backlog can include losses before subscription.
Subscriber loss only	That consumer fell behind. A fast independent subscriber can remain correct; do not infer global hardware loss from one client.
Histogram missed epochs	Reader retention is finite/latest-wins. A missing epoch is not necessarily a gap in the counting engine.
Empty phase/SE values	Significance or required sample/batch counts are unavailable, or a normalization has zero variance. Missing is not zero.
Counter in "measurement"	Diagnostic data have native-code units and contain deterministic full-scale structure. Select ADC for physical measurements.
Small filtered variance	Check actual FIR bandwidth, response and normalization against the reference. Filtering alone is not optical squeezing.
Non-Gaussian tails	Inspect clipping, quantization, drift, interference and selection before assigning a quantum origin. Retain out-of-view counts.
Late raw trigger	Read the event/capture indices and delay. Host follow-up has no pretrigger raw data, so it may miss the event waveform.

C Glossary

ADC slot

One nominal 8 ns sampling opportunity; invalid slots are counted.

Capture
A finite contiguous paired waveform, distinct from the continuous summaries.

Coverage
Union of retained raw sample intervals divided by elapsed recording time.

LO Local oscillator; an optical phase/amplitude reference supplied by the experiment.

Quadrature
A phase-dependent field observable, with normalization and mode specified.

Coherence
In the displayed spectral matrix, magnitude-squared electrical cross-coherence; distinct from the quantum definition of a coherent state.

ENBW
Window equivalent noise bandwidth; distinct from a FIR's integrated white-noise bandwidth and an analogue receiver's effective bandwidth.

Vacuum reference
LO-on detection of the specified unexcited signal mode, under a documented detector model; not simply an open or LO-off ADC input.

Shot-noise unit
A stated reference variance/density normalization; record the detection scheme, gain, LO power, band and electronic-noise treatment.

Wigner / Husimi Q
Different phase-space representations; Wigner can be negative, whereas ideal heterodyne outcomes sample a nonnegative Q distribution.

Revision / session
Configuration context / FPGA-loading identity. Both matter for joining records or interpreting counter coordinates.

References

[1] Christopher C. Gerry and Peter L. Knight. *Introductory Quantum Optics*. Cambridge University Press, 2005.

[2] D. F. Walls and G. J. Milburn. *Quantum Optics*. Springer, 2 edition, 2008.

[3] Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J. Cerf, Timothy C. Ralph, Jeffrey H. Shapiro, and Seth Lloyd. Gaussian quantum information. *Reviews of Modern Physics*, 84: 621–669, 2012. doi: 10.1103/RevModPhys.84.621.

[4] Roy J. Glauber. Coherent and incoherent states of the radiation field. *Physical Review*, 131:2766–2788, 1963. doi: 10.1103/PhysRev.131.2766.

[5] E. C. George Sudarshan. Equivalence of semiclassical and quantum mechanical descriptions of statistical light beams. *Physical Review Letters*, 10:277–279, 1963. doi: 10.1103/PhysRevLett.10.277.

[6] L.-M. Duan, G. Giedke, J. I. Cirac, and P. Zoller. Inseparability criterion for continuous variable systems. *Physical Review Letters*, 84:2722–2725, 2000. doi: 10.1103/PhysRevLett.84.2722.

[7] Ulf Leonhardt. *Measuring the Quantum State of Light*. Cambridge University Press, 1997.

[8] A. I. Lvovsky and M. G. Raymer. Continuous-variable optical quantum-state tomography. *Reviews of Modern Physics*, 81:299–332, 2009. doi: 10.1103/RevModPhys.81.299.

[9] Joseph C. Chapman, Joseph M. Lukens, Bing Qi, Raphael C. Pooser, and Nicholas A. Peters. Bayesian homodyne and heterodyne tomography. *Optics Express*, 2022. doi: 10.1364/OE.456597. Archived manuscript: arXiv:2202.03499v4.

[10] Kan Takase, Masanori Okada, Takahiro Serikawa, Shuntaro Takeda, Jun ichi Yoshikawa, and Akira Furusawa. Complete temporal mode characterization of non-Gaussian states by dual homodyne measurement. *Physical Review A*, 99:033832, 2019. doi: 10.1103/PhysRevA.99.033832. Archived manuscript: arXiv:1901.04669v1.

[11] Marco Avesani, Davide G. Marangon, Giuseppe Vallone, and Paolo Villoresi. Source-device-independent heterodyne-based quantum random number generator at 17 Gbps. *Nature Communications*, 9:5365, 2018. doi: 10.1038/s41467-018-07585-0. Archived arXiv:1801.04139v1 has a different manuscript title.

[12] Tobias Gehring, Cosmo Lupo, Arne Kordts, Dino Solar Nikolic, Nitin Jain, Tobias Rydberg, Thomas B. Pedersen, Stefano Pirandola, and Ulrik L. Andersen. Homodyne-based quantum random number generator at 2.9 Gbps secure against quantum side-information. *Nature Communications*, 12:605, 2021. doi: 10.1038/s41467-020-20813-w.