ajku 1.0: reference documentation

overview, theory, control reference, registers, test records

Rok Zitko, Jozef Stefan Institute

2026-11-13 (commit 6af1a9b)

Contents

Optical I/Q analyser — ajku 1.0 **2**
- Implemented and planned . . . 2
- Quick start without Vivado . . . 3
- Presets and specifications . . . 5
- Hardware setup and calibration . . . 5
- Verification . . . 5
- Documentation . . . 6

Measurement and implementation **6**
- What is measured . . . 6
- Exact code arithmetic . . . 6
- Continuous third- and fourth-order sums . . . 7
- Slow inputs and PPS . . . 7
- Full-rate histograms . . . 7
- Continuous FPGA spectra . . . 8
- Statistics service stream . . . 8
- Finite DDR capture . . . 8
- Host processing and coverage . . . 9

Service, client and recording format **11**
- Python . . . 11
- Continuous spectra (0.5) . . . 12
- Third- and fourth-order sums (0.4) . . . 12
- Slow inputs and PPS (0.3) . . . 12
- Continuous summary readout (0.2) . . . 12
- Raw capture metadata . . . 13
- `.iqbin` and HDF5 . . . 13
- MS1 host products . . . 13

Register map — prototype ABI 0x00040000 **14**
- Statistics record . . . 15
- Buffered statistics and histograms . . . 16
- Histogram windows . . . 16
- Timed captures . . . 17
- Continuous spectra . . . 17

Verification record **17**
- 2026-11-03 — release 1.0, build 0xfb4e4824 . . . 17
- 2026-11-03 — 0.5: continuous spectral accumulation, build 0x07bd2ef4 . . . 19
- 2026-11-02 — 0.4: continuous third/fourth-order sums, build 0x7322d68c . . . 21
- 2026-11-02 — 0.3: windows, timing, slow inputs and RTL tidying, build 0x733760cb . . . 22
- 2026-11-02 — committed baseline qualification and LaTeX/PDF manual . . . 24
- 2026-11-01 — six first-implementation priorities, host revision 4 . . . 26
- 2026-11-01 — full host FIR filtering/decimation workflow . . . 27
- 2026-11-01 — board repeat of committed code b2cd4c3 . . . 29
- 2026-11-01 — first-release host analysis, bounded pipeline and completed-buffer copy . . . 30
- 2026-11-01 — MS1 host workflow and native readout, FPGA 0.2 . . . 32

2026-11-01 — initial board qualification, build 0xd3f1cbc0 33
2026-11-01 — 0.2 buffered statistics and histograms, build 0xd3f1cbc0 35
2026-10-31 — acquisition/statistics prototype, build 0x3abc6145 35
Remaining board qualification 36

Optical I/Q analyser — ajku 1.0

A **CW I/Q measurement instrument** for two balanced receivers or another source of simultaneous I/Q voltages. It keeps both 14-bit inputs at 125 MSa/s and computes, continuously and exactly, the raw statistics through **fourth order** of every 8.4 ms block (Σ I … Σ Q^4: cumulants, excess kurtosis, $\kappa 22$ and $\langle r^4\rangle/\langle r^2\rangle^2$ over all the time, not only the captured 1–2 %), full-rate joint and marginal histograms with zoom windows, and **spectra**: every sample enters non-overlapping 1024-point Hann frames, and the FPGA accumulates per 134 ms epoch the exact power $|Z_k|^2$ and complementary product $Z_k Z_{(N-k)}$ of the FFT of z = I + iQ, from which `spectrum()` gives the auto spectra of I and Q, their cross spectrum and the complex and complementary spectra of z to 25 MHz in 122 kHz bins at 100 % duty. Finite records (also hardware-timed) are captured into DDR for PC analysis, and a PPS input and the slow inputs AI0–AI3 are recorded with every statistics block. Host science spectra extend to **25 MHz** at any resolution, with converter-skew correction. The MS1 host workflow adds live plots, saved-reference comparisons, traceable run recordings, Allan deviations of the statistics, LO-power fits and drift monitoring with follow-up captures.

Read the root equipment-safety and board information. Receiver outputs need suitable termination, attenuation and analogue anti-alias filtering. Full ADC resolution is retained; leave headroom for noise tails.

Status: release 1.0 (2026-11-03), digitally qualified on the board, including a one-hour workload without loss (Verification). PPS and slow-input accuracy have not been tested with a source. Analogue performance and optical sensitivity require characterization by the end user on their own board and optical chain. The user's manual explains operation, continuous-variable quantum optics, and those characterization procedures. From 1.0 the SD image builder includes ajku by default; images built before 1.0 do not contain it.

Implemented and planned

Available in 1.0 (FPGA and host MS1)	Later milestones
Two-channel finite DDR capture; sample indices; explicit error handling; hardware-timed starts; kept allocation	Continuous optimized bulk transfer and measured scheduled duty cycle
Continuous exact moments through fourth order, extrema, rails; PPS edges and AI0–AI3 means per block; 64-block FIFO and service statistics stream	Board-qualified sustained retention under concurrent load
Full-rate 64×64 joint and two 1,024-bin marginal histograms, configurable windows, out-of-window counts	Pre/post-anomaly (pretrigger) capture
Continuous FPGA spectra: 1024-point Hann frames, 2^14 per 134 ms epoch; I/Q auto and cross spectra, complex and complementary spectra of z, 122 kHz bins	Overlapping frames or finer FPGA resolution (needs a larger device)
Internal counter pattern for sample-order tests	Measured receiver/LO calibration and noise budget
Framework driver/client/service, binary download, full-code recordings	Optimized compiled raw transfer
Full host I/Q spectral matrix, coherence/phase diagnostics, signed/complementary spectra, complete third/fourth cumulants	Predictors and mode reconstruction
Explicit FIR pass/stop bands, integer decimation, raw/FIR views and exact saved profiles	Optical/analogue response correction
Raw/FIR amplitude/power/complex lag correlations; conditional batch-mean standard errors	Broader uncertainty models and optical validation
Captured tails/CDF/Q–Q, Gaussian residuals and batch-jackknife cumulant errors	Higher-order/mutual-information and projection extensions
Measured-LO-power sweep recording/fitting, signed noise-budget components and conditional errors	Actual optical qualification with independent references
Raw-statistics drift trends, threshold events and context-guarded host follow-up captures	FPGA pretrigger raw history / fast anomaly capture
Allan/modified deviations of the block statistics (also skewness, kurtosis, $\kappa 22$) and slow inputs; skew-corrected cross-spectra from the board's delay calibration	Skew correction of the FPGA spectra
Bounded two-job host acquisition/analysis pipeline; GIL-free completed-buffer copying	Further raw-transfer optimization and release-scale endurance

Available in 1.0 (FPGA and host MS1)	Later milestones
Host live dashboard, compatible saved references, integrated recording	Optical calibration and longer contention/endurance qualification

The FPGA statistics cover the **raw digitized response**, including frequencies outside the 25 MHz science band. Host band powers cover recorded intervals only. The latest-block register remains a snapshot. The new FIFO retains up to 64 blocks (0.537 s); use `statistics_stream()` for sustained readout through the service. FIFO overruns, slow-subscriber losses and index gaps are explicit. The FPGA spectra are raw broadband too (no anti-alias filter beyond the analogue chain), cover every complete frame without an invalid slot, and are exact integers until the host scales them. No DDC, QRNG or feedback is required for v1.

Quick start without Vivado

Use a board on which this prototype's overlay and plugin have been installed through the common setup. Until a qualified build is available, installation is a developer task; `make upload` copies files, and loading programs the FPGA. No board programming is implied by the examples below.

On the PC, from this directory:

```
export PYTHONPATH="$PWD/../common/python:$PWD/control:$PWD/tools"
# Use the current BOARD_HOST from ../config.mk for direct Python tools:
export BOARD_HOST=192.168.0.102
python3 -m rpfpga.cli --host "$BOARD_HOST" --preset acquire
python3 -m rpfpga.cli --host "$BOARD_HOST" --statistics
```

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

with CLIENT(os.environ["BOARD_HOST"]) as dev:
    print(dev.statistics())
    metadata, raw = dev.capture(2**21)  # 8 MiB, 16.78 ms; blocks until complete

save("vacuum-001.iqbin", metadata, raw)  # refuses to overwrite
v = samples(metadata, raw)               # connector volts from captured calibration
s = spectra(v, nperseg=65536)            # NumPy on the PC; Hann, 50% overlap
print(s["bin_hz"], s["enbw_hz"], s["segments"])
```

For the integrated interface, install NumPy, Matplotlib, ipywidgets and Jupyter on the PC and open **`tools/live.ipynb`**. See the MS1 measurement guide for source/processing controls, references, recording layout, stopping and reproducible command-line use. Opening the panel does not start acquisition. Optical-condition labels are supplied by the operator. The first-release checklist tracks remaining work beyond MS1; the current hardware/software is still a prototype, not a version 1.0 release.

Live summaries on a separate connection:

```
with CLIENT(os.environ["BOARD_HOST"]) as dev:
    h = dev.histogram()                # Busy until the first complete epoch
    print(h["count"], h["invalid"], h["missed_epochs"])
    sp = dev.spectrum()                # continuous spectra of the latest 134 ms epoch
    d = sp["densities"]               # V²/Hz: psd [I, Q], cross_psd conj(I)·Q, complex/complementary
    print(sp["raw"]["frames"], sp["raw"]["skipped"], d["bin_hz"], d["psd"][100])
    with dev.statistics_stream() as stream:
        packet = stream.read()         # exact block sums + calibration/context
        print(packet["lost_blocks"], packet["batch"]["dropped_blocks"])
```

By default histograms cover all native codes (256 codes per joint axis bin, 16 per marginal bin). A window in connector volts zooms them, for example onto an LO-off reference whose σ is a few tens of codes, or to one code per marginal bin for missing-code checks; spans round up to 2^k codes per bin:

```
dev.apply({"histogram": {"center_v": [0.0, 0.0], "joint_span_v": [0.02, 0.02],
                         "marginal_span_v": [0.01, 0.01]}})
h = dev.histogram()          # window, joint_window_codes, voltage_range, *_outside/below/above
dev.apply({"histogram": {"center_v": [None, None], "joint_span_v": [None, None],
                         "marginal_span_v": [None, None]}})   # back to the full span
```

The panel's **Centre ±6σ on latest block** stages such a window. Samples outside a window are counted, never dropped. Arrays are raw counts, not density-normalized. Snapshots contain 2^24 clock slots (134.22 ms), aligned to statistics blocks; invalid slots are counted separately. Three banks keep counting while one bank is read. A slow reader can miss whole snapshot epochs, which are counted rather than mixed.

Without the service worker, `dev.statistics_blocks(maximum=32)` drains the FPGA FIFO directly. It is destructive: use one drainer. The service refuses DRAIN while a statistics stream is subscribed. Direct cross-process consumers are detectable as sample-index gaps but must not be used as competing drainers.

`statistics()["higher_order"]` gives the latest block's central moments, joint cumulants, skewness, excess kurtosis, κ22/(κ20κ02) and $\langle r^4\rangle/\langle r^2\rangle^2$; for streamed or drained blocks use `ajku.model.higher_order(raw, adc)` on the PC. Every statistics result also carries `slow_inputs` (AI0–AI3 block means in volts, nominal scale until the common slow-input calibration is measured) and `pps` (edge count and the sample index of the latest edge on DIO0_P). Allan deviations of the recorded stream:

```
python3 tools/iqstability.py --host "$BOARD_HOST" --seconds 600 --record /path/to/new-stats.jsonl \
    --output /path/to/new-adev.json --plot /path/to/new-adev.png \
    --metrics mean_i,mean_q,variance_i,variance_q,relative_power,ai0
```

A capture can start at an exact sample index (`dev.sample_index()`, then `dev.capture(pairs, start_index=...)`); `record.py --timed` schedules its seeded, signal-independent cadence in sample indices this way.

For scheduled recordings:

```
python3 tools/record.py --host "$BOARD_HOST" --directory /path/to/new-run --captures 10
python3 tools/record_statistics.py --host "$BOARD_HOST" --output /path/to/stats.jsonl --batches 100
python3 tools/iqmeasure.py --host "$BOARD_HOST" --output /path/to/new-analysed-run --captures 100
 ↪  --pipeline --lags 256
```

The last command overlaps capture/download with analysis/storage, with at most two admitted jobs. It saves raw temporal correlations and batch-based errors as well as spectra. Its throughput-driven `pipeline` selection differs from `record.py`'s seeded schedule. Use actual coverage and the recorded selection policy in analysis.

The nominal 0.5 s cadence with 8 MiB records would give 3.36% raw coverage, about 16.8 MB/s and 60 GB/hour. These are **targets/calculations**: concurrent statistics and repeated histogram readout gave 2.18% coverage in the initial board test. The prototype copies each completed buffer before binary transfer; it does not yet use the optimized common C pump. Native histogram readout improves responsiveness, but full host analysis and storage add their own time. The integrated 120-second MS1 workload retained 0.90% raw coverage; consult the verification record for the precise workload. Missed scheduling deadlines in `record.py` are skipped.

The host-only digital qualification tool is reproducible with:

```
python3 tools/qualify.py --host <board-host> --seconds 10 --output /path/to/new-results.json
python3 tools/qualify_ms1.py --host <board-host> --seconds 120 --directory /path/to/new-evidence
 ↪  --dashboard
python3 tools/qualify_readers.py --host <board-host> --stall-seconds 60 --output
 ↪  /path/to/new-reader-results.json
python3 tools/qualify_features.py --host <board-host> --output /path/to/new-feature-results.json
python3 tools/qualify_spectra.py --host <board-host> --output /path/to/new-spectra-results.json
```

`qualify_spectra.py` checks the FPGA spectra: every sum of a counter epoch against the bit-exact model, consecutive ADC epochs and their alignment with the histograms, Parseval against the statistics and the FPGA band PSD against host Welch spectra of a capture at the same resolution.

`qualify_features.py` checks the 0.3 features through the service: a timed ADC capture and the refusal of a past start, slow inputs/PPS in the stream, ADC histogram windows with conserved counts, and the kept allocation around a 64 MiB request.

Use `BOARD_HOST` from the repository-root `config.mk`. This test selects the internal counter, checks transfer and summaries, cancels one busy capture, and compares actual ADC samples with their matching FPGA moment block. It returns to ADC/calibrated-zero; it does not calibrate the analogue or optical chain. The MS1 tool additionally records/reopens raw and processed data, compares saved references, probes concurrent control latency and optionally exercises the widgets. `qualify.py --maximum-pairs 8388608` limits the largest test to 32 MiB when the 64 MiB contiguous allocation cannot be satisfied; no allocation failure is hidden. The reader test deliberately stops one statistics consumer, checks that losses are reported without disrupting a fast consumer, then closes a completed capture's download early and verifies recovery. It does not terminate the board service.

Presets and specifications

- `acquire`: ADC source, outputs set to the calibrated zero codes.
- `counter`: deterministic signed ramp/complement pair, independent of ADC inputs. For digital integrity testing; no physical-voltage interpretation.
- `SAFE`: selects ADC, returns outputs to calibrated zero and aborts an active capture. Accepted AXI transactions drain before the buffer becomes reusable.

Quantity	Prototype specification
Inputs	Two synchronous native signed 14-bit channels, 125 MSa/s nominal
Science spectrum	Up to 25 MHz; actual analogue response/alias rejection unmeasured
Statistics	2^20 clock slots = 8.388608 ms; exact sums of all monomials through I^a Q^b, a+b ≤ 4 (60/74-bit third/fourth order); invalid slots explicitly counted
Statistics buffering	64 completed blocks; drop-new overflow counter; 64 batches per stream subscriber
Histograms	64×64 joint + 2×1024 marginals, three banks; 2^24-slot epochs; per-axis window centre and 2^k codes per bin (joint 1…256, marginal 1…16), out-of-window counts
Rail indications	Exact end codes; near-rail outer 128 codes on either end
Capture	32-pair multiples; default 2^21 pairs; software ceiling 2^24 pairs (64 MiB)
Buffer	Contiguous board DDR allocation, kept for later captures once drained (32 MiB when available); the first allocation can still fail
Time	64-bit sample-clock index; host wall time separately labelled; PPS on DIO0_P counted and indexed per block (±1 clock, 2-clock latency)
Slow inputs	AI0–AI3 (E2, 0…7.0 V nominal): XADC, several hundred sequences summed per statistics block; accuracy and bandwidth unmeasured
Timed capture	Start at a given sample index, at least two samples and at most 60 s ahead
Spectrum resolution	Host: set by segment/record length; FPGA: 1024-point frames, 122.07 kHz bins (Hann ENBW 183 kHz)
FPGA spectra	z = I + iQ, Hann window (Q17), radix-2^2 FFT (18-bit input, 25-bit whole-code output, bit-exact Python model), non-overlapping frames aligned to the statistics blocks; per epoch of 2^14 frames (the histogram epoch) 64-bit sums of $\|Z_k\|^2$ (1024 bins) and $Z_k Z_{(N-k)}$ (512); two banks, the latest epoch wins unless claimed; frames with an invalid slot skipped and counted; fixed-point noise 0.3 % of the ADC noise at 1.3 codes rms (model)
Outputs	Digital zero after loading; calibrated zero on apply; reference levels during common calibration

The 64 MiB software ceiling bounds peak allocation/copy/transfer memory in this prototype. It can be revisited after memory and transport qualification. There is no imposed small FFT-length ceiling on the host. A 64 MiB allocation passed initial tests but failed during a later repeat; the default 8 MiB workload passed the MS1 bounded soak. The maximum is an API ceiling, not guaranteed available memory.

Hardware setup and calibration

Use IN1 for I and IN2 for Q. The design is detector-agnostic. A PPS (3.3 V LVCMOS, rising edges, **not 5 V tolerant**) goes to DIO0_P on E1; slow references such as a receiver's DC monitor outputs or a tap photodiode go to AI0–AI3 on E2 (0…7.0 V nominal). Record the receiver, termination, coupling, attenuation and analogue filters with each experiment. The Red Pitaya inputs are high impedance; a receiver specified into 50 Ω needs an appropriate external load. A GHz-band receiver needs characterized anti-alias filtering before a 125 MSa/s ADC; cropping a spectrum cannot undo aliasing.

Connector volts use the shared board calibration. Common calibration can measure input means and drive DAC reference holds. Optical gain/quadrature-angle/delay calibration additionally needs a coherent probe; LO-power scaling and receiver compression need independent power information. Neither is inferred from a circular vacuum histogram or ADC full-scale use.

Verification

Release 1.0 passes the offline tests, nine simulations (the FFT and spectra bit-exact against the Python model; statistics, histograms, FIFO, capture and the register interface through AXI transactions against in-bench references) and the digital board qualification (2026-11-03): exact counter captures, statistics, histograms and spectra; an actual ADC block equal to the FPGA's fourteen sums; FPGA spectra within 0.4 % of host Welch spectra; a one-hour full-analysis workload with 1,657 captures and 429,551 statistics blocks without loss (control p99/max

141/334 ms); 32 MiB captures, cancellation, readers and the real-browser notebook. Build, timing and board results: docs/VERIFICATION.md. Optical and analogue performance are not measured: the end user characterizes their own receiver chain (manual, characterization chapter).

Documentation

- User's manual (PDF), LaTeX sources and build instructions
- Supported interface baseline and operating limits
- Measurement theory and architecture
- MS1 live measurements, references and recording
- Host filtering/decimation, sample coordinates and replay
- First-priority calibration, distributions and drift workflows
- Registers, protocol and binary format
- Verification and outstanding qualification
- Research wiki and approved plan, archived papers
- Original feature specification

For developers: `make test`, `make sim`, `make`; all targets use the shared workflow.

Measurement and implementation

What is measured

I and Q are voltages from simultaneous converter samples. In an appropriately calibrated optical double-homodyne arrangement their joint distribution is related to the Husimi Q distribution, including joint-measurement vacuum noise. Voltage histograms are not automatically a Wigner function or evidence of optical squeezing. Independent vacuum outcomes need not have a common signal in their cross-spectrum.

The generic analyser does not assume detector gain, wavelength, optical efficiency or LO power. Board calibration establishes connector volts. Optical interpretations need the additional receiver/hybrid calibration described in the research wiki.

Exact code arithmetic

The common ADC outputs signed 16-bit stream words with two zero low bits. Recover native codes by arithmetic division by four. Each statistical block occupies 2^20 clock slots; a missing channel-valid signal makes that slot invalid, rather than moving a later sample into its place. The first slot has a 64-bit sample index.

The pipeline registers I, Q and their three signed 28-bit products, then accumulates signed 34-bit first sums and signed 48-bit second sums. Worst cases fit even at constant full scale. A completed block includes its final sample and the next block starts without dead time. Counts, extrema, clipping and near-rail occupancy are accumulated for valid pairs only. The invalid count makes missing slots visible.

For common calibration `stream_code = K V + A`, native code `x` gives `V = (4 x - A)/K`. The host centers moments using integer numerators, for example `(N sum(x²) - sum(x)²)/N²`, before converting to volts. These are population moments; confidence intervals require temporal-dependence and calibration information. Counter-test data stay in code units. The prototype applies the DC gain/offset calibration. Optional common AC response and delay fields are preserved in recording metadata but are not automatically equalized; a calibrated frequency-dependent correction remains host follow-on work.

The latest completed block is held for a whole block (8.39 ms) and read in place, independently of the 64-block FIFO. It changes in a single clock, so a burst read that begins and ends with the same epoch word is one consistent block; the driver re-reads otherwise. There is no snapshot copy. The FIFO serializes each 32-word record into a 32-word RAM slot directly from that held record; two BRAM36 tiles hold 64 slots. A slot becomes visible only after its last word is written. The reader reads the oldest slot word by word in place and POP retires it. A full FIFO drops new blocks and increments a counter. It can absorb about 0.537 s of reader delay, not an unbounded outage. APPLY discards the partial block, flushes the FIFO and clears the latest record so calibration/source epochs cannot mix.

Continuous third- and fourth-order sums

Each block also accumulates the nine raw monomials of degree three and four, $\Sigma\, I^3 \ldots \Sigma\, Q^4$, exactly: the squares I^2, IQ, Q^2 feed four 28×14 and five 28×28 products (DSP slices, three output registers), and the sums are 60 and 74 bits wide, so even a constant full-scale input cannot overflow. Every other quantity of a slot is delayed to the same stage, and a tag travelling with each slot drops those in flight at an APPLY, so blocks keep their boundaries. The host forms central moments from integer numerators, `C_ab = Σ (N x − S_I)^a (N y − S_Q)^b` (central moment `C_ab / N^(a+b+1)`), and fourth-order cumulants such as $\kappa 40 = (N\, C40 - 3\, C20^2)/N^6$ with a single final division: no large raw sums are subtracted in floating point, however large the mean.

This makes the non-Gaussianity statistics continuous: every sample, not the 1–2 % in captures, enters κ30…κ04, the excess kurtoses, κ22 and $\langle r^4\rangle/\langle r^2\rangle^2$ (2 for circular Gaussian noise; not an optical g^2). Over a given time their statistical error is about √(1/coverage) smaller, roughly ten times. Blocks can be summed exactly for any window (the drift monitor does), and the kurtosis of any projection $X(\theta) = \delta I \cos\theta + \delta Q \sin\theta$ follows from the five fourth-order cumulants. These are raw broadband statistics of the digitized voltages: the band, the receiver response, the ADC's own nonlinearity and dither all enter, and a non-zero excess kurtosis is evidence about the measured process, not about the optical state without a calibrated model. Batch errors come from the block-to-block scatter (`tools/iqstability.py`); captured records keep their batch-jackknife errors.

Slow inputs and PPS

The same block records two slow, independent inputs. The XADC runs the common continuous sequence (as in feedback, `common/rtl/optional/rp_xadc.v`): about one sequence every 15 µs, several hundred per block. The block holds their count and the sums of the four 16-bit AI words (E2 AI0…AI3, behind the board's dividers); the host forms block means in volts through the common slow-input calibration. They are meant for slowly varying references logged against the noise statistics: a receiver's DC monitor outputs, a tap photodiode on the LO, a temperature sensor. Their bandwidth and accuracy are those of the XADC and dividers, not of the fast inputs; an LO-power monitor still needs its own calibration against an independent power meter before anything is called watts.

A rising edge on DIO0_P (3.3 V LVCMOS) is synchronized by two flops and counted; the sample index two clocks after the edge is first sampled is recorded. Every block holds the edge count and the latest edge's index at its end, so the gapless statistics stream carries the PPS timing without polling: with a disciplined PPS, `index` differences measure the board's actual sample rate (the nominal 125 MHz is the crystal's), and a UTC-aligned PPS ties sample indices to UTC to ±1 clock plus the pin's delay. Capture metadata records the latest edge before arming.

Full-rate histograms

Each axis has a window: a centre code c and a bin width of `2^s` native codes. With B bins (64 joint, 1024 marginal) the window starts at `low = c − (B/2)·2^s` and a code x falls in bin `(x − low) >> s` when $0 \le x - \mathrm{low} < B \cdot 2^s$. The joint index is `(I_bin << 6) | Q_bin`, inside only if both axes are. s = 8 (joint) and s = 4 (marginal) at c = 0 are the full ADC span (0.2's fixed bins); a narrow window resolves an LO-off or electronic-noise reference whose σ is a few tens of codes, and a 1-code marginal shows missing codes and differential nonlinearity directly. Samples outside a window are counted per epoch (joint outside, each marginal below/above), so every valid sample is accounted for. The bins are computed one register stage ahead of the banks, with valid, epoch starts and indices delayed alike. The host converts connector-volt windows to codes through the board calibration at Apply; each frozen epoch carries its own window in its metadata. Changing the window is an APPLY and starts new epochs.

Each of three banks contains 4096 joint and 2×1024 marginal 32-bit counters, using six RAMB36 tiles per bank. Eighteen tiles are deliberately used instead of the initial two-bank plan's twelve: one bank can stay claimed while the other two alternate accumulation and clearing. No new clock or multiplier is needed.

Each RAM reads a bin, registers its increment, then writes it. Forwarding from both the pending increment and last committed write handles consecutive/revisited addresses, including same-clock cross-port memory collisions. End-of-epoch metadata is published only after all pending increments reach RAM.

An epoch spans `2^24` slots, aligned to 16 complete `2^20`-slot statistics blocks. Initial clearing waits for the next statistics boundary; `gap_slots` counts those uncovered slots. With the production parameters the two non-reader banks can clear well before reuse, so a held reader does not introduce further sampling gaps. Every unclaimed frozen snapshot is replaced by the next complete epoch. CLAIM protects a snapshot until RELEASE; meanwhile new epochs that cannot be published are discarded. Both replacement and discard increment `missed_epochs`. This is continuous computation with explicit finite-retention losses, not guaranteed archival coverage. First indices, valid/invalid counts and loss counters must accompany plots.

Continuous FPGA spectra

Version 0.5 transforms every sample. The engine (`src/iq_spectra.v`) windows z = I + iQ in non-overlapping frames of N = 1024 samples, aligned to the statistics blocks, and sums over an epoch of 2^14 frames (2^24 slots, the histogram epoch, 134 ms) the power P_k = |Z_k|² of every bin and the complementary product C_k = Z_k Z_(N−k) for 0 ≤ k < N/2. Because I and Q are real, their transforms follow from Z: X_k = (Z_k + Z*(N−k))/2 and Y_k = (Z_k − Z*(N−k))/(2i), so per frame

