Timing & Charge Analysis for TI-LGADs

Iskra Velkovska

Jozef Stefan Institute, F9

Testbeam analysis meeting
11.09.2026

OVERVIEW

Motivation for timing/charge analysis of an TI-LGAD

- Per-DUT timing resolution (σ_t) from constant-fraction-discriminated pulse timing, measured against a reference signal (TREF: 1.2 x 1.2 mm^2 HPK LGAD sensor).
- How that resolution degrades as a track approaches the inter-pixel distance -> the key feature studied here
- A single “bulk” resolution number per pad, so the trench-approach curves can be read as departures from a clean baseline

Charge analysis built on the timing macro

- Every operating point (which CFD fraction to use, per channel) is scanned and fixed, on bulk data

- charge macro based on timing
- Same event / track → waveform matching
- Same “bulk region, safely away from the gap” method already used for the timing bulk-ROI

- Charge response of each pad vs. track position -> in the bulk, and approaching the inter-pad gap.
- How charge is shared between two neighboring pads near that gap
- An absolute charge scale in fC, plus normalized observables

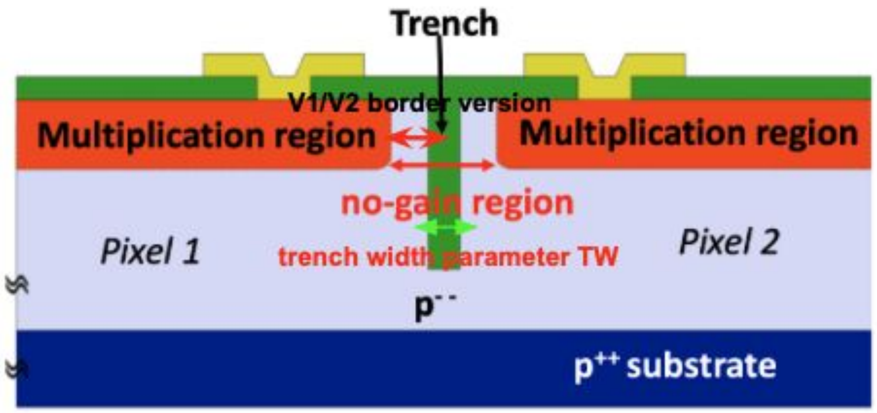

notice: bulk here is the part of the pad that is far enough from the trench

Example event showing the fixed TREF-anchored signal definition

- TREF waveform defines the time anchor.
- DUT signal is integrated in a fixed 20-bin window relative to that anchor.
- A same-length pre-signal window is used for the baseline reference.
- dt is the time difference between DUT and reference.

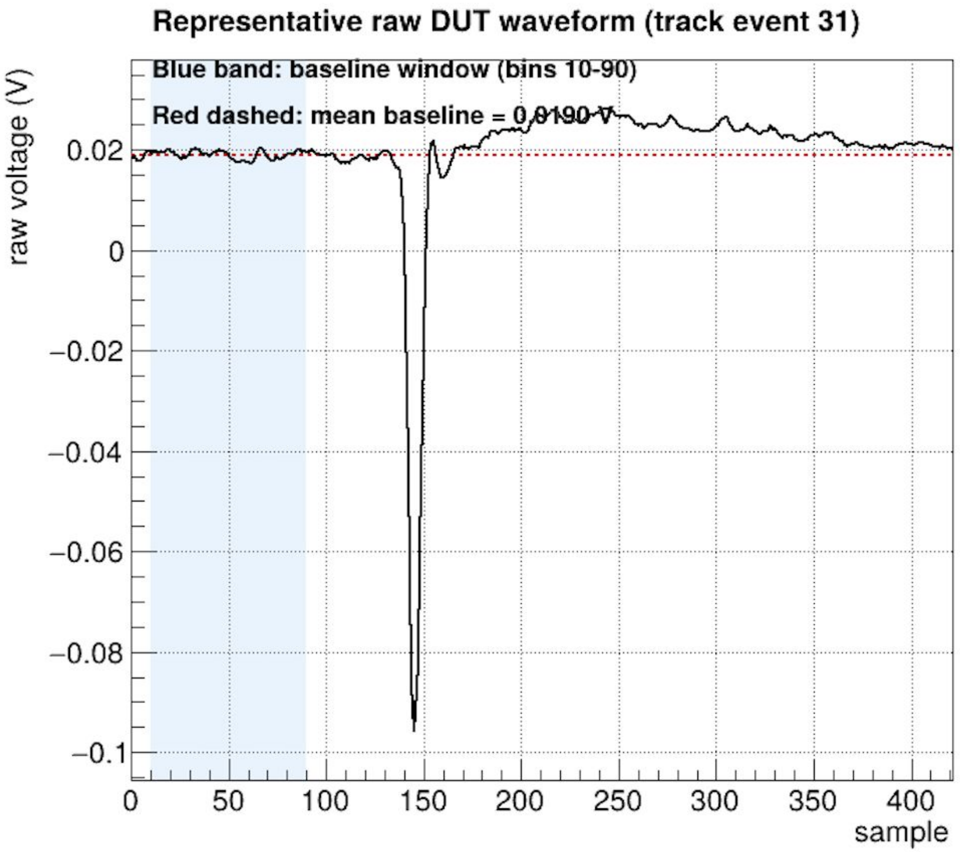

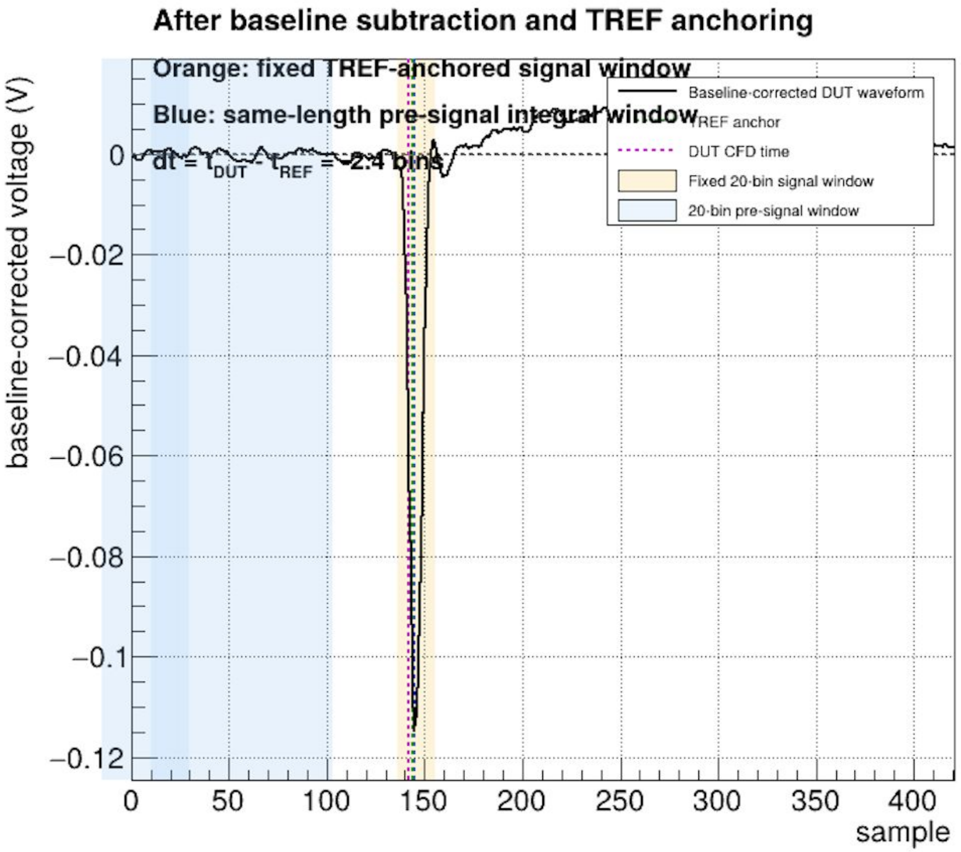

From Waveform to Charge

- Baseline -> mean of the first 100 samples (also gives the noise RMS).
- Charge = signed sum of (baseline − V) over a fixed window, converted to fC via the 0.2 ns/sample period at 5 GS/s.
- The window is not gated on the DUT's own CFD success -> every track gets a charge, including faint inter-pad events.
- Nominal window: 5 bins before + 15 after center = 4.2 ns. Short (4+11) / long (7+18) variants are recorded only as a stability cross-check, never for event selection.