```
|X_k|² = (P_k + P_(N−k) + 2 Re C_k)/4,   |Y_k|² = (P_k + P_(N−k) − 2 Re C_k)/4,
conj(X_k) Y_k = Im C_k / 2 − i (P_k − P_(N−k))/4,
```

and the sums give the epoch means exactly. `model.spectral_densities` divides by frames · fs · Σw² (one-sided auto and cross spectra, the host convention conj(I)·Q; two-sided complex and complementary spectra of z), as `tools/iqanalysis.py` does with `nperseg=1024, overlap=0, detrend=False`. Bins are 122.07 kHz; the Hann window's ENBW is 183.1 kHz. The converter offset is not removed: bins 0 and 1 hold its leakage. The Nyquist bin's complementary product is not formed (P_512 is).

For Gaussian noise each bin's estimate from F frames has a relative standard error of 1/√F: 0.78 % per epoch, 0.1 % after a minute. Non-overlapping Hann frames of white noise are independent, so F is the number of averages. A host Welch estimate at the same resolution from 1–2 % capture coverage has 26–58 times fewer, so the FPGA errors are 5–8 times smaller in the same time. Resolution is fixed at 122 kHz; finer structure (narrow lines, low-frequency 1/f noise) still needs host spectra of long captures. The spectra are raw broadband like the other FPGA products: aliasing is set by the analogue chain, and no converter-skew correction is applied.

Arithmetic (bit-exact in `ajku.model`, checked against XSim and, through the counter source, on the board): the window table is w_n = sin²(π(n + ½)/N) in Q17 (sampled at half-integer points it has the periodic Hann window's three-bin transform, Σw = N/2, Σw² = 3N/8); each product x·w keeps four fractional code bits (18-bit components). The FFT is a radix-2² single-delay-feedback pipeline (`src/iq_fft.v`, one sample per clock, 1,059-clock latency): ten butterfly stages, each adding one bit; the second of every pair multiplies by −j exactly; after each of the first four pairs a Q16 twiddle rotation (three DSP products) rounds one fractional bit away, so the output is 25-bit whole codes. The rounding noise referred to the input is about 5/3072 code² per component and sample: 0.3 % of the windowed noise power for an ADC noise of 1.3 codes rms (5 % if only quantization noise were present). The Q16 twiddles add an error about 10^{-5} of the signal (−100 dB). With Hann weights |Z_k| ≤ 2^22.5 codes, so the 64-bit sums cannot overflow in an epoch (P ≤ 2^59).

The FFT emits bins in bit-reversed order. The two members of each pair (k, N − k) are mirrored inside each octave of the output position, so the first one waits in a 512-entry pairing RAM for the second; then six 25×25 products and two banks of 64-bit read-modify-write sums (16 BRAM tiles) complete the engine. The first accumulated frame of an epoch writes its bank instead of adding, so no clearing pass is needed. Latest-wins publication, claims, missed epochs and APPLY follow the histograms (see REGISTERS); a frame with an invalid slot is skipped and counted.

Statistics service stream

The service worker drains bounded batches only while subscribers exist. Raw captures remain on command threads; the summary worker neither allocates DMA nor changes configuration. Each batch includes the session/revision, board calibration/ranges, block length and hardware drop count. Per-block `gap_before` follows exact epoch, sample-index or configuration discontinuities; missing ADC slots remain in `invalid`. A drop-count change is separately reported, because a lost new block can lie after older queued blocks rather than before the next emitted block.

The worker publishes immutable JSON to per-subscriber queues of 64 batches. A slow subscriber loses oldest whole batches; its cumulative `lost_blocks` is independent of the hardware's drop count. This backpressure cannot stall the FPGA or another subscriber. Reload/worker failure closes the stream explicitly. The stream generator does not touch registers. Histograms use a separate claimed readout transaction; the native 6144-word MMIO read measured 24.3 ms board-local in MS1 (see VERIFICATION).

Finite DDR capture

Capture starts on the sample after START reaches the writer, or, with START_TIMED, at exactly the sample whose index was written to the start-index register; the writer waits busy without AXI traffic until then. The index must be at least two samples after the START clock or the request is refused (error 5) with no writes. Timed starts make a scheduled, signal-independent sampling plan exact in sample-clock units, without host jitter; `tools/record.py --timed` uses it. Each 32-bit word holds signed native int16 I in bits 15:0 and Q in bits 31:16. Consecutive words are packed into a 64-bit beat, earliest pair low. A synchronous 2,048-beat BRAM FIFO provides 16 KiB of elasticity

(32.8 µs at 500 MB/s). It does not guarantee adequate DDR service. Its high-water mark since the last start is recorded with every capture, so the margin left under a given DDR/CPU load is measured rather than assumed.

The write-only AXI3 HP0 master uses 16-beat/128-byte bursts, four responses outstanding, constant ID zero and ordered write data. The base is 128-byte aligned, the pair count a multiple of 32, and bursts cannot cross 4 KiB boundaries. The hardware limits writes to the board's low DDR address region and validates the end address before accepting samples. The driver only supplies a live PYNQ contiguous allocation, not an arbitrary user-supplied physical address.

On overflow, invalid ADC data, abort or non-OKAY/wrong-ID response, the record is invalid. New bursts stop; an already offered AW transaction and its data remain valid until accepted. Busy remains high until all issued writes receive responses. No partial record is published as a successful contiguous acquisition. A stalled AXI target can keep the writer busy indefinitely: a software timeout is not proof that DMA has stopped.

The driver uses short framework transactions and an acquisition token. It copies the buffer only after DONE, then verifies the acquisition revision before publication. The contiguous allocation is kept for the process once drained and reused by later captures (first allocated at 32 MiB when available, so captures up to that size never allocate again); repeated allocation and release fragmented the board's CMA region, which is why a 64 MiB request could fail after smaller ones. Only a drained allocation returns to the pool; an uncertain one is quarantined as before. Cancellation cleanup checks the active physical base so it cannot abort a newer capture. If hardware ownership/drain is uncertain, a process-level quarantine keeps the allocation alive for the process lifetime. Process exit (including crashes) while DMA is active needs board qualification; this prototype makes no unattended crash-recovery claim.

Host processing and coverage

The driver keeps one immutable completed capture for binary download. The common `immutable_copy` helper can release the Python GIL while copying completed DDR into a private immutable bytes object. It is called only after DONE/drain and cache invalidation; the acquisition context is rechecked before publication. `copy_backend` and `board_copy_s` identify the actual copy path and elapsed time. This improves concurrency, not the FPGA's capture duty cycle by itself.

`tools/record.py` requests signal-independent, jittered scheduled bursts. It does not concatenate their samples: each record retains its own first sample index and session.

`tools/iqanalysis.py` uses `S_IQ = mean(conj(FFT(I)) FFT(Q))`: Q delayed relative to I has negative cross-phase. The two converter paths differ by a small delay: a 1 ns skew is 9° of cross-phase at 25 MHz. When the common loopback calibration has measured both inputs' `delay_s`, processing revision 5 adds `skew_*` products with Q's transform advanced by the difference in every segment; auto-spectra, coherence and eigenvalues are unchanged, and the uncorrected products stay. Cables and receivers are not included: only the board's converters are corrected. Auto/cross spectra of real channels are one-sided (DC and even-length Nyquist are not doubled). The spectrum of `z=I+iQ` is two-sided; the complementary spectrum is `mean(Z(f) Z(-f))`. No complex spectrum receives a one-sided factor of two.

Normalization is `fs sum(window²)` per segment. Parseval compares the full spectrum against the same detrended, windowed time-domain power before cropping to 25 MHz. Band-integrated powers describe that spectral band, not the FPGA's broadband raw statistics. Host time-domain fourth moments remain broadband unless the caller explicitly enables the separate FIR analysis branch; the raw products are never quietly relabelled as bandlimited.

The Hann window is periodic, with ENBW `1.5 fs/N`. Welch overlap reduces independent averaging; the reported effective-average estimate accounts only for window overlap under white independent input. It is not a confidence bound for arbitrary colored or drifting noise. A single segment has coherence one wherever both powers are nonzero, which is not evidence of coupling. State reconstruction, entropy/security claims, and uncertainty-aware higher-order inference remain later host work.

Temporal correlations and conditional uncertainty

Processing revision 2 adds optional **raw broadband** temporal products from one contiguous captured record. Their band is independent of the spectrum's cropped display/integration limit. A positive lag in `R_ab(k) = mean(a[t] b[t+k])` pairs channel a with a later sample of b: a positive I/Q peak means Q is delayed. Whole-record channel means are removed first. Squared-amplitude fluctuations are then centred by whole-record powers; radius-squared fluctuations use `r² = δI² + δQ²`. Complex ordinary and complementary correlations use `mean(z[t] conj(z[t+k]))` and `mean(z[t] z[t+k])`, respectively.

FFT work is divided into bounded contiguous blocks, default 16,384 pairs. Padding prevents circular wrap. Only within-block lag pairs are accumulated; pairs across block edges are deliberately omitted, and the actual denominator is saved for every lag. Partial final blocks contribute their valid pairs. No pair crosses a capture gap. Ordinary, squared-amplitude and radial-power covariances and their normalized correlations are saved, including both lag signs. Undefined normalized values for constant data remain NaN. Weak signed spectral images and radial-power fluctuations are formed before squaring, avoiding cancellation of large powers.

Uncertainty uses non-overlapping batch means at base block length, twice that length and four times that length. At least eight complete batches are required at each reported scale. For batch estimates b_j, the reported standard error is

$SE^2 = \mathrm{sum}((b_j - \mathrm{mean}(b))^2) / (B\,(B-1))$.

The five statistics are mean I, mean Q, variance I, covariance IQ and variance Q. For variance/covariance, each batch averages centred products using the record's global mean. These are asymptotic stationary estimates: the batches must be long relative to dependence **of that statistic**, including squared-amplitude dependence. Comparing scales helps diagnose insufficient batch lengths; it does not automatically establish a valid confidence interval. Incomplete tails are excluded from uncertainty batches and counted separately. Calibration uncertainty and inter-capture drift are not included.

Effective sample counts compare each statistic's single-sample influence variance with its batch-estimated SE^2. They are not shared between means and powers, are undefined for zero SE, and can exceed N for anticorrelation. They are descriptive estimates under the same stationarity/long-batch assumptions, not entropy bounds. Periodic counter data remain diagnostics, even if their batch errors are zero.

The host stores all four third-order and all five fourth-order joint cumulants of each capture (the FPGA blocks carry the same statistics continuously, above). Raw time-domain cumulants remain broadband; the optional FIR branch computes its own cumulants from filtered/decimated values. Revision 4 adds re-centred delete- batch jackknife errors/bias estimates and marginal-tail diagnostics. These still require stationarity and sufficiently long batches; they do not include optical calibration uncertainty. Predictors and broader dependence tools remain follow-ons.

Stability of the block statistics

The statistics stream is gapless at 8.39 ms per block, so its quantities (means, variances, covariance, total and relative noise power, circularity, slow inputs) form time series from milliseconds to hours. `tools/iqstability.py` computes their overlapping Allan and modified deviations with the common stability library: a block value is a frequency-type sample, integrated to a phase-type series so the shared estimators (and their noise-identification confidence intervals) apply. White fluctuations fall as $\tau^{-1/2}$ (the statistical floor of a block estimate), flicker is flat, drift and random walk rise. Gaps and context changes split the series and the longest gap-free segment is analysed. These are raw broadband statistics; the relative-power deviation, for example, is the stability of the total digitized noise power, which an LO-power monitor on AI0 can be compared with directly.

Explicit FIR bands and decimation

Processing revision 3 applies one real symmetric, unity-DC FIR identically to I/Q on the PC. It accepts pass/stop edges, peak-to-peak ripple, attenuation, integer decimation and a global sample-grid phase. A checked Kaiser design or exact saved coefficients are used; stopband onset must be no higher than the new Nyquist. Only complete FIR windows contribute, with output times referred to their centres. No filter history crosses gaps. Raw arrays and calibrated broadband summaries remain available alongside filtered products. FILTERING.md specifies the coefficient/response checks, integer coordinate equations, edge policy, normalization and reconstruction/export contracts. Digital filtering cannot undo analogue aliasing.

First-priority inference and drift

PRIORITIES.md defines the full matrix/coherence approximation, Gaussian-bin diagnostics, tail and cumulant error methods, power-dependent fits and monitoring semantics. In particular, the LO model $A+B\,P+C\,P^2$ is phenomenological: its linear term is not independently proven quantum. Capture-repeat errors and shared/independent meter errors are distinct; fits retain negative coefficients and lack-of-fit diagnostics. Raw/FIR integrated bands are fitted from each capture's band products rather than sums of supposedly independent per-bin variances.

Drift trends aggregate integer block sums before converting units. Context/loss/ index/invalid-slot boundaries split windows and reset baselines. Thresholds can request a revision-guarded finite follow-up capture. Trigger indices,

actual capture delay and absence of pretrigger raw data are preserved, so a host response cannot be mistaken for a waveform covering the event that caused it.

Service, client and recording format

Uses the common `rp-overlay.service`, TCP port 5025 and the uniform config/session API in common/README.md. Configuration is `{"source": "adc"|"counter", "histogram": {...}}`. Unsupported keys are rejected before writes. Standard presets are `acquire` and `counter`.

`histogram` sets the FPGA histogram window in connector volts, per axis `[I, Q]`: `center_v` (null: native code 0), `joint_span_v` and `marginal_span_v` (null: all 16384 codes). A span is rounded **up** to the next available window (64 joint or 1024 marginal bins of 2^k native codes, k = 0…8 joint, 0…4 marginal) and limited to the full span; the centre rounds to the nearest native code and must lie inside the input's range. The board calibration and jumper setting at Apply define the conversion; the status `histogram_window` and every histogram result report the realized windows in codes and volts. The default (all null) is the full native span of 0.2. Changing the window is an APPLY: it flushes statistics, invalidates histograms and aborts a capture.

Command	Reply
`STATS?`	`OK` + JSON: raw latest-block sums, optional connector-voltage results, calibration/ranges, revision/session
`STATS DRAIN [maximum]`	`OK` + JSON: up to 32 buffered blocks, common calibration/context and FIFO drop count
`STATS STREAM`	Dedicated connection: header, then `STATS` JSON packets from the service worker
`HIST?`	`OK` + JSON: latest claimed histogram epoch, raw counts, its window, out-of-window counts, calibration and loss/coverage metadata
`SPEC?`	`OK` + JSON: latest claimed spectral epoch: exact 64-bit sums (`raw`), densities to 25 MHz, frames/skipped/missed, calibration and context
`INDEX?`	`OK` + JSON: `sample_index` (the board's sample clock now), `sample_rate_hz`, `session`
`CAPTURE [pairs [AT <index>]]`	Long command; `OK` + metadata after complete capture and copy; default 2^21 pairs; `AT`: hardware-timed start at that sample index
`CAPTURE DATA <id>`	Separate connection: `OK` + metadata line, then exactly `payload_bytes` raw bytes, then EOF

`CAPTURE` accepts multiples of 32 from 32 through 2^24. The driver deadline is five seconds by default (plus the wait of a timed start); direct board callers can supply `timeout_s`. `AT <index>` starts the record at exactly that sample-clock index (see `INDEX?`), at most 60 s ahead; an index already passed is refused before any write, and the hardware also refuses one that passes while arming (error "timed start already passed", no record). The command claims the shared acquisition slot without holding the board lock while waiting. Other acquisition requests return BUSY; configuration edits cancel it. The normal client session wrapper protects control commands. The binary client independently checks the stream's session, UUID, format, length and completion/error fields.

Only the latest completed record is retained by the service's driver. A second completed capture supersedes it: an old identifier returns an error, never another record's bytes. A transfer that already took its immutable snapshot can finish even if a new capture completes. The stream generator does not touch hardware registers. Partial capture, timeout or uncertain transfer is not automatically retried.

Python

On the board, `DRIVER.capture(pairs)` returns metadata, followed by `DRIVER.capture_data(metadata["id"]) -> (metadata, bytes)`. On the PC, `CLIENT.capture(pairs)` performs both steps and returns that tuple; both take `start_index=` for a hardware-timed start, and `sample_index()` returns the board's sample clock. Both provide `statistics()` and the common configuration/calibration API. Both also provide `statistics_blocks(maximum=32)`, `histogram()` and `spectrum()`. The PC client provides `statistics_stream()`, a context manager whose `read()` returns a packet. `CLIENT.capture(pairs, expected_revision=...)` additionally guards CAPTURE admission through the existing REQUEST envelope. A changed revision is rejected before allocation/start; normal session and post-acquisition checks still apply. Host calibration and triggered-follow-up workflows use this to pin their context.

Continuous spectra (0.5)

`SPEC?` (`spectrum()`) claims the frozen spectral bank, reads it in one transaction (4,096 words) and releases it; `Busy` until an epoch is complete after an APPLY. `raw` holds `epoch`, `frames` (accumulated), `skipped` (frames with an invalid slot), `first_index` (u64, the epoch's first sample, equal to the histogram epoch's), `source`, `log2n` (10), `power` (1,024 integers, $\Sigma|Z_k|^2$) and `complementary` (512 `[re, im]` pairs, $\Sigma\, Z_k\, Z_{(N-k)}$): exact integers of the bit-exact model (`ajku.model.spectral_epoch`). `densities` (`model.spectral_densities`; null without frames): `frequency_hz` (0…25 MHz, 205 bins of 122.07 kHz), `psd` (`[I, Q]` one-sided), `cross_psd` (`[re, im]` of conj(I)·Q, one-sided), `signed_frequency_hz`, `complex_psd` and `complementary_psd` (`[re, im]`) of z = I + iQ (two-sided), `units` (`V^2/Hz` for the ADC, `codes^2/Hz` for the counter), `frames`, `bin_hz`, `enbw_hz`. `frames_per_epoch` (2^14), `epoch_slots` (2^24) and `missed_epochs` (since APPLY) complete the result. The reply is about 100 kB of JSON.

Third- and fourth-order sums (0.4)

Every `raw` block (latest, drained, streamed) carries the exact integer sums `sum_iii`, `sum_iiq`, `sum_iqq`, `sum_qqq` and `sum_iiii`, `sum_iiiq`, `sum_iiqq`, `sum_iqqq`, `sum_qqqq` of native codes (Python integers; JSON integers up to 2^72 in magnitude). `STATS?` adds `higher_order` for the latest block (`ajku.model.higher_order`): `central_moments` and `cumulants` keyed `"ab"` for I^a Q^b (volts^(a+b) from the board calibration, or native codes for the counter), `skewness`, `excess_kurtosis` (per channel), `kappa22_normalized` and `radial_kurtosis_ratio` ($\langle r^4\rangle/\langle r^2\rangle^2$ of the centred radius). They are population values of one raw broadband block, formed from exact integer numerators. Drained and streamed blocks carry only the sums; hosts call `model.higher_order(raw, adc)` (and `model.projection_kurtosis` for X(θ)), so the board's CPU does not convert every block.

Slow inputs and PPS (0.3)

Every statistics result and block carries the raw words 24–31 of the record (`ai_count`, `ai_sum`, `pps_total`, `pps_index`) and, derived from them, `slow_inputs`: `{"sequences", "volts": [AI0..AI3 block means], "calibrated": [...]}` through the common slow-input calibration (`rpfpga.calibration.slow`; nominal scale until measured), or null when no XADC sequence completed. `STATS?` adds `pps`: `{"edges", "index", "latency_ticks": 2}` or null before the first edge. `STATUS?` includes `pps_edges`, the latest AI words in volts (`ai_v`, `ai_calibrated`) and the capture FIFO's `capture_fifo_high_water_beats`/`capture_fifo_beats`.

Continuous summary readout (0.2)

`STATS?` is independent of the destructive block FIFO. `STATS DRAIN` consumes only whole records; `dropped_blocks` counts new records rejected by the full FPGA FIFO, modulo 2^32 since APPLY. The FIFO holds 64 records, about 0.537 s. Its first drained records may be old backlog: first-index/epoch fields, not wall-clock arrival, define their times. Configuration changes flush old-context records before new ones appear.

`STATS STREAM` requires the common service to have started the design's worker. Its first line is `OK {"format": "ajku-statistics-v1", "session":..., ...}`. Subsequent lines are `STATS {"lost_blocks":N,"batch": {...}}`. The batch contains `blocks` (each with `raw`, optional `physical`, `slow_inputs` and `gap_before`), calibration/ranges, revision/session, block slots, sample rate, `dropped_blocks`, and `fifo_loss_changed`. `lost_blocks` is the cumulative count of blocks dropped from this subscriber's 64-batch queue. It is independent of hardware drops. Per-block gap flags use exact epoch/sample-index discontinuities; any `raw.invalid` count must also be respected.

The worker drains only while subscribers exist. The service refuses `STATS DRAIN` while a stream is active; direct board processes must likewise use one drainer. Reload or worker failure ends the stream with an explicit error. The generator reads immutable JSON published by the worker and performs no MMIO. Client close ends the subscription. `tools/record_statistics.py` preserves complete packets as JSON lines; it can run alongside the separate raw-burst recorder.

`HIST?` claims, reads and releases a complete snapshot in one bounded hardware transaction. Before the first complete epoch it returns BUSY. The three-bank engine continues counting while the selected bank is read. Results include `joint` (flattened I-major 64×64 counts), `i` and `q` (1,024 counts each), `count`, `invalid`, `slots`, `epoch`, `first_index`, `missed_epochs`, and `gap_slots`, and the out-of-window counts `joint_outside`, `i_below`, `i_above`, `q_below`, `q_above`. The driver verifies that each histogram plus its outside counts equals `count` and that valid+invalid equals slots. The epoch's own window (recorded with it by the FPGA) is in `window` (`centers`, and per `joint`/`marginal`: `low`, `high`, `codes_per_bin`, `shift`, `volts`, `full_span`), with the shortcuts `joint_window_codes`/`marginal_window_codes` (half-open `[low, high)` native codes per axis) and `voltage_range`/`marginal_voltage_range` (the same edges in connector volts from the captured board gain/offset; null for the counter source).

These are counts, not normalized probability densities. (0.2's `native_range` field is replaced by these window fields.)

Missed epochs include replacement of an unread/unclaimed snapshot and discarding a new epoch while another is claimed. Startup/bank-unavailable slots are counted in `gap_slots`. Claims cannot create unreported acquisition dead time, but finite retention can discard complete epochs. The 6,144-word MMIO read uses the optional common native repeated-read path inside the same checked transaction. `STATUS?` includes `read_backend` (`native` or `python`); direct drivers opt in with `prepare_native_reads()` outside a transaction. Measured latency and the tested contention envelope are in VERIFICATION.md.

Raw capture metadata

Capture metadata contains the UUID, FPGA session, configuration revision and source, build ID, driver/common-software versions, read backend, full board calibration and jumper ranges, nominal sample rate, representation, pair count, first sample index, accepted/acknowledged counts, error/completion flags, selection method and host request/completion timestamps. Host timestamps are explicitly **not UTC-aligned ADC timestamps**. Add optical-chain and reference-condition metadata when available; absence is not a calibration. New captures also record `dac_stream_codes`: the two actual signed 16-bit converter words when armed, including calibrated-zero or a held calibration reference. These additive fields do not change the `.iqbin` payload or FPGA ABI. From 0.3, metadata also records `timed_start`, `requested_start_index` (null unless timed; then equal to `first_index`), `pps_before` (the latest PPS edge before arming, or null), `buffer_pairs` (the kept DMA allocation's size) and `capture_fifo_high_water_beats` of 2048 `capture_fifo_beats`. With common software 1.3, `copy_backend` identifies the immutable-buffer snapshot path (`native` or `python`), and `board_copy_s` records its elapsed duration after DMA completion/cache invalidation. These are performance provenance, not physical sample timestamps or new hardware-completion guarantees.

The payload is `[I0,Q0,I1,Q1,...]`, each native right-aligned signed 16-bit little-endian integer in −8192...8191. Unlike the vendor stream, the low two bits are real ADC information, not zero padding. No whitening, rejection or bit truncation.

`.iqbin` and HDF5

`control/ajku/recording.py` writes `IQ01` (four bytes), the UTF-8 JSON length (u32 little endian), the JSON, then the exact raw payload. It refuses existing files, invalid acquisitions, trailing bytes and truncated payloads when reading. The metadata limit is 1 MiB; the prototype raw limit is 64 MiB. Failed file writes can leave an incomplete file, which the reader rejects.

Host-only `iqanalysis.export_hdf5` saves the same metadata as `metadata_json` and native codes as the Nx2 `native_codes` dataset. h5py is optional and never imported by the board plugin. Host spectra and moments are analysis products, not a claim that the missing intervals between captures were measured in detail.

MS1 host products

`tools/iqmeasure.py` adds `ajku-measurement-v1` NPZ products, `ajku-reference-v1` reference directories and `ajku-run-v1` run manifests. These are host-only formats; the service still transfers native `.iqbin` data and summary packets. NPZ loading disables pickle. References include raw and processed file hashes and require matching recorded processing, source/configuration, ranges, board calibration, converter outputs and declared chain metadata before comparison. Signed excess noise is preserved; neither physical conditions nor quantum origin are inferred.

A run records raw/processed file paths and hashes, exact sample intervals, statistics packets, histogram snapshots and events. Its states are `recording`, `stopping`, `complete` and `complete_with_errors`; already admitted jobs finish after stop. Raw coverage unions retained sample intervals per session and divides by measured host elapsed time. Imported data does not count toward live coverage. See MS1.md for usage and the distinction between continuous summaries and intermittent raw recordings.

Processing revision 2 adds a `temporal` processing contract (null when disabled), full third/fourth cumulants and optional `temporal_*` NPZ arrays. Correlation arrays are indexed by signed `temporal_lag_samples/temporal_lag_s`; `temporal_lag_pairs` gives actual within-block pairs. Amplitude/power matrices have shape `(lags,2,2)`; complex and radial products have shape `(lags,)`. Batch arrays have shape `(scales,5)` in the saved `uncertainty.statistics` order; lengths/counts accompany each scale. Metadata records centring, pair policy, omitted

pairs, unused tails and the stationarity/long-batch assumptions behind estimated standard errors. NaNs stay in arrays for undefined normalized values; JSON metadata remains finite.

Processing revision 3 adds `processing.filtering` (null or an `ajku-fir-v1` descriptor), `document.filtered`, exact `fir_coefficients`/response arrays and `filtered_*` products. Original unprefixed arrays, native `.iqbin` data and FPGA summary packets remain raw. The descriptor hashes its specification, coefficients, rate, response check and edge/phase policy; loaders verify these plus derived coordinates and frequency/lag geometry. `ajku-filter-profile-v1` JSON additionally contains the coefficient list for exact replay. Filtered reference directories carry `filter.json` and its hash in the reference manifest.

Filtered geometry uses **original sample-counter coordinates**: output first/last centre, integer stride, valid count, omitted edge centres and unused input tails. It never repurposes the raw capture's `pairs`, `first_index` or sample rate. The run index adds filter ID and derived count/first/stride without changing raw coverage. `compare()` retains its raw results and adds a `filtered` comparison.

Optional `iqmeasure.export_hdf5` writes `ajku-measurement-hdf5-v1`: complete measurement JSON, original `native_codes`, every NPZ product under `products/`, and full `filtered_samples` with rate, uint64 first-centre index, stride and geometry attributes. A root `complete` marker is set only after all datasets have been written/flushed; consumers must check it. The `.iqbin`/service wire formats and FPGA ABI do not change. See FILTERING.md for formulas and usage.

Processing revision 5 adds `processing.skew_correction` (`requested`, `applied`, `delay_s`, `measured_delay_s`, `source`, `scope`). When the board calibration holds measured `delay_s` for both inputs at their jumper settings (common loopback delay), their difference is Q's skew relative to I, and every raw/FIR branch adds `skew_*` arrays: `skew_cross_psd`, `skew_spectral_matrix`, `skew_cross_phase_rad`, `skew_complex_psd`, `skew_complementary_psd`, `skew_band_covariance_matrix` (`skew_band_covariance` in the summary). Unprefixed products stay uncorrected. Time- domain products, moments and histograms are never corrected. Loaders check that these arrays exist exactly when `applied` and agree with the uncorrected magnitudes.

Processing revision 4 adds Hermitian `spectral_matrix`, eigenvalue and conditional coherence/phase products, `band_covariance_matrix`, and optional `diagnostic_*` arrays/metadata (also under `filtered_*`). Diagnostic contracts record binning, sigma thresholds, batch sizes and estimator versions; out-of-view counts, subset sizes, bias/SE arrays and applicability limits remain explicit.

`tools/iqdrift.py` and `tools/iqstability.py` also accept the higher-order metrics `skewness_i/q`, `excess_kurtosis_i/q`, `kappa22_normalized` and `radial_kurtosis_ratio` (None for 0.2/0.3 records, which lack the sums); drift windows aggregate the exact sums.

`tools/iqstability.py` writes `ajku-stability-v1` JSON: per metric the overlapping Allan and modified deviations (`tau_s`, `dev`, `lo`, `hi`, `n`) of the block series on its longest gap-free segment, with units, block interval, segment span, context and the number of segments (lost blocks, invalid slots and context changes split the series).

LO-level/sweep manifests (`ajku-lo-level-v1`, `ajku-lo-sweep-v1`) link complete raw and processed records to declared independent power readings and their provenance. `ajku-lo-calibration-v1` NPZ reports store signed polynomial fits, observed/model clearance, statistical and meter-error propagation and optional conditional Q/C. Unknown powers are ineligible and synthetic data require explicit opt-in.

Monitored runs additionally store `drift.jsonl` (`ajku-drift-window-v1` windows) and context/gap/baseline/trigger/outcome events. Triggered records use `selection="drift_followup"`; capture metadata includes cause indices, actual delay, late/overlap flags and `pretrigger_raw_pairs=0`. The run index retains selection and trigger metadata. Automatic jobs are bounded to one, and busy/stale/failed/skipped outcomes are distinct. Stopping retains already admitted work before sealing the run. These are host follow-ups, not hardware pretrigger records. See PRIORITIES.md for the complete workflow contract.

Register map — prototype ABI 0x00040000

Common sysid: 0x40000000, design `OIQS`, version 0.5. Design AXI-Lite block: **0x43C00000**, 4 KiB. All offsets below are bytes; words are little-endian u32. Use the checked driver rather than accessing an unconfigured FPGA.

Offset	Access	Meaning
0x000	R	Register ABI
0x004	W	Pulse command: bit 0 START, 1 ABORT, 2 SNAPSHOT (latch the sample index), 3 APPLY, 4 POP_STATS, 5 CLAIM_HIST, 6 RELEASE_HIST, 7 START_TIMED, 8 CLAIM_SPEC, 9 RELEASE_SPEC

Offset	Access	Meaning
0x004	R	bit 0 busy, bit 1 done; error mask in bits 15:8
0x008	RW	Next capture physical base, reset 0
0x00c	RW	Next capture pair count, reset 2^21
0x010	RW	Shadow source: bit 0 ADC=0/counter=1; reset 0
0x014/018	RW	Shadow OUT1/OUT2 signed 16-bit stream code, sign-extended writes; reset 0
0x020	R	Accepted sample pairs in this capture
0x024	R	Pairs in bursts whose AXI responses arrived; inspect errors too
0x028/02c	R	Capture first sample index, low/high
0x030/034	R	Sample-clock index latched by SNAPSHOT, low/high
0x038	R	Statistics block length (2^20 in the real design)
0x03c	R	Active capture's latched physical base, for ownership-aware cleanup
0x040	R	Statistics FIFO: complete blocks queued (bits 15:0), reflecting a POP accepted at the same clock
0x044	R	Dropped new statistics blocks since APPLY, modulo 2^32
0x048	R	Statistics FIFO capacity (64)
0x050	R	Histogram bit 0 ready, bit 1 claimed, bit 2 accumulating
0x054	R	Missed histogram snapshots since APPLY, modulo 2^32
0x058/05c	R	Histogram uncovered clock slots since APPLY, u64
0x060	RW	Histogram read index, 0…6143
0x064	R	Claimed histogram word; increments index; synchronous memory read
0x068/06c	R	Joint side (64) / marginal bins (1024)
0x070	R	Histogram epoch clock slots (2^24)
0x080–0bc	R	Frozen histogram metadata (16 words)
0x0c0	RW	Shadow histogram centres: bits 13:0 I, 29:16 Q (signed native codes); reset 0
0x0c4	RW	Shadow histogram bin widths (log2 codes per bin): 3:0 joint I, 7:4 joint Q, 11:8 marginal I, 15:12 marginal Q; reset 0x4488 (full span)
0x0c8/0cc	RW	Timed-capture start sample index, low/high
0x0d0	R	Capture FIFO high-water mark (64-bit beats) since the last START
0x0d4	R	Capture FIFO capacity (2048 beats, 16 KiB)
0x0d8	R	PPS rising edges since loading, modulo 2^32
0x0e0/0e4	R	Latest slow-input words: AI1:AI0, AI3:AI2 (16-bit XADC words)
0x100–1fc	R	Latest complete statistics block (55 words; 0x1dc–0x1fc read zero), read in place
0x200–2fc	R	Oldest queued statistics block (55 words of a 64-word slot; the rest read zero), read in place; synchronous memory read
0x300	R	Spectra: bit 0 ready, bit 1 claimed, bit 2 accumulating
0x304	R	Missed spectral epochs since APPLY, modulo 2^32
0x308	RW	Spectral read index, 0…4095
0x30c	R	Claimed spectral word; increments the index; synchronous memory read
0x310	R	Spectral frame length N (1024)
0x314	R	Frames per spectral epoch (2^14)
0x318	R	Words per spectral bank (4096)
0x340–37c	R	Frozen spectral metadata (8 words; 0x360–0x37c read zero)

Error bits: 0 invalid base/length, 1 aborted, 2 FIFO overflow, 3 invalid ADC sample, 4 AXI response code/ID error, 5 timed start index not at least two samples after the START clock. A failed acquisition can acknowledge a prefix of the requested record; it remains invalid. START while busy is ignored. APPLY commits source, DAC and histogram-window shadows, clears partial statistics, invalidates histograms and spectra and aborts an active capture. SNAPSHOT only latches the sample-clock index; it does not touch acquisition.

The AXI-Lite slave accepts address and data independently and honors byte strobes. Unsupported writes return SLVERR; unsupported reads return zero. Read responses and write responses remain stable until accepted. Reset/programming must not be used as a substitute for orderly DMA drain during normal operation.

Statistics record

The same 55-word record is the latest block (0x100) and each FIFO slot (0x200).

Word	Meaning
0	Completed-block epoch, modulo 2^32; continues across APPLY
1	Valid pair count
2–3	First clock-slot index, u64
4–5	Sum I, sign-extended i64

Word	Meaning
6–7	Sum Q, sign-extended i64
8–9	Sum I^2, i64 (nonnegative)
10–11	Sum IQ, i64
12–13	Sum Q^2, i64 (nonnegative)
14–17	Minimum I, maximum I, minimum Q, maximum Q; signed i32
18–19	I/Q rail counts (codes −8192 or +8191)
20–21	I/Q near-rail counts (code < −8064 or >= +8064)
22	Invalid slot count
23	Source: ADC=0/counter=1
24	Slow-input (XADC) sequences completed in the block
25–28	Sums of the AI0, AI1, AI2, AI3 16-bit words over those sequences
29	PPS rising edges since loading, modulo 2^32, at the block's end
30–31	Sample index recorded for the latest PPS edge, u64 (0 before the first)
32–39	Sum I^3, sum I^2Q, sum IQ^2, sum Q^3: each i64 (sign-extended 60-bit sums)
40–54	Sum I^4, I^3Q, I^2Q^2, IQ^3, Q^4: each three words, i96 little-endian (sign-extended 74-bit sums)

All sums are exact for any input: 34-bit first, 48-bit second, 60-bit third and 74-bit fourth-order sums of 2^20 native codes (a constant −8192 gives −2^59 and 2^72). The products come from a six-stage pipeline (squares, then the nine third/fourth-order monomials, 28 DSP48 slices); every slot's values, valid flag, index and slow-input/PPS data travel together, and slots in flight at an APPLY are dropped, so a block still starts with the first slot after it.

Valid count + invalid count must equal the block length. After reset/APPLY the latest record is all zero until a block completes. For an all-invalid block, minimum/maximum sentinel values are +8191/−8192 and physical results are absent.

The latest record changes in one clock at a block end. Read it as a burst, then read word 0 again: equal epochs prove one consistent block (the driver retries otherwise). There is no copy to latch.

The PPS input (DIO0_P) passes a two-flop synchronizer; the recorded index is 2 sample clocks after the clock edge that first saw the pin high (±1 clock for an asynchronous edge), plus the pin's input delay.

Buffered statistics and histograms

Read the oldest queued record at 0x200…0x2d8 while 0x040 is nonzero, then POP_STATS retires it. Each word read is a synchronous RAM read (the response two clocks later: the RAM output and, since 1.0, one more register before the AXI read data); reading does not change the queue. Reading an empty queue returns zeros. New blocks are discarded when full. APPLY resets the queue and its drop counter, without changing the statistics epoch sequence.

CLAIM_HIST freezes the current latest snapshot if one exists. Check both ready and claimed before reading metadata or counters. RELEASE_HIST consumes/releases that bank. Unclaimed snapshots may change at every epoch; a claim is essential for a consistent copy. Keep the entire claim/read/release under one framework transaction to serialize software readers and configuration changes. APPLY invalidates all histograms and clears the loss/gap counters; the histogram epoch sequence continues.

Histogram metadata words: 0 epoch, 1 valid count, 2–3 first slot index (u64), 4 invalid count, 5 source (ADC=0/counter=1), 6 epoch slots, 7 valid samples outside the joint window, 8/9 I below/above its marginal window, 10/11 Q below/above, 12 the epoch's bin widths (as 0x0c4, limited to the full span), 13 its centres (I in 15:0, Q in 31:16, sign-extended), 14–15 zero. Counter indices 0…4095 are joint bins in I-major/Q-minor order; 4096…5119 are I marginals; 5120…6143 are Q marginals. A data read advances the index once when its AXI read address is accepted, not when the read response is acknowledged. Out-of-range or unclaimed reads return zero. The 32-bit pointer saturates at 0xffffffff.

Histogram windows

Per axis a, with centre c (0x0c0) and bin-width exponent s (0x0c4), a window of B bins (64 joint, 1024 marginal) has the low edge `low = c − 2^(log2 B − 1 + s)` and covers native codes `low ≤ x < low + B·2^s`, bin `(x − low) >> s`. The joint sample is inside only if both axes are. s = 8 (joint) and s = 4 (marginal) with c = 0 are the full native span. Valid count = sum of joint bins + joint outside = sum of a marginal + its below + above counts, for every epoch.

Timed captures

Write the start index to 0x0c8/0x0cc, base and pair count as usual, then START_TIMED. The writer is busy (no AXI traffic) until the sample whose index equals the start index, records from exactly that sample, and reports it as the first index. The index must be at least two samples after the clock that accepts START_TIMED, else error 5 (no writes, done). ABORT or APPLY while waiting ends the capture without writes.

Continuous spectra

The engine windows z = I + iQ with a Hann table (Q17, $w_n = \sin^2(\pi(n + ½)/N)$, symmetric, at most $2^{17} - 1$), transforms every non-overlapping 1024-sample frame (radix-2^2 FFT, 18-bit input with four fractional code bits, 25-bit output in whole codes) and accumulates per bin k the power $P_k = |Z_k|^2$ ($0 \le k < 1024$) and the complementary product $C_k = Z_k Z_{(1024-k)}$ ($0 \le k < 512$; $C_0 = Z_0^2$, the Nyquist bin's product is not formed). `ajku.model.spectral_epoch` reproduces every sum exactly; THEORY gives the arithmetic.

Frames follow the statistics blocks (1024 frames per block). An epoch is 2^{14} frames (2^{24} slots): the histogram epoch, starting at the second block boundary after a reset or APPLY. A frame with an invalid slot is skipped and counted. Two banks: one accumulates, the other holds the last published epoch. CLAIM_SPEC freezes it (if ready), RELEASE_SPEC consumes it. An unclaimed epoch is replaced by the next one (counted missed); an epoch that ends while the bank is claimed is discarded (counted missed), and acquisition continues. For about 120 ns after an epoch ends nothing is ready (its last writes drain): a refused claim is retried. Keep claim, reads and release in one framework transaction. APPLY invalidates the bank and clears the missed count; the epoch sequence continues.

Metadata words (0x340…): 0 epoch, 1 accumulated frames, 2–3 first slot index (u64, equal to the matching histogram epoch's), 4 skipped frames, 5 source (ADC=0/counter=1), 6 frames per epoch, 7 frame length. Accumulated + skipped = frames per epoch.

Bank words (index at 0x308, data at 0x30c): 64-bit two's-complement sums, low word first. Words 0…2047: P_0…P_1023; 2048…3071: Re C_0…C_511; 3072…4095: Im C_0…C_511. A data read advances the index when its address is accepted. Out-of-range or unclaimed reads return zero; the index saturates at 0xffffffff.

Verification record