In general:

- TREF CFD defines the event time anchor
- Same-length pre-signal window used for baseline reference
- $dt = t_{DUT} - t_{REF}$
- The DUT charge is integrated in a fixed 20-bin window placed relative to the TREF CFD anchor, to avoid the pedestal bias of peak-centering while retaining the main part of the DUT pulse for the selected events in this run

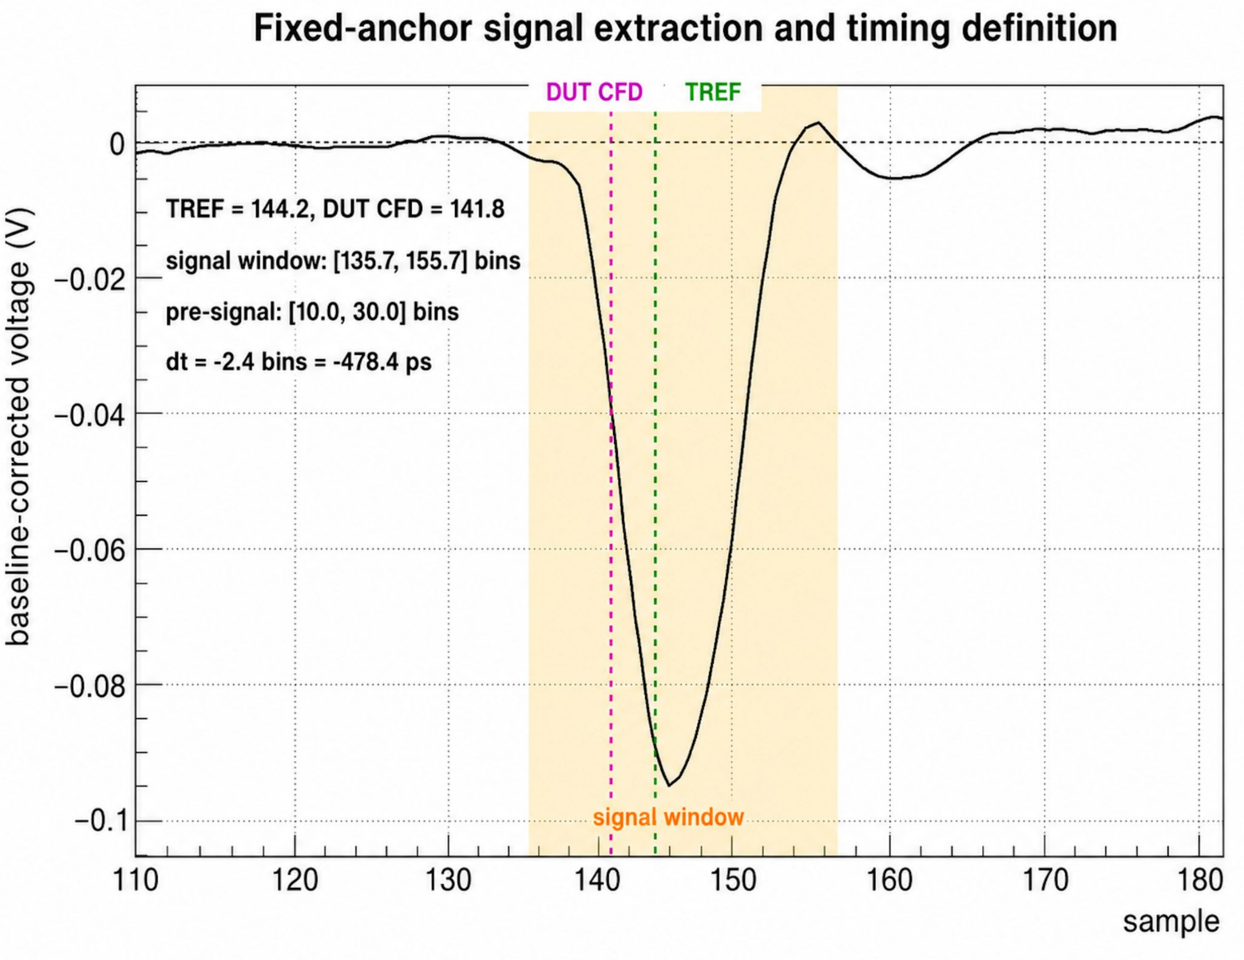

threshold = k_CFD × A_peak

timestamp = linear interpolation of the crossing between the two samples straddling threshold

A pulse is only accepted if its peak clears 5σ of that channel's own baseline

Calibrating TREF: the Triplet Method

$$\sigma_{TREF}{}^2 = ½(\sigma_{1R}{}^2 + \sigma_{2R}{}^2 - \sigma_{12}{}^2)$$

three pairwise timing widths (TREF-DUT1, TREF-DUT2, DUT1-DUT2) solve algebraically for each individual channel's own resolution

- TREF's own resolution is extracted from three pairwise widths among TREF and two calibration DUTs, using the same three-channel variance decomposition throughout this analysis.
- k_{TREF} is scanned, both calibration DUTs are held at a fixed, standard fraction so only TREF's own operating point moves.
- A scan point is only eligible as "best" if it clears BOTH a relative-error floor and a minimum-statistics floor
- Result: one k_{TREF} and one σ_{REF}, per group, and used as the fixed reference for every subtraction in every later stage.

with tracking included $\sigma_{TREF} = 40.5 \pm 1.8$ ps

Frozen once, reused everywhere

- This is the ONE place σ_{TREF} is measured in the entire macro.
- Later resolves an independent "anchor triplet" using the frozen k's, purely as a cross-check
- Every quoted σ_{DUT} (the bulk ROI, the 1D profile, the 2D map) subtracts the same σ_{TREF}, in quadrature

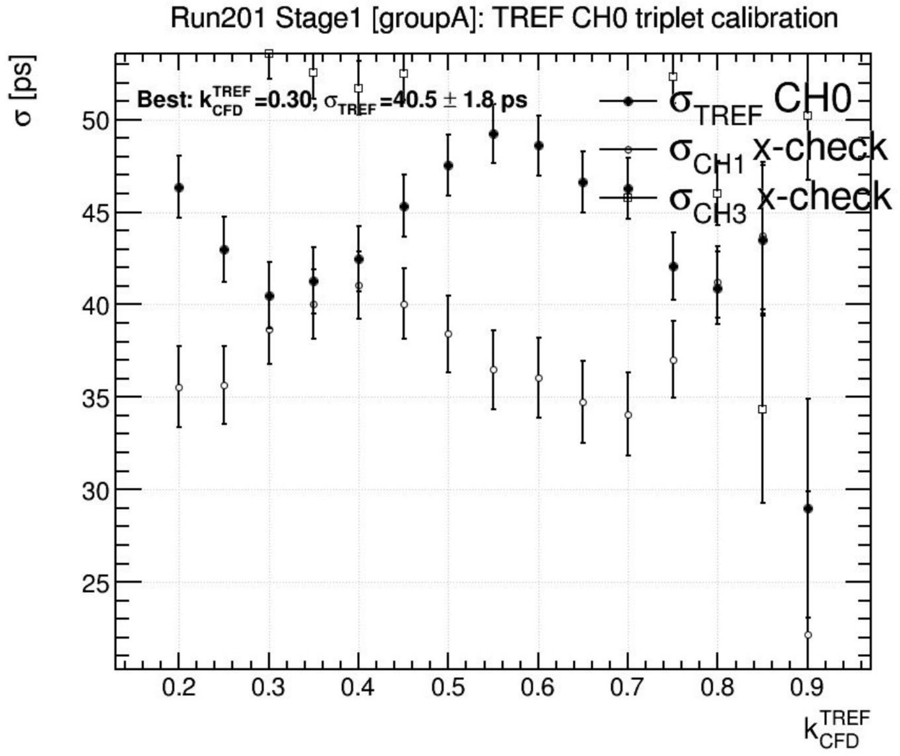

Per-Channel Scan, and Fixing How "Best k" Is Chosen

$\sigma_{DUT}(k) = \sqrt{(\sigma_{pair}(k)^2 - \sigma_{TREF}{}^2)}$ ε_CFD(k) = N_timestamped(k) / N_preselected