Newest results first. Planned targets are not measurements. Release 1.0 passed the digital board qualification below (2026-11-03; 0.2 and the digital checks of 0.3–0.5 before it); PPS/slow-input accuracy is untested; analogue response, anti-alias rejection and optical calibration remain unmeasured (an end-user task).

2026-11-03 — release 1.0, build 0xfb4e4824

The user declared the feature set complete and asked for version 1.0: rebuild and re-verify, a soak test, the documentation and the project page. FPGA/driver **1.0**, register ABI **0x00040000** (unchanged), host processing revision **5**, common **1.3**, Vivado 2025.1. The tested software is commit `a0a2eca` plus the uncommitted 1.0 changes (version strings, the read register below, the common calibration cache below, documentation); every installed driver/framework file was hashed against that working tree before and after the board sequence.

Build and timing

The version change alone (the sysid word) moved placement enough to miss the +0.25 ns setup target. Three implementation settings were tried before an RTL change:

Build (all 1.0 sources but the last)	WNS / WHS (ns)	Worst path
default settings (as 0.5)	+0.177 / +0.017	`iq2` product to a DSP input, routing only
placement ExtraTimingOpt + post-place AggressiveExplore	**−0.026** / +0.020	statistics-FIFO BRAM to `s_axi_rdata`, 10 LUT levels
post-place AggressiveExplore	+0.140 / +0.017	the same FIFO path, 11 LUT levels
register the RAM read word, default settings: 0xfb4e4824	**+0.275 / +0.027**	capture writer's burst counters (as 0.5)

The FIFO path appeared among the ten worst in all three builds: its data is the BRAM latch, which synthesis placed at the head of the whole read multiplexer. `iq_core` now registers the selected RAM word (FIFO, histogram or spectral) once more before the AXI read register; such reads take one more clock. The implementation settings are back to 0.5's (post-route AggressiveExplore only).

Build	LUTs	FFs	Slices	BRAM36	DSP	WNS / WHS (ns)
0.5, 0x07bd2ef4	9,440	10,723	3,513	40	57	+0.280 / +0.039
1.0, 0xfb4e4824	**9,421**	**10,333**	**3,400**	**40**	**57**	**+0.275 / +0.027**
caps / device	10,000 / 17,600	16,000 / 35,200	3,300 / 4,400	48 / 60	48 / 80	≥ +0.25

The 390 fewer flip-flops are in the spectral engine (3,467 → 3,083, its RTL unchanged; more of its delay registers packed into SRLs), compared per module from the placed checkpoints. Setup/hold TNS 0, no critical warnings, no unconstrained internal endpoints, the same 18 outputs without output delays and the `pps_i` false path; routed DRC: the same DSP pipelining advisories (30 DPIP-1, 11 DPOP-1, 18 DPOP-2) and the SmartConnect one. DSPs and slices exceed the project caps as accepted for 0.5.

- `make sim`: **nine benches passed** (`tb_core` reads through the added register).
- `make test`: **132 ajku tests** passed (3 skipped, Python 3.14.7); common suite **324 tests** (20 skips, Python 3.12.9), with the site tests once the page was complete.
- `python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed**; fresh and restored builds both **0xfb4e4824**, the fresh one again at WNS +0.275 ns. Evidence: `.build/verification-pf167u54/`.

Board qualification

On 192.168.0.102 (PYNQ 3.0.1), after `make install-service`: the stages of the 0.2 release baseline, plus the spectra and the 0.3 features, run by one sequence with host `/proc` and read-only board OS samples every 30 s, then a real-browser notebook check. The existing board calibration was retained; no analogue or optical characterization.

First attempt: failed. `make test-hw` and the spectra, features, 32 MiB transport and priorities stages passed, but the one-hour workload exceeded its control envelope: p99 223.6 ms (≤ 250) but one of 18,381 probes took **515.6 ms** (≤ 500); 1,682 captures, 429,495 blocks and 8,366 histograms were retained with no loss. The sequence stopped there. Control latency had grown since 0.2 (p99/max 181/388 ms): an idle `STATUS?` took 46 ms against 6 ms for `INDEX?`, because one status merged the board calibration file onto its defaults seven times (`SNAPSHOT?`, the probe, nine times), each ~5 ms on the board's CPU (the merge deep-copies at every level). Common `Board.calibration()` now caches the merged result by the file's stat stamp (`files.read_cached`, which already guarded the session and safety files) and returns a fresh copy per call; `ranges()` reads its file through the same cache. A new common test covers copies, a same-size rewrite and deletion. Idle `STATUS?` fell to 19 ms and `SNAPSHOT?` to 26 ms. The first attempt's receipts are kept as `board-1.0-attempt1-*`.

Second attempt (after the common change): passed, every stage:

- Source preflight: all **31** installed driver/framework files equal the tested working tree; again at the end. Source, final.
- `make test-hw`: every check, including the exact counter spectra (1,024 powers and 512 complementary sums of a 16,384-frame epoch) and an ADC epoch's Parseval power (ratio 1.00010). Receipt.
- `qualify_spectra.py`: an exact counter epoch (2,048 sums); eight consecutive ADC epochs, none skipped or missed, on the 2^24 grid; Parseval 1.00003; FPGA/host band PSD **1.0042 / 1.0006**; `SPEC?` 111–141 ms. Receipt.
- `qualify_features.py`: exact timed start, refused past start, 105 streamed blocks with slow inputs, conserved narrow window, 64 MiB allocated and reused. Receipt.
- `qualify.py --seconds 10 --maximum-pairs 8388608`: captures to 32 MiB; **1,200/1,204/1,243** blocks in the three phases, zero FIFO drops; the timed cancellation aborted at the first attempt (4,594,915 of 8,388,608 pairs); one ADC block equal to the samples in all fourteen sums. Receipt.
- `qualify_priorities.py --seconds 15`: 1,791 exact blocks, the recorded trigger decisions and windows replayed exactly, one follow-up (141.8 ms) and one busy skip, stale-revision rejection, unknown power refused for fitting. Receipt.
- `qualify_ms1.py --seconds 3600 --dashboard --pipeline --lags 128 --filter --diagnostics --diagnostic-bins 32` and the same for 120 s with 32 MiB captures:

Quantity	One hour, 8 MiB captures	120 s, 32 MiB captures
Captures / elapsed recording	**1,657 / 3,603.473 s**	14 / 139.902 s
Verified statistics blocks / histograms	**429,551 / 8,820**	16,675 / 385
FIFO loss / subscriber loss / invalid slots	**0 / 0 / 0**	0 / 0 / 0
Raw coverage	0.771 %	0.672 %
Independent SNAPSHOT queries	23,702	—
Control median / p99 / maximum	**43.3 / 141.2 / 334.0 ms**	— / 138.9 / 245.1 ms
Capture/download median, analysis median	0.702 s, 1.745 s	2.383 s, 9.327 s
Board immutable-copy median	65.5 ms, native	—
Matched-reference captures	1,654 (+3 before the reference)	—

Maximum FIR trace discrepancy 1.05e−11 native codes; the ADC run (3 captures, 1,171 blocks, no loss) and the headless-widget phase completed. One hour, 32 MiB.

- `qualify_readers.py --stall-seconds 60`: the stalled subscriber lost 6,760 blocks while the fast one kept 7,197 with no loss and the FPGA FIFO dropped none; control p99/max 48.6/83.1 ms; disconnect after the first 4096 bytes of a completed 8 MiB transfer, then a new acquisition. Receipt.
- Resources (one-hour counter phase, 106 steady samples after a five-minute warm-up): board service RSS 72.9–81.2 MiB, first/last-quarter medians 73.5/74.7 MiB; host analysis RSS 161.6–292.9 MiB (193.9/167.1); board free CMA 46.0–46.5 MiB (the driver keeps its 32 MiB DMA allocation since 0.3; 0.2 showed 82–95 MiB). Sampled finite-run ranges, not bounds. Sequence.
- Real browser (Notebook, headless Chromium 153 through Playwright 1.48; the harness of 0.2 with the 0.3 button name ”Apply edits”, processing revision 5 and the FPGA spectrum summary): source staging/discard; two ADC captures (1,996 blocks, 10 histograms) with raw/FIR reference save and comparison; the three displayed views equal to their expected renders (PNG SHA-256); ”Spectrum epoch …: 16384 frames, 0 skipped” in the live summary; a monitored run (443 blocks, 28 trend windows, one follow-up after 470.3 ms, no pretrigger pairs); a labelled synthetic fit; normal cleanup. The frontend’s startup `schema` TypeError of 0.2 recurred. The first try stopped at the old button name before any acquisition (receipt). Receipt.

The board was left running 1.0 on the ADC source with full-span histograms. This is digital qualification of a finite duration, not optical, analogue or endurance qualification.

2026-11-03 — 0.5: continuous spectral accumulation, build 0x07bd2ef4

After a resource estimate (recorded in `wiki/fpga-resource-budget.md`) the user chose ”the smallest that still fits (N=1024, 16-bit FFT)”: continuous non-overlapping 1024-point Hann spectra of z = I + iQ, summed per 2^24-slot epoch. FPGA/driver **0.5**, register ABI **0x00040000**. This changes the approved ”no FPGA FFT in v1” boundary and exceeds the DSP and slice caps by the user’s choice. Instead of the Xilinx FFT core the engine uses a custom radix-2^2 single-delay-feedback FFT, so the repository’s bit-exact Python model stays the reference and `make sim` needs no IP libraries. Measured out of context (xc7z010-1, routed at 125 MHz):

FFT, N = 1024	Slices	LUTs	FFs	BRAM36	DSP	WNS
Xilinx FFT 9.1, 16-bit (estimate basis, placed)	1,108	2,563	5,125	4	13	+3.78 ns (placed)
custom, first version	549	1,711	1,153	4.5	20	+0.568 ns
custom, kept (BRAM delay lines through the output register, DSP post-adds)	**506**	**1,573**	**1,445**	**4.5**	**12**	**+1.928 ns**

- `make sim`: **nine benches passed**. `tb_fft` compares every output of the FFT for N = 16, 64 and 1024 with `model.spectral_fft` (random full-scale frames, a constant and an alternating extreme frame, back to back after a pre-roll frame). `tb_spectra` (N = 16, 16 frames per epoch) compares five published banks word by word with `model.spectral_epoch` and checks an invalid slot (15 frames, 1 skipped), latest-wins replacement (missed 1), an epoch ending while the bank is claimed (discarded, missed 2, claimed data intact), the next epoch written fresh into the bank the discarded one used, and a clear mid-epoch (nothing ready; the next epoch starts at the second block boundary). `tb_core` checks the ABI, the spectral geometry registers, a claim retry, the metadata, bank words of a constant input against the model, index readback, out-of-range and unclaimed reads (zero), release and APPLY invalidation. A deliberately corrupted expected word made `tb_spectra` fail.
- `make test`: **132 ajku tests** passed (3 skipped, Python 3.14.7). New: the FFT against exact DFTs (N = 16, 64), extreme frames inside the RTL widths, the window table, frame products, epochs and bank words,

the XSim vectors being current; with NumPy the fixed-point noise (0.3 % of the noise power at 1.3 codes rms), a tone's levels, cross-phase ($-2\pi f\tau$ for a delayed Q) and complex-spectrum ratio against `iqanalysis.spectra(nperseg=1024, overlap=0, detrend=False)`, white-noise normalization against NumPy, and volt scaling; driver tests through the fake core (exact sums, claim retry, release after a failed read, frame accounting, counter source), `SPEC?` over the network, and the dashboard drawing FPGA spectra. Common suite: **323 tests, OK, 20 skips** (Python 3.12.9).

- Build (Vivado 2025.1):

Build	LUTs	FFs	Slices	BRAM36	DSP	WNS / WHS (ns)
0.4, 0x7322d68c	6,752	7,231	2,568	26	31	+0.351 / +0.034
0.5, 0x07bd2ef4	**9,440**	**10,723**	**3,513**	**40**	**57**	**+0.280 / +0.039**
caps / device	10,000 / 17,600	16,000 / 35,200	3,300 / 4,400	48 / 60	48 / 80	≥ +0.25

In context the spectral engine takes 2,561 LUTs, 3,467 FFs, 12 RAMB36 + 4 RAMB18 and 26 DSPs (FFT 12, products 12, window 2). The estimate (~10,050 LUTs, ~13,470 FFs, 39.5 BRAM, 58 DSPs, 3,050–4,000 slices with the Xilinx core) was close on everything but the FFs. LUTs (560 below the cap), FFs, BRAM and timing pass the caps; DSPs (57) and slices (3,513, 80 % of the device) exceed them, as accepted. The worst path is unchanged in kind: the capture writer's burst counters (+0.280 ns; the post-route pass found nothing to improve). Setup/hold TNS 0, no critical warnings, no unconstrained internal endpoints, the same 18 outputs without output delays; routed DRC: DSP pipelining advisories (30 DPIP-1, 11 DPOP-1, 18 DPOP-2) and the SmartConnect no-routable-loads one.

- `python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed**; fresh and restored source-only builds both **0x07bd2ef4**, the exported overlay. Evidence: `.build/verification-f4g7srb_/`.

Board: digital checks of 0.5

The user then asked for the board test. `make install-service` installed 0.5 (build **0x07bd2ef4**, reused) and restarted the service on 192.168.0.102 (common 1.3).

- `make test-hw` (counter): **every check passed**: the 0.4 checks unchanged (exact captures to 8 Mi pairs, the kept 32 MiB allocation, the exact timed start, 474–475 XADC sequences per block, all fourteen sums of every drained block, histograms with readout 25.2 ms, the held bank, the zoomed windows), and new: **every sum of a counter spectral epoch** (16,384 frames; 1,024 powers, 512 complementary products) equal to the bit-exact model (board-local readout 65 ms), spectral and histogram epochs aligned, and an ADC epoch's Parseval power within 0.06 % of the statistics' mean $I^2 + Q^2$. Receipt.
- `tools/qualify_spectra.py` (new): an exact counter epoch (2,048 sums) through `SPEC?`; **eight consecutive ADC epochs** of 16,384 frames, none skipped or missed, first indices exactly 2^24 slots apart and on the histogram epochs' grid; Parseval ratio 1.00033; the FPGA/host mean PSD over 1–25 MHz **1.0040 (I) / 1.0006 (Q)** against Welch spectra of a 2 Mi-pair capture at the same resolution (2,048 segments; tolerance 2 %); `SPEC?` 112–117 ms over the network, conversion included. Receipt.
- `tools/qualify.py --seconds 10 --maximum-pairs 8388608`: the first two runs **failed the cancellation case** ("completed before cancellation", busy never observed; receipts 1, 2). The case polled `STATUS?` until it saw the 32 MiB DMA busy, then applied. With a statistics stream active `STATUS?` took 70–86 ms (45 ms without), the transfer lasts 67 ms, and an apply reaches the hardware 40–60 ms after it is sent: the poll missed the transfer in 4 of 6 trials under the stream and in none of 6 without it. Applies sent 0–100 ms after a capture request always cancelled it, so cancellation itself was not at fault; the 0.4 pass under the same load was a favourable race, and "busy seen, request cancelled" could not tell a DMA abort from a cancelled copy after the transfer. The case now schedules a hardware-timed capture, applies once `INDEX?` (6 ms) reaches the start and proves the abort from the capture's own record (abort flag, 0 < accepted < pairs); an apply that arrives after the transfer is retried (up to six attempts, all recorded). Trials: 5/5 aborts at the first attempt without the stream, 5/5 within 1–3 attempts with it, stopping at 5.4–8.3 M of 8.39 M pairs. The repeat **passed**: exact counter captures to 32 MiB; **1,205/1,208/1,235** statistics blocks in the three phases with zero FIFO drops; an abort at the first attempt (7,277,328 pairs accepted); one complete actual ADC block whose fourteen FPGA sums equal those of the captured samples. Receipt.
- `tools/qualify_features.py`: as for 0.4 (exact timed start, refusal of a past start, 97 streamed blocks with slow inputs, conserved narrow ADC window, 64 MiB allocated and then reused). Receipt.
- `tools/qualify_ms1.py --seconds 120 --pipeline --lags 128 --filter --diagnostics --diagnostic-bins 32`: **60 captures, 14,695 statistics blocks, 325 histograms, no FIFO/subscriber loss, no invalid slots**; coverage 0.816 %; control p99/max **218/332 ms** (envelope 250/500 ms; 0.4 gave 202/399 ms); analysis median 1.75 s (0.4: 1.59 s). Receipt.

- A live ADC spectrum identifies the uncontrolled IN2 signal noted for 0.3 and 0.4 ($\sigma \approx 0.51$ V): a tone near **10.0 MHz** (strongest bins 9.89/10.01/10.13 MHz, 0.256 V^2 in the band). IN1 carries it 55 dB lower (crosstalk or common pickup, not separated) besides its DC offset, which the Hann window leaks into bin 1. Its source was not identified. With no reader, each unread epoch replaced by the next counts as missed (about 7.5 per second), as for the histograms.

The board was left running 0.5 on the ADC source with full-span histograms.

2026-11-02 — 0.4: continuous third/fourth-order sums, build 0x7322d68c

After the 0.3 board test the user asked for ”A”: exact third- and fourth-order sums in every statistics block. FPGA/driver **0.4**, register ABI **0x00030000**, 55-word records in 64-word FIFO slots. This changes the approved ”second-order FPGA processing first” boundary at the user’s request. After the local checks below the user approved the board test: the digital board checks passed (below).

- `make sim`: **seven benches passed**. `tb_stats` now checks all nine third/fourth-order sums against 80-bit references over 20 blocks of mixed rails, small codes and invalid slots (with the slow-input and PPS fields), and a full 2^20-slot block at I = −8192, Q = +8191 ($\Sigma I^3 = -2^{59}$ and $\Sigma I^4 = 2^{72}$, the extreme values of the 60/74-bit sums); `tb_core` reads the latest record’s and a FIFO slot’s higher-order words, the words beyond the 55-word record (zero) and the new ABI; `tb_stats_fifo` the 55-word records, zero tails of 64-word slots, wrap, overflow and reading during writes.
- `make test`: **120 ajku tests** passed (3 skipped under Python 3.14.7, 4 under 3.12.9 for missing optional packages). New: the integer cumulants against exact `Fraction` arithmetic at a large mean offset (central moments, all five fourth-order cumulants, excess kurtosis, $\langle r^4\rangle/\langle r^2\rangle^2$, projections, volt scaling, degenerate and empty blocks), the driver’s `higher_order`, the record round trip, drift metrics and exact window aggregates (and 0.3 records without the sums), stability metrics (block excess kurtosis of 16 Gaussian codes near −6/17), and `qualify.exact_sums` against the model at the rails. Common suite: **323 tests, OK, 20 skips** (Python 3.12.9).
- Builds (Vivado 2025.1):

Build	LUTs	FFs	Slices	BRAM36	DSP	WNS / WHS (ns)
0.3, 0x733760cb	4,679	5,599	2,039	24	3	+0.474 / +0.031
0.4 first, 0x4b7eedf1 (not kept)	6,710	7,169	2,591	26	31	+0.255 / +0.022
0.4 second (not kept)	6,739	7,194	2,567	26	31	**−0.036** / +0.033
0.4, 0x7322d68c	**6,752**	**7,231**	**2,568**	**26**	**31**	**+0.351 / +0.034**

The first build’s worst path was the AXI read-address compare into the histogram BRAM enables; registering that read moved the limit to the BRAM output through the read-data mux into the AXI read register (−0.036 ns, rejected). The kept build registers the selected histogram word once more (histogram MMIO words now take three extra clocks) and enables post-route `phys_opt_design -directive AggressiveExplore` (`design(impl_properties)`, as feedback). Its worst path is a DSP-to-DSP route inside the fourth-order products (+0.351 ns after routing and the post-route pass); setup/hold TNS 0, no critical warnings, no unconstrained internal endpoints, the same 18 outputs without output delays. Routed DRC adds DSP pipelining advisories (24 DPIP-1, 5 DPOP-1, 8 DPOP-2) to the SmartConnect one. The 28 new DSP slices are the nine products; 181 LUTs are shift registers (the delayed slot data). All caps pass: slice headroom to 3,300 is 732, DSP headroom to 48 is 17.

- `python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed** (with the new `design(impl_properties)`); fresh and restored source-only builds both **0x7322d68c**, the exported overlay. Evidence: `.build/verification-bu6ob44_/`.

Board: digital checks of 0.4

`make install-service` installed 0.4 (build **0x7322d68c**, reused) and restarted the service on 192.168.0.102.

- `make test-hw` (counter): **every check passed**, now including all fourteen sums of every drained block against the model through fourth order; the kept 32 MiB allocation, the exact timed start, 474–475 XADC sequences per block, histograms (readout 25.1 ms), the held bank and the zoomed windows as in 0.3. Receipt.
- `tools/qualify.py --seconds 10 --maximum-pairs 8388608`: exact counter records to 32 MiB; **1,202/1,211/1,266** statistics blocks in the three phases with zero FIFO drops (55-word records); cancellation of a busy 32 MiB DMA; and one complete actual ADC block whose **fourteen FPGA sums (through Σ Q^4) equal those of the captured samples exactly**. Receipt.

- `tools/qualify_features.py`: as for 0.3 (exact timed start, refusal of a past start, 98 streamed blocks with slow inputs, conserved narrow ADC window, 64 MiB allocated and then reused). Receipt.
- 120 s MS1 workload: **65 captures, 14,664 statistics blocks, 322 histograms, no FIFO/subscriber loss, no invalid slots**; coverage 0.886 %; control p99/max **202/399 ms** (envelope 250/500 ms; 0.3 gave 198/299 ms); analysis median 1.59 s. Receipt.
- A live `STATS?` on the uncontrolled ADC inputs took 29 ms including the board-side cumulants; Q's excess kurtosis was **−1.50**, that of a sinusoid, consistent with the large periodic signal seen on IN2 (σ ≈ 0.5 V; its waveform was not measured otherwise). A statistics stream read 156 blocks/s while catching up its backlog. Receipt. The service journal was not readable without root; no stall was visible in the stream or loss counters.

The board was left running 0.4 on the ADC source with full-span histograms.

2026-11-02 — 0.3: windows, timing, slow inputs and RTL tidying, build 0x733760cb

The user asked for the low-cost improvements from the design review: Allan deviations, histogram zoom, converter-skew correction, XADC monitors, a PPS input, timed captures, RTL tidying and a kept DMA allocation. FPGA/driver **0.3**, register ABI **0x00020000**, host processing revision **5**, common **1.3**. Builds used Vivado 2025.1, xc7z010clg400-1. After the offline/build checks below the user asked for the board test: digital board checks passed on **192.168.0.102** (PYNQ 3.0.1); PPS and slow-input accuracy were not tested (nothing wired to DIO0_P or E2).

Board: digital checks of 0.3

`make install-service` installed the 0.3 overlay, plugin and common software and restarted the service with ajku active (build **0x733760cb**, unchanged by the workflow's rebuild).

- `make test-hw` (counter source): the **first run** right after the service restart captured every record exactly, but the 32 MiB reserve allocation was not obtained: each larger capture fell back to its own size and the kept-allocation check failed. A minute later 8/16/32 MiB allocations succeeded (64 MiB failed, as in 0.2), also through the driver; the cause of the transient failure was not established. The **repeat passed every check**: exact records of 32…2^21 pairs into one kept allocation of 2^23 pairs (FIFO high water 15–16 of 2048 beats); a hardware-timed capture beginning exactly at its requested sample; **474–475 XADC sequences per 2^20-slot block** (17.7 µs per sequence); exact counter statistics and full-span histograms (readout 26.7 ms); a 0.75 s held bank with 5 missed epochs and no new gaps; zoomed windows (1-code marginals, 8-code joint bins on the anti-diagonal) with the exact expected bin and out-of-window counts. First run, repeat.
- Thinking through that fallback showed that a failed larger request lost the kept allocation (it was released before the attempt). The driver now allocates the previous size again before raising; a unit test covers it, and on the board a failed 64 MiB request was followed by an 8 MiB capture that still used a kept 32 MiB allocation. Receipt (an ad-hoc predecessor of `tools/qualify_features.py`).
- `tools/qualify.py --seconds 10 --maximum-pairs 8388608` through the service: exact counter records 32 pairs…32 MiB (32 MiB download 29.9 MB/s, end to end 23.2 MB/s); three 10 s phases with **1,206/1,213/1,212** statistics blocks and zero FIFO drops (32-word records); cancellation of an observed busy 32 MiB DMA and a correct next capture; one complete actual ADC block matching all FPGA sums. Receipt.
- `tools/qualify_features.py` (new): a timed 8 MiB ADC capture over TCP began exactly at its index (0.25 s lead); an index already passed was refused before any write; 93 streamed blocks all carried slow inputs (open pins read about 0.014–0.025 V on AI0–AI2 and 1.54 V on AI3, nominal scale: not a calibration) and PPS count 0. The inputs carried an uncontrolled signal (σ 0.51 V), so the ±6σ window rounded up to the full span; a narrow window (8-code joint, 1-code marginal bins) conserved every count, the whole I range falling in 91 single-code bins and most Q samples outside. 64 MiB allocated this time and the next capture reused it. Receipt.
- `tools/qualify_ms1.py --seconds 120 --pipeline --lags 128 --filter --diagnostics --diagnostic-bins 32`: **62 captures, 14,552 statistics blocks, 321 histograms, no FIFO/subscriber loss, no invalid slots**; raw coverage 0.85 %; control p99/max **198/299 ms** (envelope 250/500 ms; 0.2's same workload gave 171/185 ms); analysis median 1.62 s; ADC run and incompatible-reference rejection passed. Receipt.

The board was left running the 0.3 service on the ADC source with full-span histograms.

Simulation and offline tests

- `make sim`: **seven benches passed** (`tb_sysid`, `tb_stats`, `tb_capture`, `tb_core`, `tb_stats_fifo`, `tb_histograms`, new `tb_xadc`). New checks: the 32-word record's XADC sequence count and AI sums and its

PPS fields against a reference over 20 blocks; the FIFO read in place (twice, unchanged), POP, wrap, overflow, a push while serializing, reading during the next write, clear; histogram windows off-centre and zoomed on both axes with every bin and all five out-of-window counts against a reference, the epoch's recorded geometry, rails and zero outside narrow windows; timed starts (exact first sample by the data pattern, the shortest legal lead, a start one sample ahead and one already passed refused with error 5 and no writes, abort while waiting without writes), the FIFO high-water mark (and 64/64 when overflowing); in the core: ABI, new registers, PPS count and index latency (2 clocks) and record fields, the latest record read in place, POP status at B acceptance, decoded histogram index (bounds, readback, saturation), a zoomed window applied by APPLY with its outside/below/above counts and zero bins; `rp_xadc` with the UNISIM XADC model (AI0..AI3 within 2 LSB at strobes, the 60 µs step followed).

- `make test`: **117 ajku tests passed**, 3 skipped (h5py) under Python 3.14.7; the same 117 with the optional host environment (Python 3.12.9, h5py, ipywidgets), 3 skipped (allantools). New: kept/grown/fallback DMA allocation, timed captures (driver, protocol `INDEX?/CAPTURE n AT i`, early/late/far refusals before writes, the hardware's own refusal not published), re-read of a record torn by a completing block, slow inputs/PPS through calibration, histogram windows from connector volts (rounding, centre, register words, derive round trip, validation before writes, window readout and conservation with outside counts), skew-corrected spectra (a 0.3 sample delay removed from the cross phase, magnitudes/coherence unchanged, products only with both delays measured, saved/loaded/validated), Allan deviations (white block means within 25 % of σ/√m, gap/invalid/context/loss segmentation, units, JSON/figure output), the panel's window controls, `record.py --timed`.
- Common suite: **323 tests, OK, 20 skips** (Python 3.12.9). Under Python 3.14.7 the unchanged `test_buffers` refcount assertion fails (`sys.getrefcount` 1 != 2: 3.14 counts borrowed references differently); not caused by these changes.
- `scripts/test_hw.py` gained checks of the kept allocation, an exact timed start, slow-input sequences per block and zoomed counter histograms (anti-diagonal of 8 codes per bin, 1-code marginals, outside counts), whose expected counts were checked against the integer model. It has **not** been run on a board.

Build, resources and timing

Build	LUTs	FFs	Slices	BRAM36	DSP	WNS / WHS (ns)
0.2, 0xd3f1cbc0	4,153	6,837	2,216	24	3	+0.367 / +0.027
0.3 first, 0x52933dd8 (not kept)	4,625	5,489	1,985	24	3	+0.256 / +0.046
0.3 second, 0xb6fa96dd (not kept)	4,664	5,554	2,049	24	3	**+0.180** / +0.022
0.3, 0x733760cb	**4,679**	**5,599**	**2,039**	**24**	**3**	**+0.474 / +0.031**

All builds: one MMCM, setup/hold TNS 0, bitgen without errors or critical warnings. 0x733760cb also uses the XADC and 252 LUTs as distributed RAM: every LUTRAM cell is the histograms' three 512-bit epoch metadata entries (listed from the routed checkpoint, `.build/lutram-0x733760cb.txt`), read through a mux instead of a 512-bit copy. The flip-flop reduction is the removed copies (latest snapshot, FIFO readout, FIFO holding register, histogram metadata), net of the new fields. The first build's worst path was the 64-bit timed-start compare after an adder (fixed by registering start − 1); the second's was the statistics `pipe` enable fanning out to every accumulator, now `max_fanout`; the kept build's worst path is a histogram BRAM read into the AXI read data (+0.474 ns). Routed DRC: the same six DSP input-pipeline, three DSP output-pipeline and one SmartConnect advisories. No unconstrained internal endpoints; 18 outputs without output delays as before, and `pps_i` without an input delay under its false path (asynchronous, synchronized). All caps pass; slice headroom to the 3,300 cap is 1,261 (was 1,084).

`python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed** (fresh recipe build including the optional common `rp_xadc.v` path, unchanged-build reuse, missing-export recovery, vendor HDL content change, design RTL change, failed build blocking upload, cleanup, saved-project regeneration). The fresh and restored source-only builds both have ID **0x733760cb**, matching the exported overlay. Evidence: `.build/verification-jb7slz87/verification_logs/`. The project was recreated from the recipe for the new ports, so `ajku.xpr` and the saved block design changed.

Remaining 0.3 board checks

With wiring: a PPS (a GNSS receiver or GPSDO, 3.3 V) into DIO0_P to check edge counts, the index spacing against 125 000 000 and the latency; known DC voltages on AI0–AI3 against the nominal/common slow-input calibration; a terminated-input or LO-off record with a zoomed window against the captured full-code histogram; longer stalled-reader and endurance runs with the 32-word records.

2026-11-02 — committed baseline qualification and LaTeX/PDF manual

The user approved committed-baseline qualification followed by release-interface documentation and the lab manual. **Throughput is accepted as a measured operating limitation.** Physical characterization is an end-user task for their board and optical chain; no new analogue/optical calibration was performed here. The existing board calibration and jumper-range declarations were identical before and after.

The sequence began on **2026-11-01** and finished after midnight local time. Receipt filenames for that sequence retain its start date; each contains exact Unix times. Acquisition/analysis used commit **60635cb6868a44ea754b5d69b3b90ea161** FPGA/ driver **0.2**, build **0xd3f1cbc0**, processing **4**, common **1.3**, board **192.168.0.102**. All **31** installed driver/framework source files matched the commit before and after the sequence. The subsequent host GUI correction below is identified separately; the installed acquisition service continued to match the commit. Preflight, qualified-board check, final check.

Complete one-hour workload

`qualify_ms1.py --seconds 3600 --pipeline --lags 128 --filter --diagnostics --diagnostic-bins 32 --dashboard` used default 8 MiB captures, 65,536-point Hann spectra, the default 137-tap factor-two FIR, 16,384-pair diagnostic batches and pointwise coherence alpha 0.05. Statistics, histograms, references and independent control probing ran concurrently. Every published capture/product was reopened; all raw counter samples, every retained statistics block and histogram, and the FIR products were checked. Reopening/verification occurs after the recording and is outside its coverage denominator.

Quantity	Measured result
Captures / elapsed recording	**1,561 / 3,603.180 s**
Verified statistics blocks / histogram snapshots	**429,517 / 8,548**
FIFO loss / subscriber loss / invalid slots	**0 / 0 / 0**
Raw retained time / coverage	**26.189 s / 0.727%**
Independent SNAPSHOT queries	**20,663**
Control p99 / maximum	**181.3 / 388.3 ms**
Capture/download median	**0.791 s**
Analysis median	**1.869 s**
Board immutable-copy median	**65.3 ms**, native backend
Matched-reference captures	**1,558**, plus three admitted before reference availability

The p99 ≤250 ms / observed maximum ≤500 ms test envelope passed; this is not a hard real-time guarantee. Only the initial statistics boundary flag was present. Maximum direct-model FIR trace discrepancy was **1.09e−11 native-code units**, reflecting numerical consistency rather than analogue accuracy. Three subsequent ADC/reference captures passed with 1,066 statistics blocks and no losses/invalid slots. The separate headless-widget phase retained two captures and 1,769 blocks; its FIFO counter was already **144** before subscription and remained **144 in all 370 packets**, with no subscriber losses. It is not included in the counter-phase zero-loss statement. One-hour receipt.

Bounded sizes, cancellation, monitoring and readers

- `qualify.py --seconds 10 --maximum-pairs 8388608` passed the capture-size sequence through **32 MiB**, concurrent summary/raw phases, observed-busy cancellation and a correct subsequent capture, plus exact agreement between actual ADC samples and an overlapping complete FPGA statistics block. Digital receipt.
- The complete **32 MiB** raw/FIR/temporal/distribution repeat retained **17 captures, 15,850 blocks and 355 histograms in 133.033 s**, with zero statistics loss/invalid slots. Coverage was **0.858%**; control p99/max **149.0/190.6 ms**. Capture/download median was **2.578 s**, analysis **7.110 s**. Allocation succeeded throughout this bounded repeat; it is not a guarantee of future contiguous availability. 32 MiB receipt.
- The priority repeat retained **1,792** exact statistics blocks, reproduced all saved trigger/window decisions, completed one follow-up and counted one busy skip. Its measured follow-up delay was **235.8 ms**; source/revision rejection and the subsequent guarded capture passed. Synthetic LO fits passed; actual unknown-power ADC records remained ineligible for fitting. Priority receipt.
- A 60-second stalled subscriber reported **6,595 lost blocks**, while the fast subscriber retained **7,203** verified blocks in total without subscriber/FIFO losses. Control p99/max was **78.1/107.0 ms**. Closing a completed binary transfer after its first 4096 bytes and then acquiring again also passed. The service was not terminated during DMA. Reader receipt.

Resource observations

Host `/proc`/disk-space and read-only board OS memory observations were sampled approximately every 30 seconds. After excluding the first five minutes of the one-hour counter phase, **106** valid observations gave:

- Board RSS **71.80–80.11 MiB**; first/last-quarter medians **72.12/72.21 MiB**.
- Host RSS **149.3–289.8 MiB**; first/last-quarter medians **195.4/190.8 MiB**.
- Board free CMA total **82.46–95.30 MiB**; this is not the largest contiguous block.

These are sampled finite-run ranges, not instantaneous maxima or a proof against all future memory growth. Board file-descriptor enumeration was unavailable to the unprivileged observer; RSS/CMA observations succeeded. Filesystem free-space changes include other host activity and must not be treated as recording-only throughput. Per-stage `*-resources-*.jsonl` files preserve the observations; the manual's derived provenance records the calculation and input hashes. Sequence and commands.

Real browser: discovered backlog, corrected lifecycle and repeat

The first real Notebook/Chromium run found a host UI callback backlog that headless widgets had not exposed: repeated summary reads could enqueue plots faster than the notebook rendered them, starving button and cleanup messages. The local test kernel was stopped only after confirming the board capture was idle; the acquisition service remained running. That attempt and its unfinished host run are retained in the initial browser receipt.

Common `ManagedPoller` now permits **one outstanding UI publication**. A slow UI backpressures reads; invalidated delivery releases the worker, and Close does not wait for an undeliverable queued callback. Three new tests cover delayed publication, latest pending requests/cancellation, disposal and publisher failure. The continuous statistics recorder remains independent of GUI polling.

The first corrected probe caught a test-harness issue: Playwright 1.48 treats numeric `index=0` as an empty selection. Its invalid source was rejected before hardware writes; the harness now uses the actual string DOM option values and waits for readback. Retained probe. A subsequent functional run passed, but fixed screenshot delays could capture a previous view; that receipt establishes acquisition/recording/cleanup, not exact view-image identity.

The final browser repeat additionally waited for **displayed PNG SHA-256 values to match the expected raw matrix, raw distribution and filtered-distribution renders**. It passed source staging/discard, two ADC captures with **2,006 statistics blocks / 12 histograms**, raw/FIR reference save/comparison, and a monitored run with **494 blocks / 31 trend windows / three histograms / one completed follow-up**. Both recordings had zero FIFO/subscriber losses and zero invalid slots. Follow-up delay was **305.1 ms**, with no event overlap or pretrigger raw pairs. Explicitly labelled synthetic fitting and normal cleanup passed. The frontend logged a startup `schema` TypeError; it remained in the receipt and did not prevent those exercised workflows.

These browser measurements identify the host GUI correction by source hash; they do not relabel the earlier one-hour service baseline. The starter notebook's cleanup cell is now commented so Run All does not immediately close the dashboard. The actual notebook executed through a Jupyter kernel; tests after the correction passed: **105 ajku tests, no skips; 323 common tests, 20 skips**. Five site tests passed too. Logs: `.build/manual-host-tests-final.log`, `.build/manual-common-tests.log`, `.build/manual-site-tests.log`, `.build/manual-notebook-execution.log`.

Manual, export and final state

The **37-page LaTeX/PDF manual** covers operation and the six workflows, then coherent states/displacement, squeezing/loss, non-classical states, Wigner/Husimi distributions, homodyne/heterodyne vacuum noise, CW sideband/temporal modes, shot-noise normalization, worked exercises and end-user electrical/optical characterization procedures. Computed figures pass normalization/variance and single-photon Q-moment checks; the qualification table/resource plot derive from the actual receipts. The final PDF has no overfull boxes, unresolved references or text outside the checked page bounds; pages and figures were visually inspected. Manual/build sources.

A real ADC capture's full HDF5 export was reopened: **262,144 native pairs, 131,004 FIR pairs and all 127 product arrays** matched exactly, including the completion marker and integer centre/stride metadata. The publication receipt records the PDF hash/size, host environment and unchanged board calibration/ranges. The final board was ADC/calibrated-zero, idle, without capture errors or a claimed histogram. Publication/export receipt. The interface baseline is RELEASE.md. No optical performance or active-DMA process-crash safety is inferred from these results.

2026-11-01 — six first-implementation priorities, host revision 4

Implemented the original specification's §11 workflows: measured-LO-power level/ sweep software, explicit spectral matrices and pointwise coherence/phase diagnostics, captured distributions/tails, raw/FIR amplitude and power correlations, complete third/fourth cumulants with batch-jackknife errors, and continuous raw-statistics drift monitoring with guarded host follow-up captures. Usage and statistical assumptions are in PRIORITIES.md. This implements the requested features; it does not establish optical calibration or a version 1.0 release.

These tests used the working-tree implementation after commit **bd87dc7**, host processing revision **4**, FPGA/driver **0.2**, build **0xd3f1cbc0**, and common software **1.3** on **192.168.0.102**. The service installation reused the existing bitstream. No RTL, resource or routed-timing change is introduced by these host workflows; the earlier hardware receipts remain applicable.

Numerical, provenance and lifecycle checks

- **105 host tests passed, no skips**, using `make test` with the host virtualenv (NumPy, Matplotlib/ipywidgets and h5py available). **320 common tests run, OK, 25 skipped**. Logs: `.build/priorities-tests-reviewed.log` and `.build/priorities-common-tests.log`. The common suite includes workflow and design-uniformity checks. Standard-library-only host imports also pass.
- Direct leave-batch recomputation checks the re-centred cumulant jackknife. Independent quantiles/tail counts, histogram conservation, Gaussian bin integrals, zero variance, singular covariance and insufficient tail events are covered. Error estimates remain conditional on stationarity/dependence assumptions.
- Matrix Hermiticity, cross-phase sign, DC/Nyquist exclusions and the one-segment significance limit are checked. Raw/FIR archive geometry, histogram counts and reference contracts reject inconsistent products.
- Known signed polynomial coefficients and physical units, finite-difference meter Jacobians, common-scale covariance cancellation and unknown-error reporting pass. Ineligible powers/captures are rejected rather than silently omitted. Complete synthetic waveforms independently exercise both raw and FIR analysis branches.
- Drift persistence, hysteresis, gaps/context resets, exact aggregation, bounded capture admission, stale revisions and stopping with an admitted capture pass. Cancelled calibration levels retain completed products and remain ineligible.
- Headless GUI tests cover the new panels and recording modes. Final review removed an unnecessary board connection from the saved-level fitter; a regression test completes the fit/archive/plot while any attempted board clone would fail.

Calibration and drift workflow on the board

`priorities-workflows-2026-11-01.json` records the following results from `tools/qualify_priorities.py`:

- **Five synthetic power levels / 20 captures** recover the known variance model $0.00014901161193847656 + 0.2\ P + 25\ P^2$ within the test tolerances after the declared raw science-band/FIR response. The synthetic origin is saved, and `optical_qualification` is false. These are not optical power measurements.
- A **15.007 s** counter monitor retained **1,791 exact statistics blocks** and **112 trend windows**, with zero FIFO/subscriber loss and zero invalid slots. Two simultaneous rule detections yielded **one completed 8 MiB capture and one explicitly counted busy skip**, with no failed/late captures.
- Actual sample-clock follow-up delay was **260.594 ms** after the triggering interval; event overlap was false and **pretrigger raw pairs were zero**. Control latency was **113 ms p99 / 132 ms maximum** over 102 probes. This is a host-reactive workflow measurement, not a deterministic trigger-latency guarantee.
- Replaying the complete saved statistics reproduces every trend window and trigger ID. A stale revision is refused before capture; the subsequent correctly guarded capture succeeds. Two actual ADC captures at an explicitly unknown LO power are retained, and fitting them is refused. No noise level is relabelled as watts.

Combined analysis, readout and recording workload