*N_preselected is fixed ONCE per channel — independent of which k is being scanned; tracks landing on the pad with good S/N ~**1000 events***

- For each of the 12 DUT channels, k_{DUT} is scanned against that group's now-fixed TREF; each point's σ_{pair} has σ_{TREF} quadrature-subtracted to give $\sigma_{DUT(k)}$.
- Best-k is only eligible if its CFD efficiency clears 92 % of a k-independent preselected population (scanning with different k)

For the later spatial analysis (after k_{CFD} is frozen): statistics per spatial bin are much smaller >>

$$\sigma_{68} = ½\,(Q_{84} - Q_{16})$$

the central-68% half-width of the Δt distribution — a percentile, not a fitted parameter

- A spatial bin in the 1D profile or 2D map typically has 60-200 tracks — enough for a percentile, but a known-fragile regime for a binned χ^2 Gaussian fit (most of a 120-bin histogram sits empty).

- very bin still carries a Gaussian-fit cross-check column alongside σ_{68} in the numeric output tables.
- If the two disagree by more than 20% of σ_{68} in a given bin, that bin is explicitly flagged

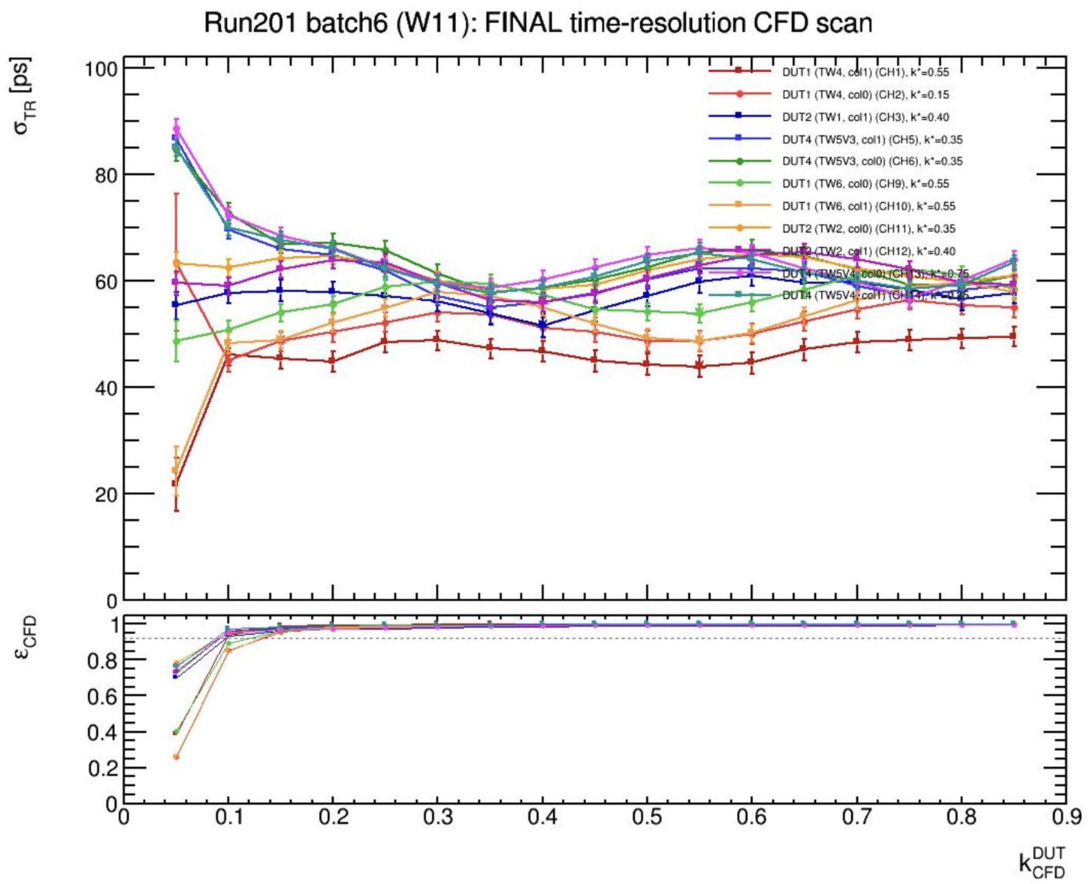

Timing resolution evaluation for W11 V2–TW4, unirradiated TI-LGAD biased @ 150 V

- Added at the per-event level: a signed waveform area (V·ns), an anchored peak amplitude, and their ratio (area/peak, in ns -> an effective local pulse width), all binned and plotted vs. distance from the trench next to the timing panels.
- The center it integrates around is not the DUT's own CFD-found peak -> it's a TREF-predicted position: a channel-specific DUT-minus-TREF peak-bin offset, calibrated once on clean bulk tracks (SNR ≥ 8, robust median/MAD outlier rejection, ≥120 clean differences required) and then held fixed.

Pad	NCFD	kCFD	σt [ps]
Pad 0	978	0.15	**42.3 ± 3.0**
Pad 1	897	0.55	**40.6 ± 3.6**

Gaussian cross-check:

Pad 0: 40.6 ± 2.7 ps
Pad 1: 38.8 ± 2.9 ps

- CFD timing efficiency $\varepsilon_{CFD}(x,y)$ — is this bin's population actually being timestamped, or mostly failing?
- Selected-track occupancy N(x,y) — is a grey/degraded region actually sampled by real tracks, or just empty?
- Every quantitative bin requires ≥120 CFD-valid tracks
- A separate, finer binning (5×10 µm) grid, built for visualization: σ_{68} is computed from the 120 nearest CFD-valid hits on the SAME physical side of the trench, within an 80 µm cap.