The optimized one-minute `qualify_ms1.py` run used two-job pipelining, **8 MiB** captures, **65,536-point** Hann spectra, **±128** temporal lags, default factor-two FIR filtering, **32×32** captured joint bins, **16,384-pair** diagnostic batches and pointwise coherence alpha **0.05**, alongside statistics, histograms and independent control probing. Both raw and filtered products were analysed and recorded.

- **30 captures**, **7,502 exact statistics blocks**, **171 histograms** in **62.967 s**, with **zero FIFO/subscriber loss and zero invalid slots** in the combined counter phase; only the initial stream-boundary flag was present.

- Control **167.5 ms p99 / 190.7 ms maximum** met the existing 250/500 ms test envelope. Median analysis was **1.732 s**, capture/download **0.761 s**.
- Actual retained raw coverage was **0.7993%**. The initial implementation retained 22 captures at about **0.56%**; the repeat follows a contiguous-channel layout optimization in distribution analysis. Optional complete diagnostics have a substantial host cost; these results do not establish the 3–5% coverage target.
- Three actual ADC reference captures, signed reference subtraction and raw/FIR validation passed. Two headless-widget captures rendered the new matrix and distribution views. That later widget phase recorded cumulative FIFO loss **115 in all 373 packets**, accrued before subscription; no increase occurred during recording, and subscriber loss was zero.

Receipts: initial workload, optimized workload. Raw/processed evidence and PNGs are in `.build/priorities-combined-optimized/`.

GUI and CLI completion checks

- The real-service headless GUI recorded **316 ADC statistics blocks**, **20 drift windows**, two histograms and one completed follow-up, with no statistics loss or invalid slots. It also completed a two-capture unknown-power level and an explicitly labelled synthetic fit. Raw/FIR matrix/distribution, drift and fit PNGs were inspected; browser interaction was not tested. GUI receipt.
- The live drift CLI (`--seconds 3 --pairs 262144 --nfft 4096 --diagnostics --diagnostic-bins 32 --filter`, using the two counter rules from the workflow test) retained **371 statistics blocks**, completed one capture and counted one busy skip. Its measured follow-up delay was **252.293 ms**, again with no raw event overlap or pretrigger history.
- CLI replay matched all **112 windows** and trigger IDs from the earlier 15-second run. `iqsweep.py record` retained two actual ADC captures at unknown power; the fitter refused them. `iqsweep.py fit --allow-synthetic` reopened the five synthetic levels and produced an explicitly synthetic NPZ/filtered plot, with quantum origin unverified. CLI receipt.

The final board check found ADC/calibrated-zero, no active capture or claimed histogram, and zero capture errors. Optical power/receiver calibration, longer stress qualification, active-DMA process-crash ownership and the bounded operating limits remain separate from completion of these six software workflows. The scope checklist keeps broader research wishes distinct.

2026-11-01 — full host FIR filtering/decimation workflow

Implemented optional, single-pass symmetric Kaiser low-pass filtering and integer decimation in host processing revision **3**. Raw and `filtered_*` products coexist; the original `.iqbin`, FPGA statistics/histograms, hardware/driver **0.2**, build **0xd3f1cbc0** and common **1.3** remain unchanged. No FPGA build/deployment was needed. This completes the explicit digital analysis-band implementation gate, not the remaining first-release features or optical calibration.

Numerical, provenance and lifecycle checks

- **86 host tests passed, no skips**, with NumPy, SciPy comparison helpers, Matplotlib/ipywidgets and optional h5py available. **320 common tests run, OK, 25 skipped**. Standard-library-only imports pass under `python3 -S`. Logs: `.build/filter-tests-final.log`, `.build/filter-common-tests.log`.
- FIR results match independent direct convolution across chunk sizes including single-output chunks, every tested phase, factor-one filtering and non-power- of-two decimation. Overlapping captures agree at common valid centres. Constant inputs remain constant; read-only/Fortran-order input arrays are not mutated.
- Tests verify impulse centres and delay signs, passband gain, rejection of a tone that would alias into the passband, white-noise variance gain, correct PSD units at the reduced rate, Parseval, and integer coordinates above **2^63**. No padded edge values or cross-capture state enter the output. Counter-wrap input is refused.
- Designs and imported profiles are checked for rate/alias constraints, real odd symmetric unity-DC taps, coefficient/descriptor hashes and actual sampled response. A self-consistent hash on a falsely labelled all-pass filter is not accepted. Short records/invalid decimation fail before a live capture request.
- Raw products remain numerically unchanged when enabling the optional branch. Filtered moments, temporal lag seconds/raw-clock offsets and saved geometry match direct references. Loaders reject inconsistent coordinates, frequency grids, missing/malformed products and tampered coefficient/reference files.
- Accepted settings freeze exact coefficients before requests. In-flight work retains its original filter while later edits invalidate its current-display status. GUI raw/FIR views, exact reference-filter staging, profile CLI replay, signed raw/FIR comparisons and unchanged raw-coverage accounting pass.

- DC-separated filtering and offset-stable centring/detrending prevent large DC offsets or constant data from generating misleading numerical noise. This is floating-point numerical treatment, not whitening or additional ADC resolution.
- Complete-waveform HDF5 export preserves native codes, all analysis arrays, float64 FIR samples and uint64 centre/stride metadata. Its `complete` marker is set after datasets are written/flushed; an injected mid-export failure leaves it false. It is not a power-loss durability guarantee.

Board run: default factor two

`qualify_ms1.py --seconds 120 --pipeline --lags 256 --filter --dashboard`, with default 8 MiB captures and 65,536-point Hann analysis, completed on **192.168.0.102**. The default design was **137 taps**, pass/stop **20/25 MHz**, decimation **2**, phase **0**, and output rate **62.5 MSa/s**. Its numerical response check measured **85.358 dB** stopband attenuation and **0.000861 dB** peak-to-peak passband ripple; one-sided FIR noise bandwidth **22.030 MHz**. These describe the digital coefficients, not the analogue receiver/ADC response.

- **81 captures**, **14,536 exact statistics blocks** and **325 histograms** in **121.952 s**, with **zero FIFO/subscriber losses and zero invalid slots** in the combined counter phase; initial stream-boundary flag only.
- Every record produced **1,048,508** valid FIR pairs. Reopened products were checked against a separate direct-convolution model of the complete counter period, including means/covariances. Maximum trace discrepancy was **9.10e−12 native-code units** (floating-point consistency).
- Independent control **173 ms p99 / 194 ms maximum** met the existing 250/500 ms test envelope. Median full raw-plus-FIR/temporal analysis was **1.256 s**, capture/download **0.771 s**; raw coverage **1.11%**. Keeping both branches adds computation/storage; decimating on the PC does not reduce wire bytes.
- Three actual ADC captures/reference checks passed, including direct FIR trace comparisons within **6.66e−16 V numerically**, not a voltage-accuracy claim. Headless widgets rendered the FIR response, raw overlays and filtered temporal panels and retained two captures plus 1,632 statistics blocks. Their FIFO counter was **101 in all 344 packets**, accumulated during the unsubscribed interval between phases; no increase occurred while recording. Subscriber loss was zero. Browser interaction was not tested.

Receipt: factor-two workflow. Raw/processed evidence and the inspected dashboard PNG are in `.build/filter-board-2x/`.

Board run: factor three and nonzero phase

With 1 MiB captures, 4096-point FFTs, ±64 lags, pass/stop **5/12 MHz**, decimation **3** and phase **2**, a 30-second pipeline run passed. Its **99 taps** gave a sampled **84.750 dB** stopband and **0.000853 dB** ripple, at **41.666667 MSa/s**.

- **108 captures**, **3,653** exact statistics blocks and **82** histograms in **30.669 s**; zero statistics loss/invalid slots and control **150/157 ms p99/max**.
- Phase-dependent output lengths **87,348 and 87,349** both occurred and were verified, with every output centre on the declared raw-counter residue. The independent counter-model trace discrepancy was at most **8.19e−12 codes**. Actual ADC trace checks agreed within **5.14e−16 V numerically**.
- Raw coverage was **0.74%** for this smaller-record workload. This is an alignment/ response/retention check, not an improved-throughput operating point.

Receipt: factor-three workflow.

CLI, reference replay and HDF5

The live CLI acquired a two-record filtered run and saved a reference/profile, then acquired two more records using that exact saved profile and matched the reference. Both raw and filtered signed subtraction retained negative bins. The full HDF5 export was reopened: all native codes and all **1,048,508** filtered pairs matched their source/reconstruction exactly, including the complete marker, integer first centre **4002158764186** and stride **2**. Reference `filter.json` and the separately exported profile were identical. Export/replay receipt.

Usage, formulas, archive layout, transition-band interpretation and revision-2 reference migration are documented in FILTERING.md. The board remains on ADC/calibrated-zero with no active capture/claim. Raw transfer goals, remaining distribution/predictor/FPGA engines and broader release qualification remain tracked in `wiki/first-release.md`.

2026-11-01 — board repeat of committed code b2cd4c3

Tested commit **b2cd4c31b97e645412e061c91d7287f3f414e73c** on **192.168.0.102**, FPGA/driver **0.2**, build **0xd3f1cbc0**, common software **1.3**. All 31 tracked driver/framework source files on the board matched the commit's Git blob hashes before and after testing. The existing FPGA session and service were retained; no deployment was needed. Receipts: source/identity preflight and final source/board state.

Digital checks and allocation limits

- `make test-hw`: exact counter captures, buffered moments and every histogram bin passed. Board-local histogram readout was **24.6 ms**. Holding a claimed bank for 0.75 s counted five missed epochs without new acquisition gaps. FIFO drops during this direct, unsubscribed test are expected. Hardware receipt.
- The full-ceiling `qualify.py --seconds 10` sweep again **failed the 64 MiB contiguous allocation** (`Allocate failed: 4294967295`) after correct smaller records through 32 MiB. It did not reach its concurrency phases. A subsequent processor-memory snapshot showed `CmaFree` 107668 kB of 131072 kB, and `MemAvailable` 344640 kB; those totals do not establish one allocatable aligned 64 MiB region. The allocator's underlying cause remains undetermined. Failed sweep, memory snapshot.
- `qualify.py --seconds 20 --maximum-pairs 8388608` then **passed** all captures through 32 MiB, three concurrent-readout phases, observed-busy DMA cancellation, a correct subsequent capture, and exact ADC/FPGA agreement for all five sums of one complete 2^20-sample block. The three phases retained 2401, 2387 and 2402 exact statistics blocks, with zero FIFO/subscriber loss. The raw/statistics/ histogram phase retained **2.42%** raw coverage. Isolated 8/32 MiB downloads were 29.3/30.5 MB/s; capture-plus-download was 19.9/22.1 MB/s. Bounded digital receipt.

Ten-minute combined acquisition, analysis and storage

`qualify_ms1.py --seconds 600 --pipeline --lags 256 --dashboard` used default 8 MiB captures, 65,536-point Hann spectra, full cumulants, raw temporal products, reference comparisons, statistics/histogram recording and independent control probing. Plots were checked in a separate final widget phase, not during the principal throughput measurement.

Quantity	Measured result
Completed captures / elapsed run time	**643 / 601.883 s**
Verified statistics blocks / histogram snapshots	**71,740 / 1,479**
FIFO loss / subscriber loss / invalid slots	**0 / 0 / 0**
Retained raw time / wall-time coverage	**10.78775 s / 1.79%**
Independent SNAPSHOT queries	**3,151**
SNAPSHOT median / p99 / maximum	**83 / 189 / 355 ms**
Histogram read/check/store p99 / maximum	**261 / 395 ms**
Capture/download median	**0.761 s**
Analysis median, including temporal products	**0.767 s**
Board immutable-copy median	**65.3 ms; native backend throughout**

All published raw/processed files were reopened and checked. There was only the initial statistics boundary flag; 640 captures used matched references and three were admitted before the reference became available. The independent SNAPSHOT probe met its declared p99 ≤250 ms / maximum ≤500 ms envelope. The histogram column includes host verification/storage and is a different operation.

The additional actual-ADC reference run and headless widgets also passed: three ADC captures/563 statistics blocks, then three widget captures/1587 blocks/nine histograms, without losses or errors. Negative excess remained signed and changing processing rejected the reference comparison. Ten-minute receipt.

A separate 32 MiB pipeline/temporal repeat passed with **19 captures in 64.848 s**, **7,723** exact statistics blocks, **167** histograms, no statistics loss/invalid slots, **1.97%** raw coverage and control **171 ms p99 / 186 ms maximum**. Larger-buffer receipt.

The first ten-minute launcher was interrupted by the host command timeout after about one second; it is not counted as a successful run. Its incomplete directory remains in `.build/board-b2cd4c3-10min/`. The board reported its issued capture fully acknowledged, idle and error-free, and was returned to SAFE before the fresh run in `.build/board-b2cd4c3-10min-retry/`. This was not a board-service crash test.

Stalled reader and completed-transfer disconnect

New host-only `tools/qualify_readers.py --stall-seconds 60` checks subscriber isolation through the real service. The slow connection's receive buffer is set before the TCP handshake and remains unchanged during recovery.

- After **60.120 s** without reading, the slow subscriber reported **6,608 lost blocks** and an observed sample-index gap. Its retained counter blocks were correct. The fast subscriber retained **7,171** blocks during the pause and **7,211** in total, with **zero subscriber loss and zero hardware FIFO loss**. Independent control p99/max was **61 / 84 ms**. The loss indication appeared after 46 resumed packets, in about 61 ms.
- Closing an **already completed 8 MiB capture's binary download** after checking its first 4096 bytes left the following capture correct and hardware idle. The concurrent statistics reader retained 62 exact blocks without loss. This tests immutable completed-data transfer cleanup, not abrupt process death during DMA.
- An initial probe shrank the receive buffer after connecting and enlarged it before recovery; its five-second read timed out. The fast reader stayed correct, but slow-reader loss reporting was not established by that attempt. The revised test sets the buffer before connecting and does not resize the zero-window socket. No service/driver/FPGA change was made for the successful repeat.

Initial probe, successful reader/transfer checks.

Final state: ADC source, calibrated-zero outputs, no active capture, no capture error and no claimed histogram. These results extend the bounded operating evidence; **64 MiB availability, 3–5% analysed raw coverage, optical qualification, hours-long endurance and active-DMA process-crash ownership remain open**.

2026-11-01 — first-release host analysis, bounded pipeline and completed-buffer copy

This advances the first-release checklist; it is not a version 1.0 completion claim. Hardware/driver remain **0.2**, build **0xd3f1cbc0**. Host processing revision **2** and common software **1.3** add faster numerical kernels, a bounded two-job acquisition/analysis pipeline, raw temporal products and a GIL-free copy of completed DMA buffers. Board: **192.168.0.102**; host Python 3.12.9 / NumPy 2.3.5. The board remains on PYNQ 3.0.1 / Python 3.10.

Numerical and ownership checks

- Reused moment powers and channel-contiguous real FFTs replace repeated full-record exponentiation and strided complex transforms. A cold-process profile of the same saved 8 MiB counter record measured `analyze()` **1.110 → 0.364 s**, including initial imports. Distribution processing was **0.528 → 0.058 s** and spectra **0.451 → 0.168 s**. These are host measurements, not guaranteed capture rates. Logs: `.build/v1-profile-{baseline,optimized}.log`.
- Independent direct complex FFTs, odd/even transform lengths, SciPy Welch/CSD, Parseval, known delays and a **1e-16-power weak image** agree. Complete third/fourth cumulants and chunk-boundary moments match direct reference calculations.
- Temporal tests compare every lag pair with a direct summation oracle, including partial blocks, lag direction and complex correlations. They distinguish power coupling from near-zero linear correlation, preserve weak radial fluctuations, and verify batch errors directly. An AR process produces different effective counts for means and variances. Constant/short records report undefined or unavailable estimates rather than inventing certainty.
- Pipeline tests cover the **two-admitted-job bound**, request-pinned context, stop with pending work, connection/transfer/analysis failures and early close. Cleanup-log failures cannot strand later jobs. Backpressure delays acquisition; it does not hide discarded records. Selection is throughput-driven, not fixed-cadence.
- `rpfpga.buffers.immutable_copy` fills a private new bytes object through libc while releasing the GIL, only after DMA drain/cache invalidation. It retains the source export and performs no MMIO. Tests check immutable results, buffer/C-API reference lifetimes, fallback and cancellation during copying: a foreign edit cannot publish the old context. Failed copies publish no record.
- Final host suite: **66 tests passed, no skips**, including HDF5 after installing optional h5py 3.16.0 in the ignored host UI environment. Common suite: **320 tests, OK, 25 skips**. Standard-library-only imports/copy pass under `python3 -S`. Receipts: `.build/v1-tests-final.log`, `.build/v1-common-tests.log`.

Measured combined workloads

All rows include statistics streaming, frequent histogram reads, host analysis and storage, references and independent control probing. Control requests wait 100 ms after replies; histograms wait 300 ms after read/check/record.

Every retained counter payload, statistics block and histogram is checked; published raw and processed files are reopened and verified.

Workload	Captures / elapsed	Raw coverage	Control p99 / max	Outcome
Original MS1, 8 MiB, serial, no temporal products	65 / 121.8 s	0.90%	171 / 185 ms	Baseline
Faster kernels + pipeline, 8 MiB, common 1.2, no temporal products	146 / 120.8 s	**2.03%**	179 / 235 ms	Passed
32 MiB pipeline + temporal products, Python snapshot	21 / 64.6 s	2.18%	**337 / 358 ms**	Failed p99 ≤250 ms gate
32 MiB pipeline + final temporal products, native snapshot	19 / 65.1 s	**1.96%**	**168 / 183 ms**	Passed
8 MiB pipeline + final temporal products, native snapshot	**336 / 301.2 s**	**1.87%**	**186 / 225 ms**	Passed

The final temporal workload uses ±256 lags and 16,384-pair blocks, with errors at up to three batch scales and direct radial-power calculation to avoid cancellation of weak fluctuations. These are development workload receipts, not isolated copy-speed benchmarks. Native copying improves control concurrency; it does not by itself make every workload faster or meet the 3–5% coverage goal.

- The **five-minute 8 MiB run** retained **35,905 exact statistics blocks** and **760 histograms**, with **zero FIFO/subscriber loss and zero invalid slots**; only the initial stream-boundary flag. Every copy reported the native backend. Median capture/download **0.757 s**, full analysis **0.753 s**, board copy **65.8 ms**. Processing and transfer overlap; their durations must not be added to estimate pipeline cadence.
- The final **32 MiB run** retained **7,747** exact statistics blocks and **166** histograms without statistics loss/invalid slots. Median board copy was **257 ms** while other control threads could run. Median host analysis was **2.98 s**; larger records are not a universal throughput improvement.
- An initial 32 MiB attempt stopped because the qualifier assumed every capture after its second had seen the newly saved reference. A slot can be reused while reference saving proceeds. The qualifier now checks the request-pinned reference ID and reports pre-reference admissions (three here). Early close recorded its unprocessed admitted jobs as failures. The subsequent latency failure above was real and was not accepted by relaxing the envelope.
- The 3–5% coverage goal, compiled raw TCP transport and hours-long memory/network/ process-failure qualification remain open. Zero statistics loss does not imply retention of every waveform or histogram epoch.

Archived receipts: 8 MiB / five minutes, 32 MiB. Intermediate reports remain in `.build/v1-pipeline-8mib/` and `.build/v1-pipeline-32mib-temporal{,-repeat}/`. Raw/processed evidence remains in the corresponding `.build/v1-native-*` directories.

ADC, interface, digital hardware and rebuild

- Actual ADC reference/temporal processing and incompatible-processing rejection passed. The five-minute qualifier's additional ADC run retained three captures and 568 statistics blocks without loss. Inputs were uncontrolled electrical signals, not optical calibration; negative reference excess remains signed.