W5, V2, TW1-2, 0e14 n_{eq}/cm^2, angled at 10˚

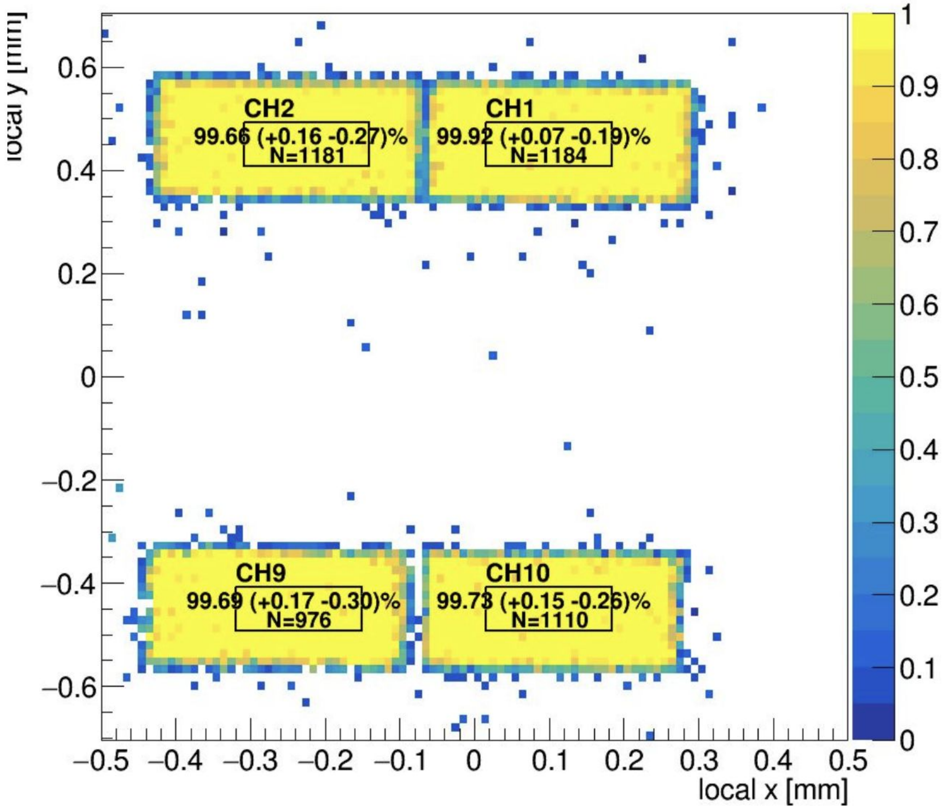

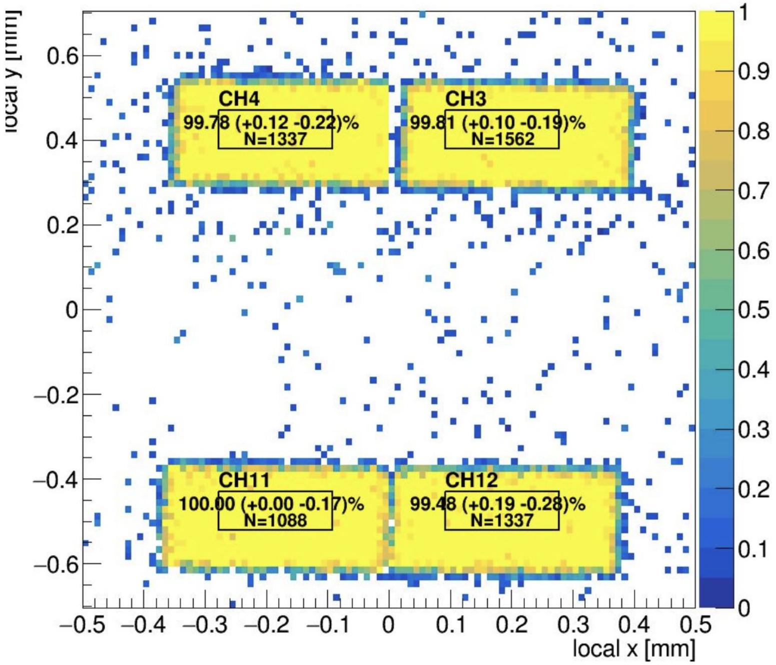

W5, V2, TW2-3, 25e14 n_{eq}/cm^2

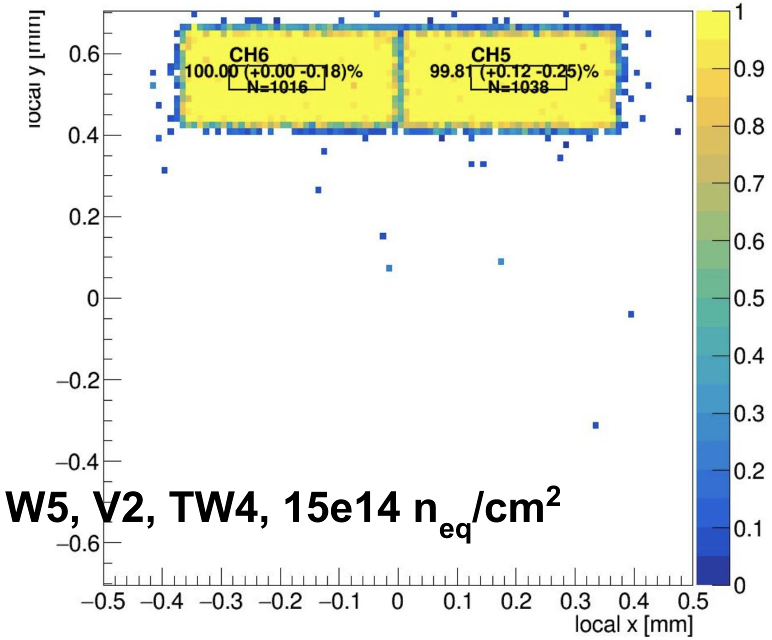

W5, V2, TW4, 15e14 n_{eq}/cm^2

Wafer 11 designs, unirradiated;