- Headless real-board widgets displayed the new correlation/error panels and saved three captures, 1,585 statistics blocks and nine histograms in 13.0 s, without errors or losses. The PNG was inspected; browser interaction remains untested. CLI `--captures 3 --pipeline --lags 256 --save-reference ...`, then a two-capture matching reference comparison, both completed. Cumulative FIFO counters retain earlier unsubscribed losses. Logs: `.build/v1-cli-pipeline-{reference,compare}.log`.
- `make test-hw` passed counter captures/moments/histograms with common 1.3; histogram readout **24.2 ms**, held-reader epoch count/gaps correct. A separate `qualify.py --seconds 5 --maximum-pairs 8388608` passed through 32 MiB, observed-busy cancellation/drain and the next capture. One complete actual ADC block again matched all five exact FPGA sums. Receipt: digital repeat.
- The shared deployment workflow invalidated the local project fingerprint and rebuilt despite unchanged RTL. The first shell invocation timed out; its logged retry completed build/export/deployment. Inspected implementation `runme.log`, utilization, routed timing and DRC: **4,153 LUTs, 6,837 FFs, 2,216 slices, 24 BRAM36, 3 DSPs, one MMCM; WNS +0.367 ns, WHS +0.027 ns, setup/hold TNS 0**. Build ID remains **0xd3f1cbc0**. The ten existing DRC advisories and 18 outputs without output-delay constraints remain; no new errors/critical warnings or unconstrained internal endpoints. No new FPGA engine was added. Log: `.build/v1-native-copy-deploy-retry.log`. Deployment retained the running FPGA session and installed common 1.3.

The board is left on ADC/calibrated-zero with no active capture or claimed histogram. Remaining release work is in `wiki/first-release.md`; the full manual follows implemented scope and qualification, not a premature version bump.

2026-11-01 — MS1 host workflow and native readout, FPGA 0.2

The first usable CW host workflow is implemented and digitally board-tested: calibrated burst spectra/moments, provenance-matched saved references, integrated run recording and shared-framework live controls/plots. Board **192.168.0.102**, PYNQ 3.0.1, common software **1.2**, FPGA/driver **0.2**, build **0xd3f1cbc0**. `make install-service` reused the current bitstream and deployed the software. No RTL or FPGA-resource change was made; the routed 0.2 results below still apply.

Native readout and repeat digital tests

- The optional common C reader uses the existing PYNQ mapping within checked Overlay transactions. Compilation is outside the hardware lock. Model/Python fallback, side-effecting repeated reads, validation-before-read, stale/loading gates, cancellation and lock exclusion pass the nine native-reader tests.
- Board-local full histogram readout decreased from **99.6 ms to 24.3 ms**. Every counter histogram bin and buffered moment still matched. Holding a bank for 0.75 s counted five missed epochs with no new acquisition gaps. Receipt: native readout.
- A repeat of the full 64 MiB transfer sweep **failed allocation** with `Allocate failed: 4294967295`, after correct 32-pair, 1 MiB, 8 MiB and 32 MiB records. It did not reach its concurrency phases. The earlier successful 64 MiB test remains valid, but the API ceiling is not a memory-availability guarantee. The failure's root cause within the allocator was not established.
- Repeating with `--maximum-pairs 8388608` passed through **32 MiB**, observed-busy DMA cancellation/drain, a subsequent 32-pair capture, and actual ADC/FPGA moment agreement for one complete 2^20-sample block. The isolated 8/32 MiB downloads were 30.3/30.1 MB/s; capture-plus-download was 20.2/21.8 MB/s.

The bounded transport repeat used three 20-second phases:

Workload	Verified statistics blocks	FIFO/subscriber loss	Other result
Statistics only	2402	0 / 0	Initial stream-boundary flag only
Statistics + histograms	2396	0 / 0	103 histograms; TCP call latency 37–70 ms
Statistics + 8 MiB raw + histograms	2423	0 / 0	29 correct captures; **2.39% raw coverage**

The native path improves histogram/control occupancy; it does not accelerate raw DDR copying or TCP export. Evidence, including the failed sweep: selected transport receipts. Full reports remain in `.build/ms1-native-transport{,-32mib}.json`.

Full analysis/storage/control workload

`tools/qualify_ms1.py --host 192.168.0.102 --seconds 120 --directory .build/ms1-board --dashboard` exercised the actual `LiveSession` and `RecordingRun`:

- **65 × 8 MiB** exact counter captures, all processed and written with file hashes; every raw/processed product was reopened and checked after recording.
- **14,511** exact statistics blocks and **324** verified full histograms in **121.785 s**. Zero FIFO drops, subscriber losses or invalid slots; only the initial statistics stream-boundary gap flag. Unread histogram epochs still expire normally; zero statistics loss does not imply retention of all epochs.
- **1.09052 s** of raw ADC-clock time retained: **0.8954% wall-time coverage**. Median capture/transfer **0.809 s**, median host analysis **0.944 s**; persistence and other work add further time. This workload includes higher moments, spectra, reference comparisons, compressed processed files and independent control probing. It must not be described as 3–5% coverage or continuous full-rate export.
- Separate control connection: **674** `snapshot()` requests, with 100 ms waits after replies. Median **71.7 ms**, p95 **146 ms**, p99 **171 ms**, max **185 ms**. This meets the test's declared p99 ≤250 ms/max ≤500 ms envelope. The common aspirational 100 ms responsiveness target is **not** met by every query.

- Separate histogram connection: 300 ms waits after each read/check/record. Median full operation **62.3 ms**, p99 **181 ms**, max **206 ms** under load. These numbers include host checking/storage, unlike the board-local 24.3 ms.

These are bounded workload measurements on the existing host/Ethernet connection, not a hard real-time or unattended-endurance guarantee. Receipt: MS1 workflow; raw/processed evidence remains in `.build/ms1-board/`. The archived JSON preserves all report fields.

References, ADC workflow and interface

- Counter and ADC references were saved, reloaded with raw/processed checksums and compared with new captures. Actual ADC data use the recorded board calibration (`loopback + dmm`) and DAC stream words `[204, -60]`. The ADC inputs were not controlled for an optical experiment; no sensitivity or vacuum claim is made.
- ADC subtraction retained **13,171 negative excess-PSD values**; no quantum-noise inference was returned. Changing FFT processing was rejected as incompatible. The three-record ADC run retained 649 statistics blocks with zero losses.
- Live ipywidgets/poller integration against the board generated plots, summaries and complete recordings. The corrected stop/idle check repeated with **four captures, 1,654 statistics blocks and ten histograms in 13.606 s**, with no errors, invalid slots or subscriber/FIFO loss. The PNG was rendered headlessly; interactive browser behaviour has not been tested in this session. The first receipt's widget stage includes 70 s of unnecessary test-helper joins on reusable idle workers; the repeat waits for idle before closing them.
- The standalone host CLI acquired/saved a reference, then recorded a compatible comparison successfully. Its FIFO counter retained **485 historical drops** from the same revision's earlier unsubscribed interval; the CLI does not reset hardware configuration to erase backlog/loss history. `complete` means file/job completion, not zero loss. Receipt: widget/CLI repeat.
- Offline host suite: **50 tests, OK, one h5py skip**, using the host ipywidgets environment. Cases include spectral normalization, signed subtraction, incompatible/tampered references, exact coverage accounting, stop-with-admitted work, stale-context suppression, rejected source revisions, widgets, and checking archived run products for corruption. `.build/ms1-tests-final.log` is the receipt.
- Common framework/workflow/uniformity suite: **316 tests, OK, 25 skips** in the initial root `make test` invocation. The command exceeded the shell's four-minute limit during feedback tests; remaining designs were then run directly. Suites passed for analog_echo (18), atomic_clock (223), awg (83), feedback (72), lockin (67), mult_add (121) and synth (82), with their optional-dependency skips. Timetagger's wider test attempt ran 58 tests but failed importing its existing **Python 3.14** `ttcore` binary under **Python 3.12** (`PyThreadState_GetUnchecked`); this unrelated local build artifact was not modified. The repository-wide suite is therefore not claimed fully green. Logs: `.build/ms1-repository-tests.log` and `.build/ms1-remaining-design-tests.log`. A narrowed root invocation also tripped the uniformity test because inherited `DESIGNS=...` replaced its expected all-design list; direct per-design invocations avoid that invocation issue. No FPGA simulation/source-build gate was repeated for this software-only change.

The usage guide is MS1.md, with milestone scope in `wiki/MS1.md`. The board remains on ajku, ADC source and calibrated-zero DAC words, with no active capture or histogram claim after testing.

2026-11-01 — initial board qualification, build 0xd3f1cbc0

Board: STEMlab 125-14 / Zynq-7010, PYNQ 3.0.1, Linux 5.15.19, at **192.168.0.102**. The repository-root `config.mk` now supplies this address to all design Make workflows (commit `849a851`); per-design overrides and command-line precedence were verified by the 307-test common suite (25 skips) and 34 LED tests. The previous timetagger acquisition was already stopped; it was put in SAFE and its service stopped before programming ajku through the shared workflow.

Loading and exact digital checks

- `make test-remote`: passed preflight, safe driver loading, sysid/build/ABI checks, power-up registers and the `acquire` preset. `make install-service` installed and started the ajku worker. The FPGA bitstream was reused, not rebuilt; ID remains **0xd3f1cbc0**.
- `make test-hw`: exact counter pairs for 32, 4096, 262144, 2097152, then 32 samples; exact buffered first/second moments; every joint/marginal histogram bin matched. A second run held one histogram for **0.75 s**: its epoch stayed fixed, **five missed epochs** were counted, and **no new acquisition gap slots** appeared.

- The board-local full histogram read took **99.6 ms**. Initial direct-test 8 MiB capture/copy/**board-side checking** took 1.066 s (1.097 s on the second run). That timing includes expensive NumPy checking on the ARM; it is not the raw capture or network-transfer time. FIFO drops during these direct tests are expected because no statistics subscriber drains the queue while checking.
- Through the TCP service, records up to **16,777,216 pairs / 64 MiB** matched the complete counter pattern. This exercises full-rate 500 MB/s capture for a 134.22 ms window; it does not demonstrate continuous full-rate export.
- A configuration edit during an **observed busy 64 MiB DMA capture** returned CANCELLED, the writer became idle, and the next 32-pair record was correct.
- One complete **2^20-sample block of an actual ADC recording** was identified by its exact sample indices. All five FPGA sums matched the corresponding raw samples bit-for-bit. The inputs were not controlled analogue calibration signals; this is an acquisition/statistics consistency check, not a bandwidth/noise-floor test.

The internal counter/complement pattern repeats every 16,384 pairs. These checks establish agreement with that pattern and the hardware completion/error counters; they are not an independent absolute-time or process-crash integrity proof.

Host transfer and concurrent readout

Measured on the existing Ethernet connection using `tools/qualify.py`. MB/s below is decimal; MiB denotes the binary recording size. Capture/copy/control includes allocation and command turnaround. Download is separately timed through EOF.

Record	Capture/copy/control	Download	Download payload	End-to-end payload, excluding host checking
1 MiB	59 ms	68 ms	15.3 MB/s	8.2 MB/s
8 MiB	137 ms	292 ms	28.8 MB/s	19.6 MB/s
32 MiB	471 ms	1.098 s	30.6 MB/s	21.4 MB/s
64 MiB	940 ms	2.166 s	31.0 MB/s	21.6 MB/s

Three approximately ten-second phases ran with a statistics subscriber:

Phase	Verified statistics blocks	FIFO / subscriber loss	Other observations
Statistics only	1201	0 / 0	Initial stream-boundary flag only
Statistics + repeated histogram queries	1207	0 / 0	36 correct histograms; host call latency 112–150 ms
Statistics + 8 MiB raw captures + histograms	1199	0 / 0	13 correct raw records and histograms; **2.18% wall-time raw coverage**

The last phase's host histogram calls and raw requests were sequential, while the FPGA histogram engine and service statistics drain ran throughout. Its raw download rate was 19.5–22.1 MB/s and raw capture-plus-download rate 13.2–14.6 MB/s, before host checking and histogram time. This is below the proposed 20 MB/s total payload operating point when those activities are combined; **3–5% coverage is not a blanket guarantee for the default 8 MiB mode**.

A separate run kept the statistics subscriber active during 8, 32 and 64 MiB captures. All payloads and **964 statistics blocks** were correct, with zero FIFO or subscriber loss and no gap beyond the initial stream boundary. End-to-end raw rates were **14.3, 16.1 and 16.5 MB/s**, respectively (excluding host checking). Thus larger records can approach 3.2–3.3% raw coverage in that particular workload, at the cost of more memory and longer turnaround. These are short runs, not an hours-long endurance qualification.

Evidence: initial board checks, held histogram, host qualification, large captures with statistics. The host verifier itself rejects corrupted payloads, moments and histograms in offline tests: `make test` now runs 34 tests, one h5py-dependent skip. The on-board, non-root Python client initialization used by the notebook also passed via the local service (`adc`, 2^20 valid samples, zero invalid slots). The notebook now uses that local service by default, with `BOARD_HOST` as an explicit PC override; its installed predecessor was verified unchanged and backed up by the common upload workflow before replacement. The board is left running ajku with the ADC source and calibrated-zero DAC words.

2026-11-01 — 0.2 buffered statistics and histograms, build 0xd3f1cbc0

Implemented the independent 64-block statistics FIFO, service-owned statistics stream, and three-bank 64×64 joint / 2×1024 marginal histograms. Histogram epochs span 2^24 slots and align to 16 complete statistics blocks. Reader claims protect one bank while the other two continue counting/clearing; losses are explicit.

- `make test`: **33 tests run, OK, one h5py-dependent test skipped**. New cases exercise FIFO order/overflow/APPLY flushing, histogram ordering/calibrated edges, count conservation, release after a failed read, actual localhost summary streams, revision changes, reload termination, subscriber overflow and independent drainers' sample-index gaps. Standard-library-only imports also pass with `python3 -S`.
- Initial synthesis caught a real resource failure: the histogram clear/increment code inferred extra memory ports and **25,344 distributed-memory LUTs**. The core alone needed 35,243 LUTs and could not fit. It was not accepted as a working build. A single write statement with pre-muxed address/data fixes the port inference; an isolated 4096-bin bank then maps to exactly four RAMB36 tiles. The whole-design resource/timing reports, not that isolated check, determine acceptance.
- `make sim`: **all six benches passed** (`tb_sysid`, `tb_stats`, `tb_capture`, `tb_core`, `tb_stats_fifo`, `tb_histograms`). New checks cover whole-frame FIFO overflow/wrap/flush, constant and alternating histogram addresses, invalid slots, all-bin reference comparisons, bank reuse and source changes, held-reader epochs without acquisition gaps, and statistics-boundary alignment. AXI tests also exercise a POP status read at write-response acceptance and a histogram read concurrent with RELEASE. The production-length full-scale moment test still passes.
- `make`: **full build and export passed**, Vivado 2025.1, xc7z010clg400-1, normal common board recipe. Build ID **0xd3f1cbc0**. Whole-design placed utilization: **4,153 LUTs (23.60%), 6,837 FFs (19.42%), 2,216 slices (50.36%), 24 BRAM36 (40%), 3 DSP48 (3.75%), one MMCM**. All resource caps pass. The 24 BRAMs are four capture, two statistics FIFO and eighteen histogram tiles; **zero LUTRAM**.
- Routed timing: **WNS +0.367 ns, TNS 0; WHS +0.027 ns, hold TNS 0**; pulse-width slack +1.845 ns. The +0.25 ns setup-margin target passes.
- Inspected implementation `runme.log`, placed utilization, routed timing and routed DRC. Bitgen completed with zero errors and no critical warnings. The DRC advisories remain the six DSP input-pipeline, three DSP output-pipeline and one SmartConnect no-routable-loads warnings. There are no unconstrained internal endpoints; the common converter constraints still report 18 outputs without output delays. These results do not constitute analogue/converter qualification.
- Common framework/workflow suite: **298 tests run, OK, 25 skipped**.
- `python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed**, including content-based invalidation for vendor/design edits, failed-build upload rejection, and clean saved-project regeneration. Initial and restored source-only builds both have ID **0xd3f1cbc0**. Receipts are in `.build/step02-verify-build.log` and `.build/verification-_7mimo0l/verification_logs/`. No board upload/test has been performed.

2026-10-31 — acquisition/statistics prototype, build 0x3abc6145

- `make test`: **24 tests run, OK, one skipped** (integer model, calibrated moments, input validation, finite capture ordering/metadata, cancellation and ownership-aware cleanup, safe loading/ABI, recording format, live localhost protocol/binary client, spectral signs/normalization and comparison against SciPy). Additional cases cover stale DMA quarantine, failed allocation, ADC calibration means, correct power-up/reference output readback in volts, and truncated TCP data. NumPy and SciPy were available. The h5py round-trip test was skipped because h5py is not installed; HDF5 export is therefore not exercised in this environment. The package also imports successfully under `python3 -S`, without site packages.
- `make sim`: XSim 2025.1 passed `tb_sysid`, `tb_stats`, `tb_capture`, `tb_core`. Cases include signed rails, invalid slots, exact block boundaries, AXI address/data backpressure, repeated captures, abort with an unaccepted address, FIFO overflow, missing ADC data, error responses, invalid addressing/length, independent AXI-Lite address/data arrival, WSTRB, atomic snapshots and calibration reference shadows. A full **2^20-slot** block at I=−8192, Q=+8191 also matches all five exact sums, exercising the high accumulator bits at the production block length.
- `make`: complete recipe creation, synthesis, placement, routing and bitstream export passed using Vivado 2025.1, `xc7z010clg400-1`, board definition `redpitaya.com:redpitaya-125-14:part0:1.0`. Export build ID `0x3abc6145`.
- Whole-design placed utilization: **1,969 LUTs (11.19%), 3,422 FFs (9.72%), 979 slices (22.25%), 4 BRAM36 (6.67%), 3 DSP48 (3.75%), one MMCM**. The FIFO mapped into block RAM; there is no distributed RAM, only two LUT shift registers. All values are below the agreed whole-build caps.
- Routed timing: **WNS +0.773 ns, setup TNS 0, WHS +0.013 ns, hold TNS 0**; pulse-width slack +1.845 ns.

No unconstrained internal endpoints or clocks.

- Inspected `ajku.runs/impl_1/runme.log`, the placed utilization report, routed timing summary and routed DRC. Bitgen reports zero errors. Routed DRC has six DSP input-pipeline advisories, three DSP PREG advisories and one SmartConnect no-routable-loads warning; no errors or critical warnings. The common converter constraints retain 18 outputs without output delays and clock/input-delay methodology warnings. Timing success applies to the specified constraints, not a newly qualified converter output interface.
- Common design uniformity: all five `test_designs.py` tests passed.
- Full common suite: **298 tests run, OK, 25 skipped** (optional dependencies, root-only SD-image mounting and/or explicitly slow tests). This includes the new design's fake-tool workflow, SD-image installation checks and uniformity.
- `python3 ../common/tests/verify_build.py --design ajku`: **all eight stages passed**: fresh recipe build, unchanged-build reuse, missing-export recovery, vendor HDL content change with preserved modification time, design RTL change, failed-build upload rejection, cleanup and saved-project regeneration. Both the initial and restored source-only builds have ID `0x3abc6145`. Evidence is retained in `.build/verification-11atn6pt/verification_logs/` and `.build/verify-build.log`. The verifier uses network sentinels; no board upload or hardware test was run.

Remaining board qualification

1. Extend repeated captures and abort/drain tests to long runs and deliberate CPU, memory and network contention, including allocation pressure.
2. Extend the bounded MS1 control/readout envelope to slow or stalled clients and deliberately heavier contention. Improve the combined transport/readout operating point before promising the proposed 20 MB/s payload or 3–5% coverage for arbitrary concurrent workloads.
3. Test client disconnect and controlled process failures while preserving DMA buffer ownership. A quarantined allocation is not crash recovery.
4. Characterize the electrical path and analogue anti-alias filtering; then add a coherent optical probe and independent power readings for optical calibration.

The 0.2 statistics FIFO and histograms have bounded retention. The initial latency and concurrent-readout measurements above do not establish sustained operation under every workload. Compiled raw transfer, histogram zoom and anomaly capture remain future milestones. See `../wiki/implementation-plan.md`.