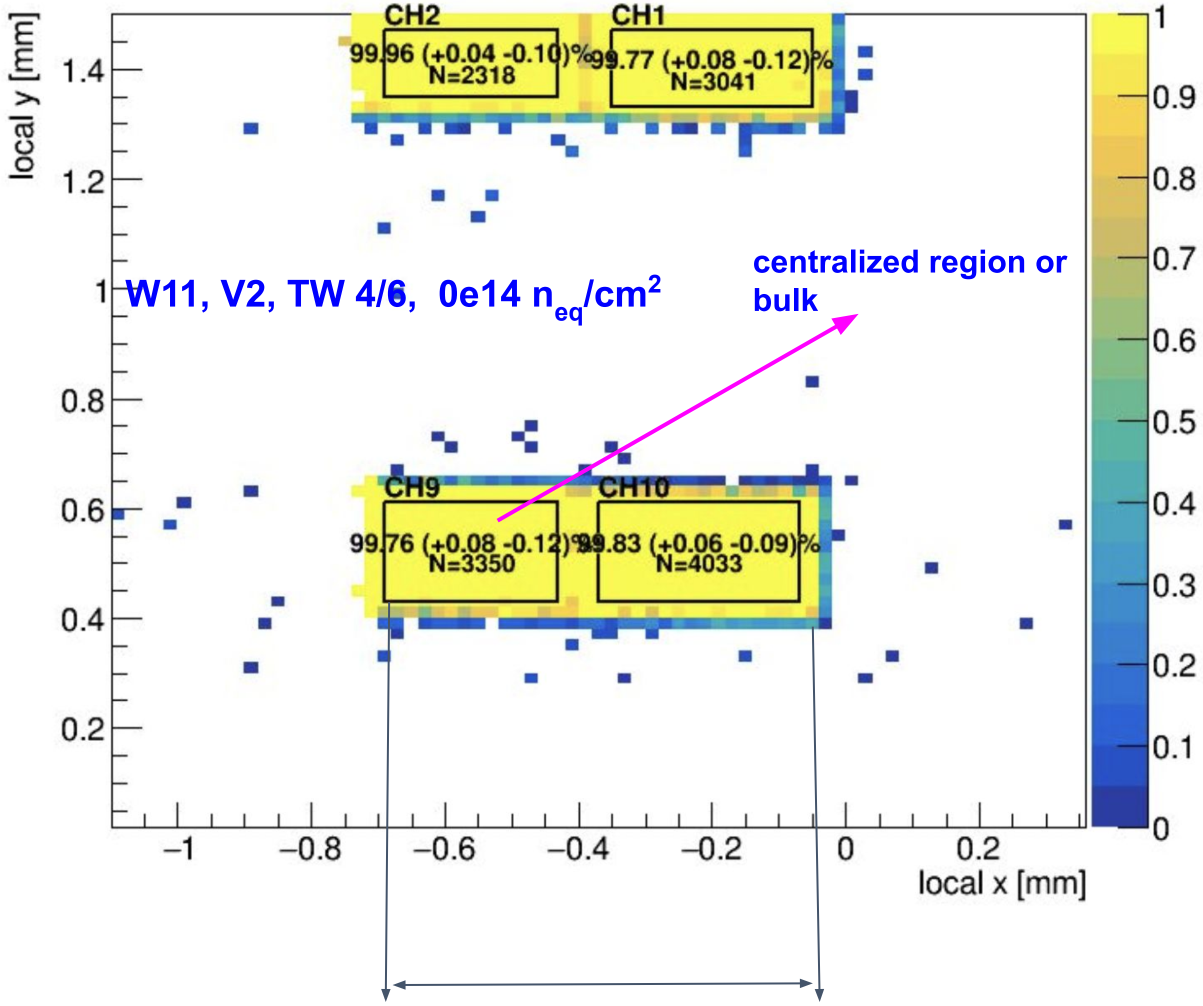

Region where time resolution is calculated and shown on the next slides

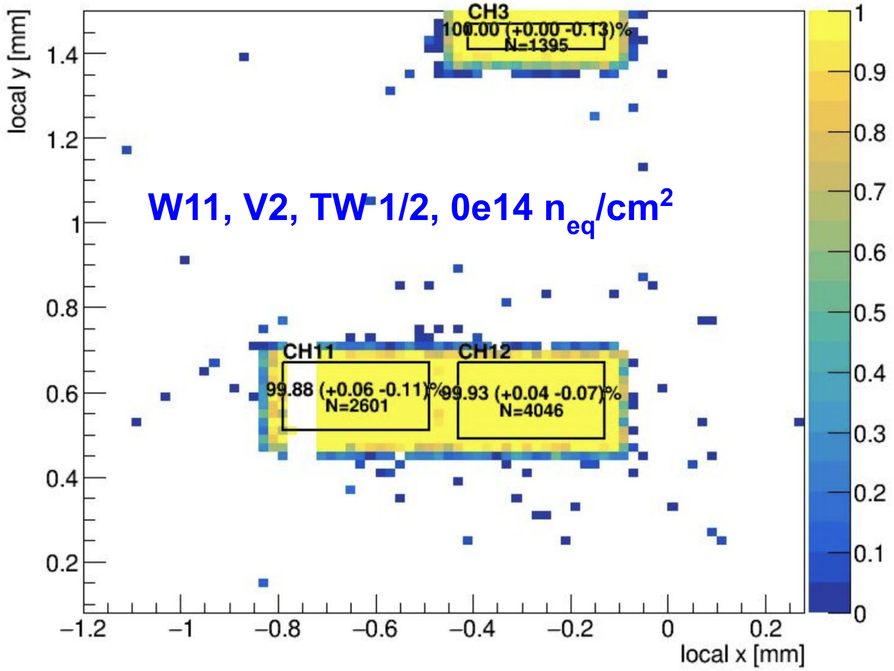

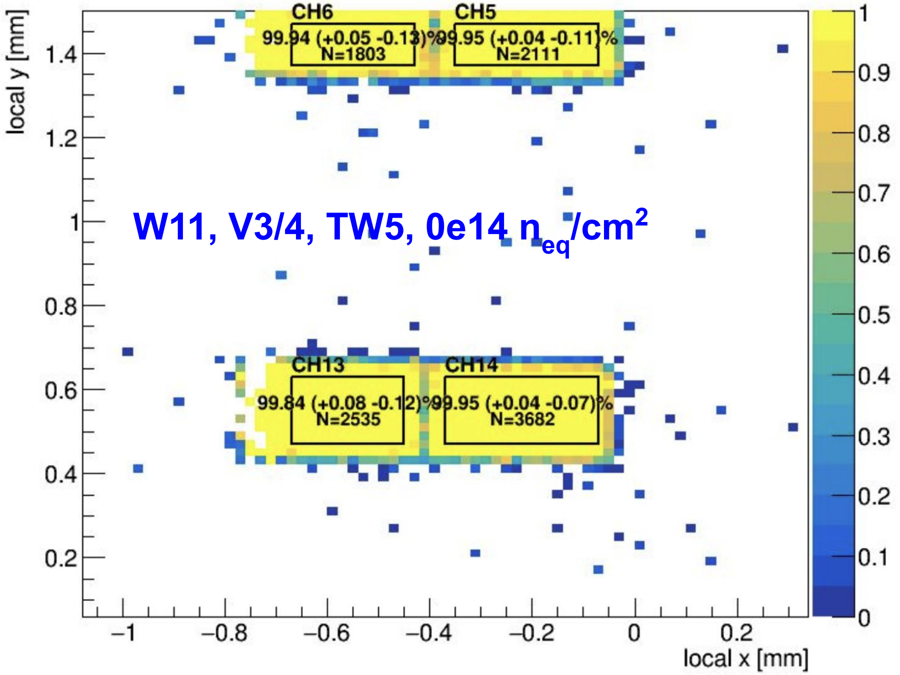

IPD results for wafer 11 and 5

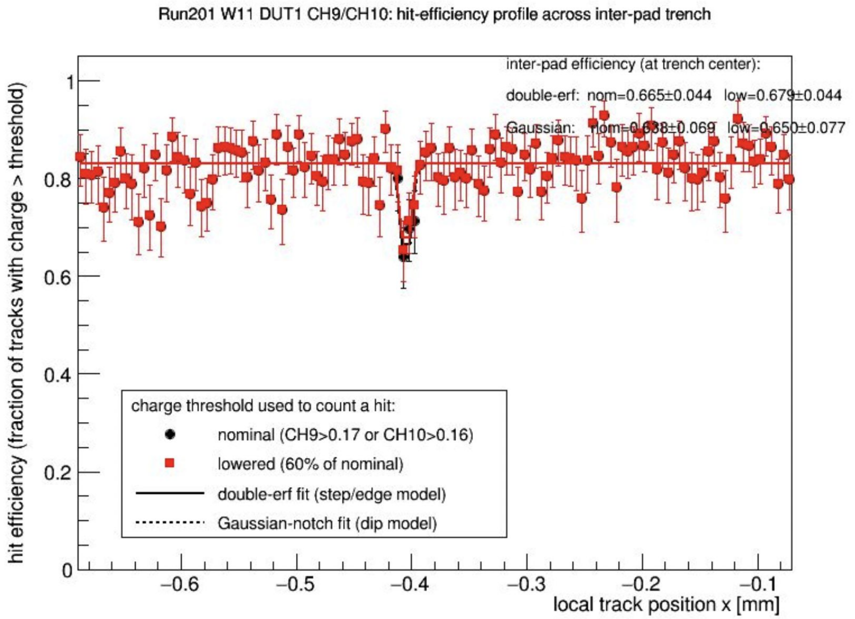

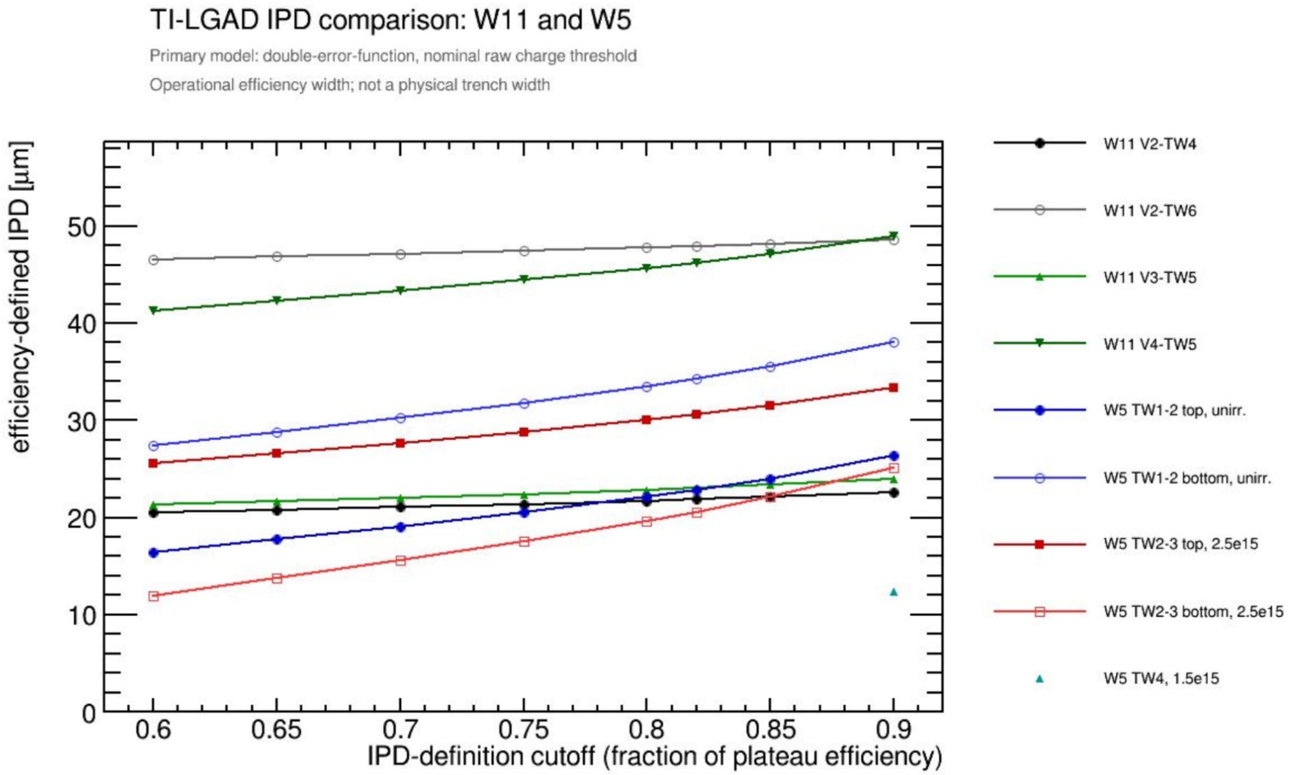

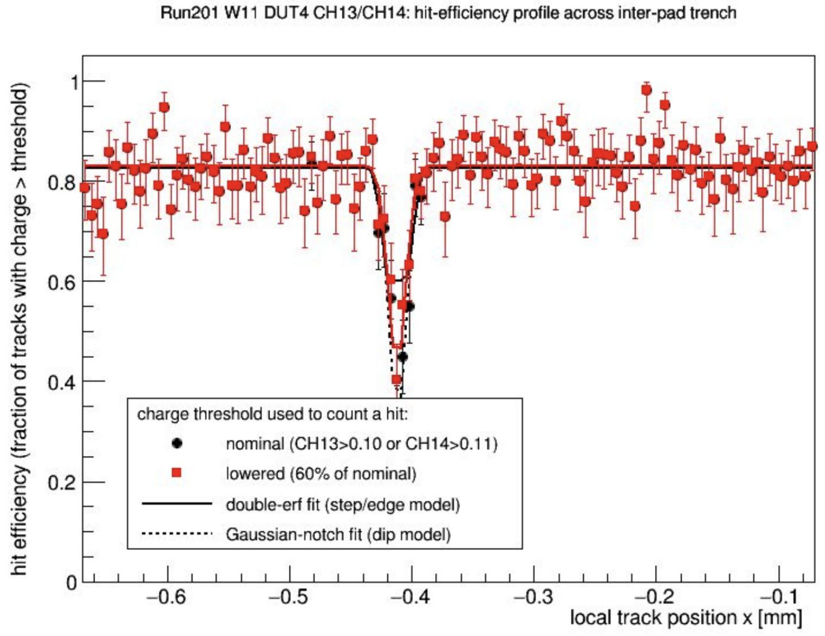

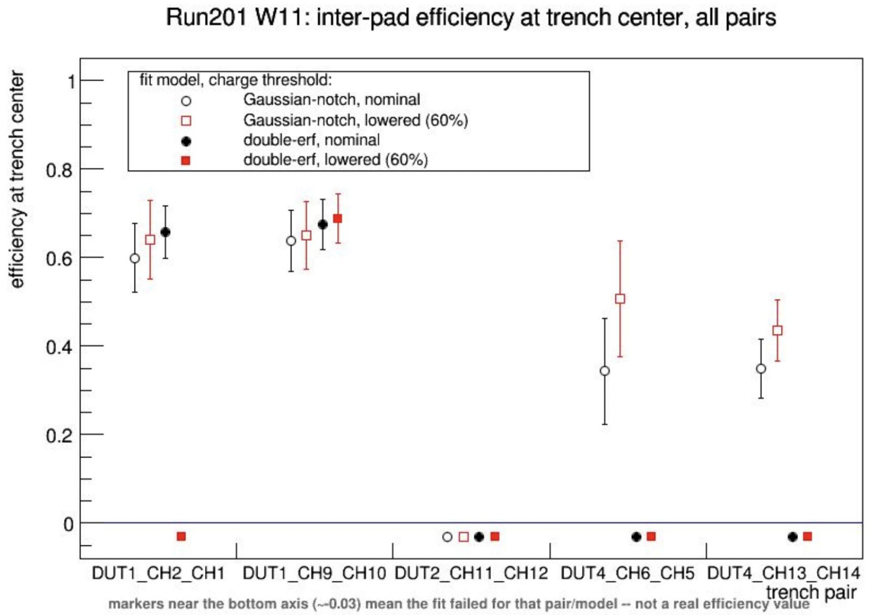

IPD is defined as the distance between the two double-erf crossings at a chosen fraction of the plateau efficiency

IPD for TI-LGAD from wafer 11 versions V2 and V3/4

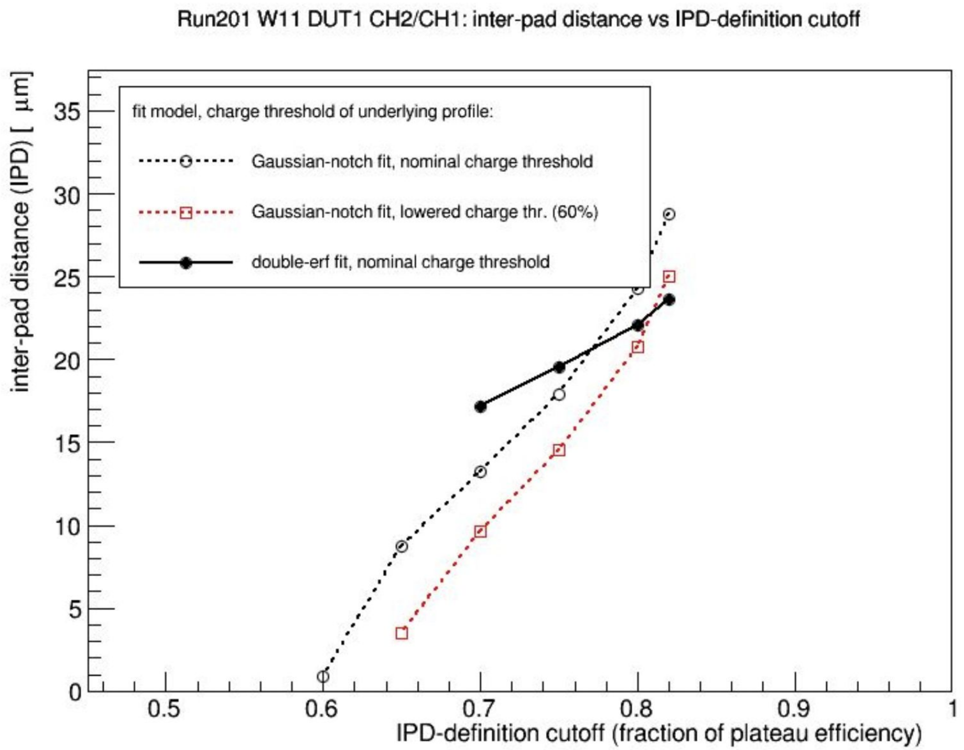

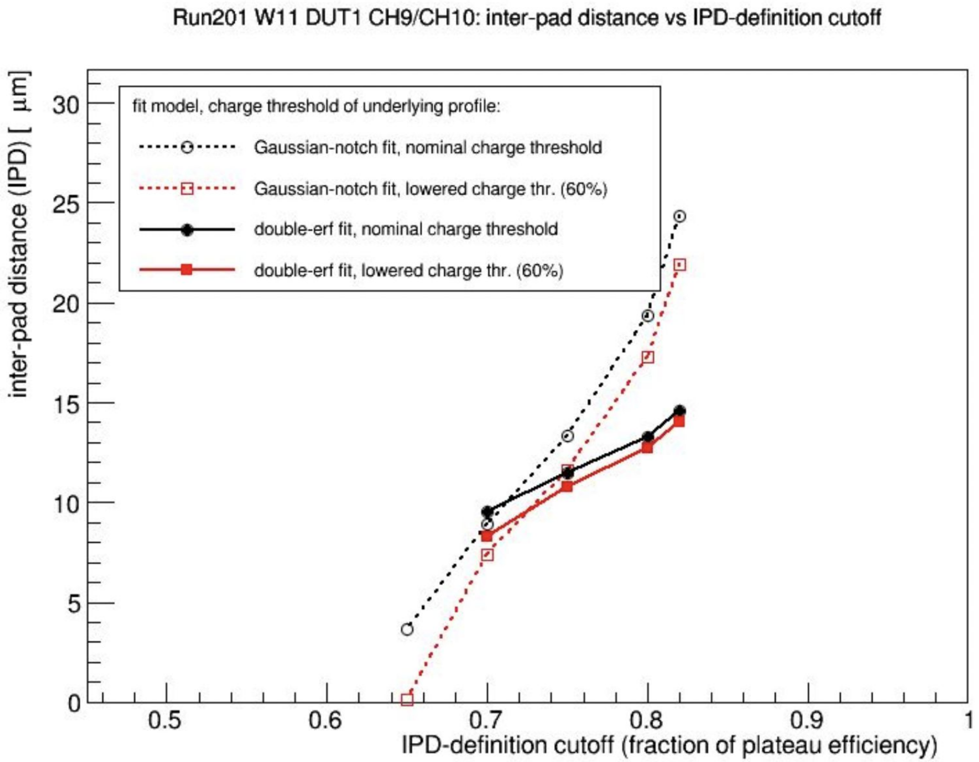

The lowered-threshold curve (0.6 Q_{thr}) is used to assess the sensitivity of the extracted IPD to the hit-threshold definition.

Efficiency vs threshold and IPD for TI-LGAD from wafer 5 and 11

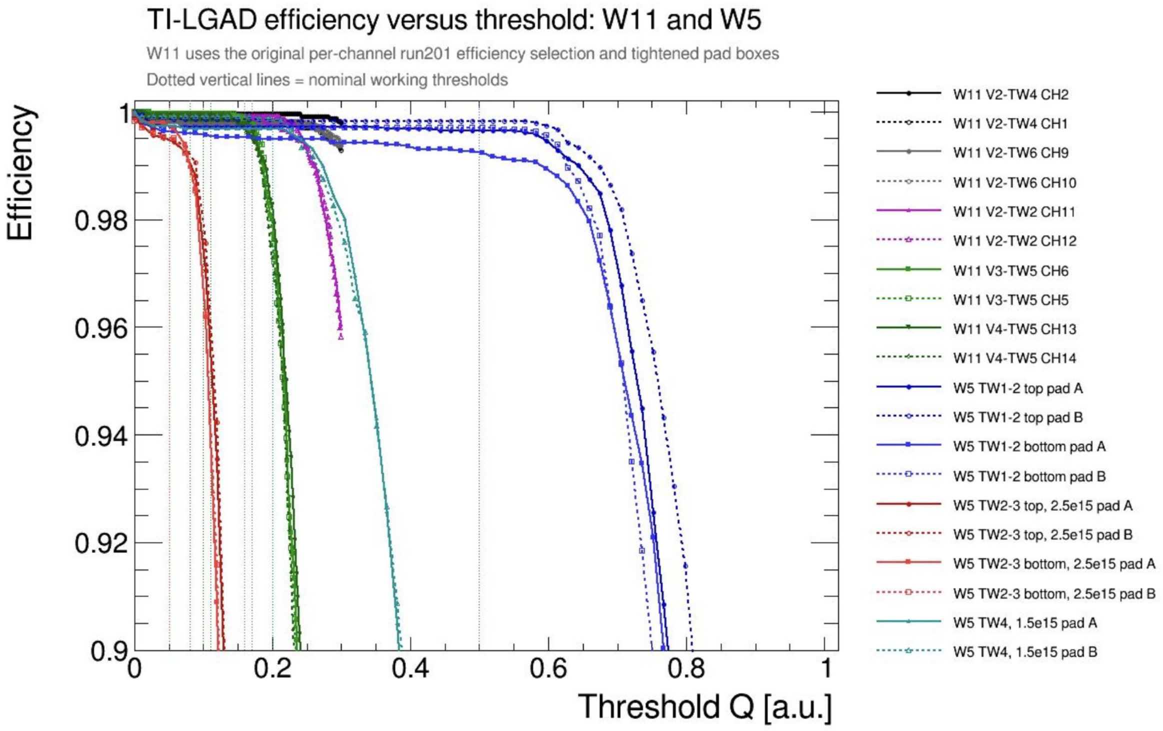

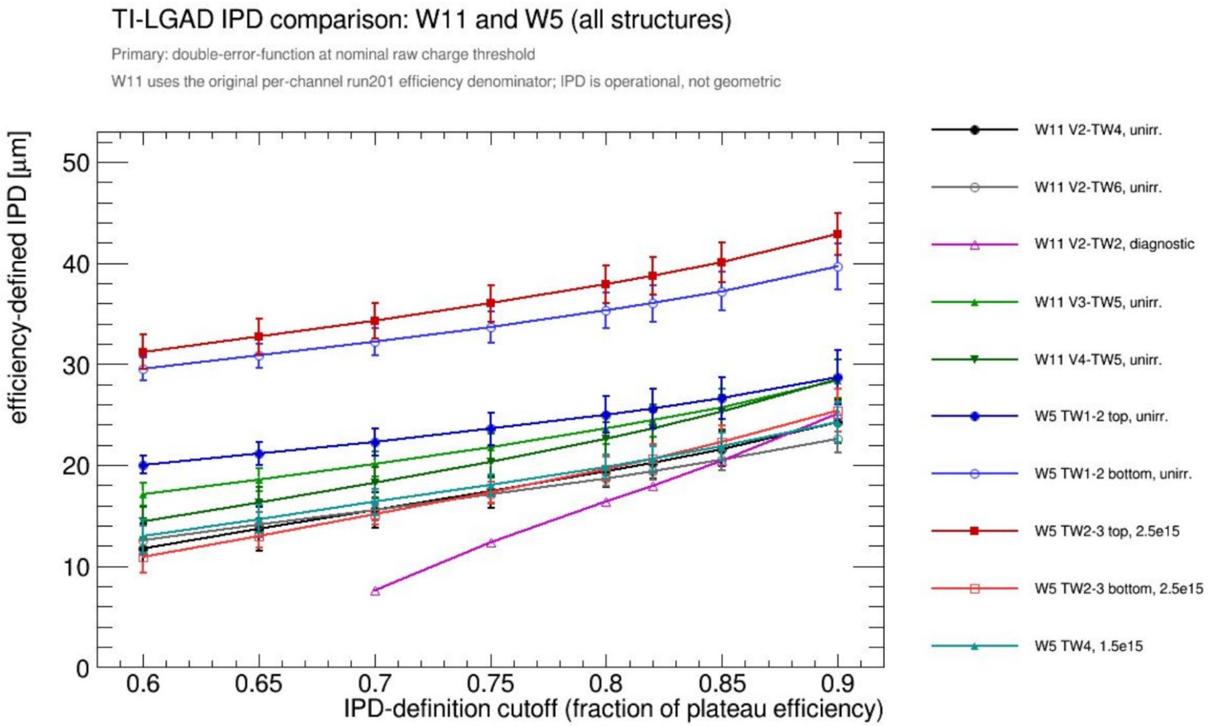

Spatial timing performance — W11 V2–TW4, unirradiated, 150 V

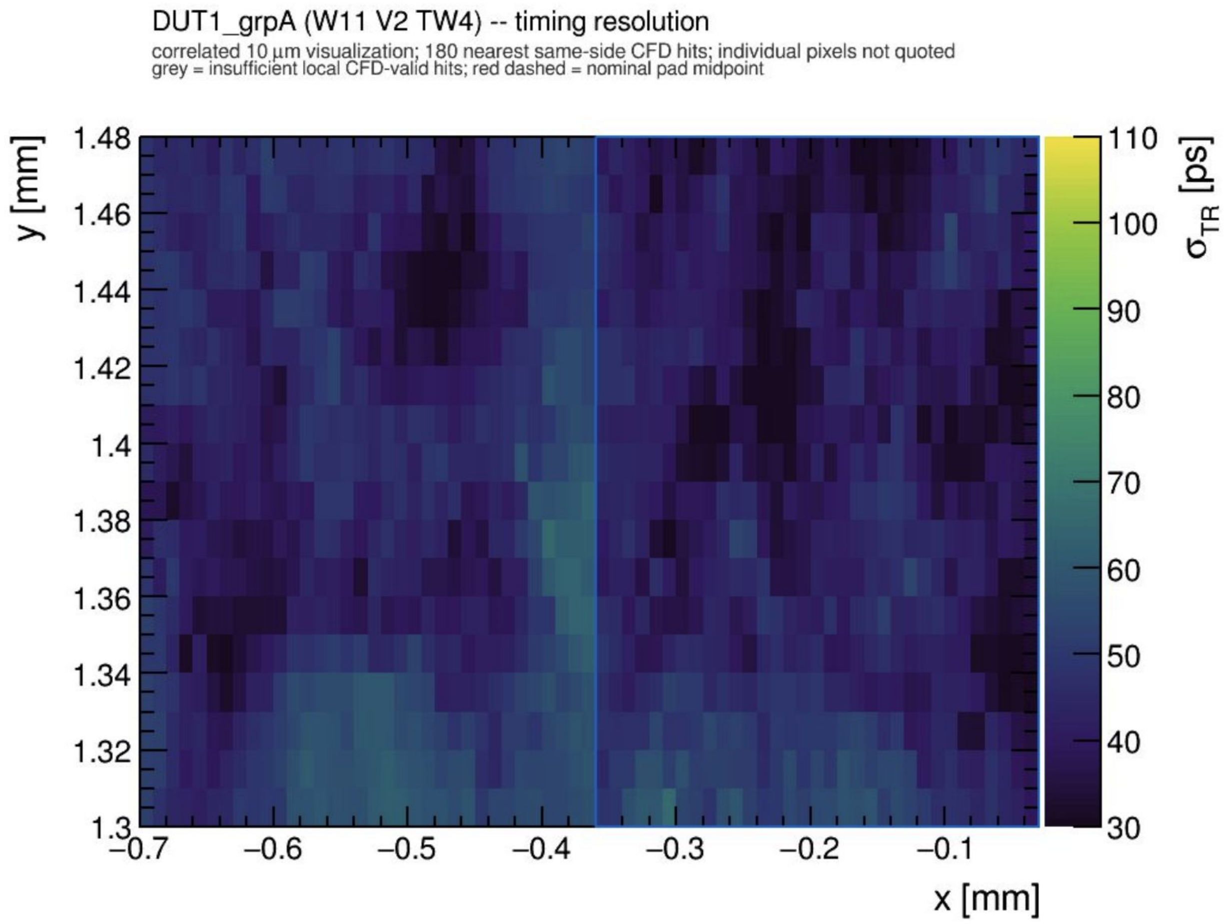

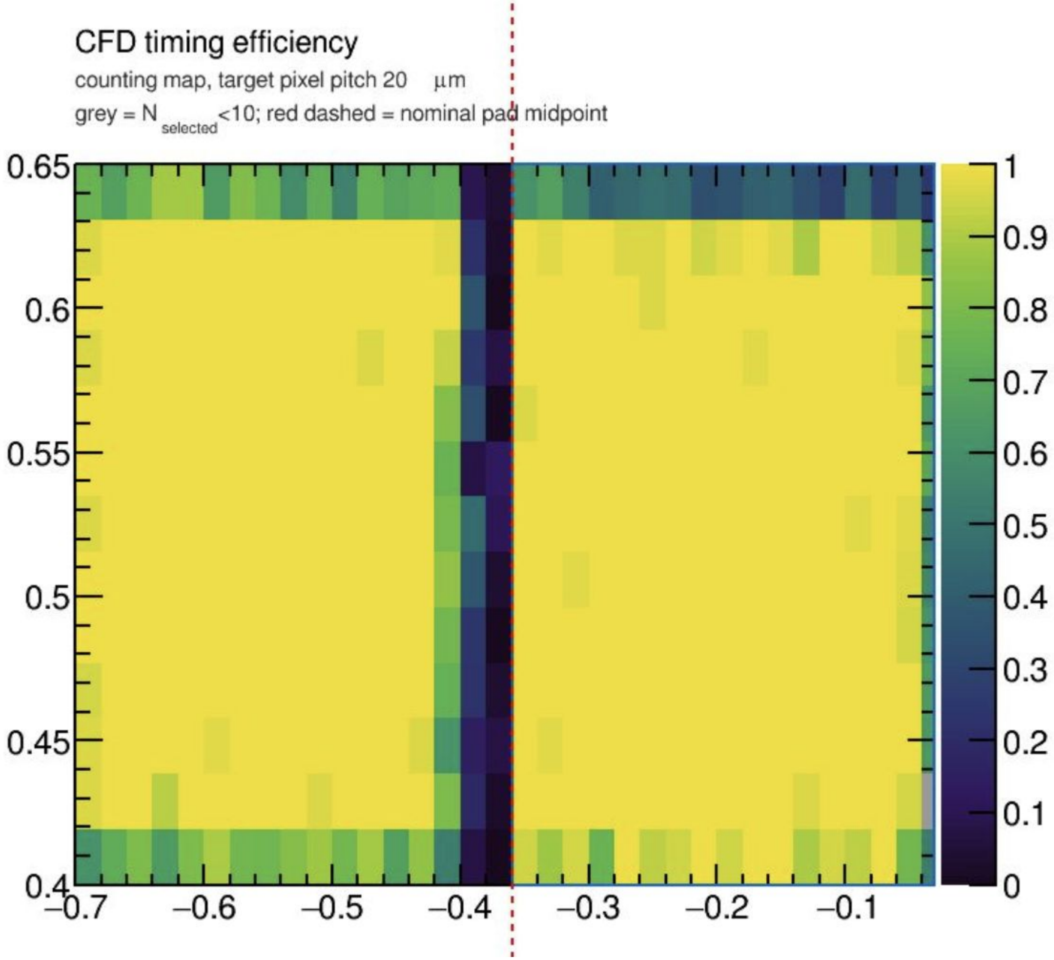

- **Bulk timing resolution: ~40–42 ps.**
- **CFD efficiency decreases locally in the inter-pixel region, accompanied by a degradation of σ_t.**
- **Away from the inter-pixel region, timing recovers to the clean-bulk level.**

Results for wafer 5 in the backup slides

Charge Spectrum

Kernel Density Estimate approach (Close to the inter-pixel region)

- Gaussian kernel density estimate over the charge sample.
- Bandwidth estimator: Silverman's rule × $n^{-0.2}$, with a grid search + parabolic refinement for the peak.
- resampling gives the uncertainty.
- Used for: the fine 5 µm spatial profile and the 2D local maps, where a bin's population is often modest.

Landau ⊗ Gaussian fit

- Explicit physical model: Landau energy-loss fluctuation convolved with Gaussian electronic noise.
- Returns MPV, Landau width, Gaussian σ, and χ^2/ndf — with real error bars.
- Accepted only if the fit converges with good covariance and χ^2/ndf < 10; otherwise the KDE mode is reported instead
- Used for: the bulk spectra — this is the primary quoted number for each pad's undisturbed response.

$$Q[\text{fC}] = \frac{V_{\text{peak}}[\text{mV}]}{5.67},$$

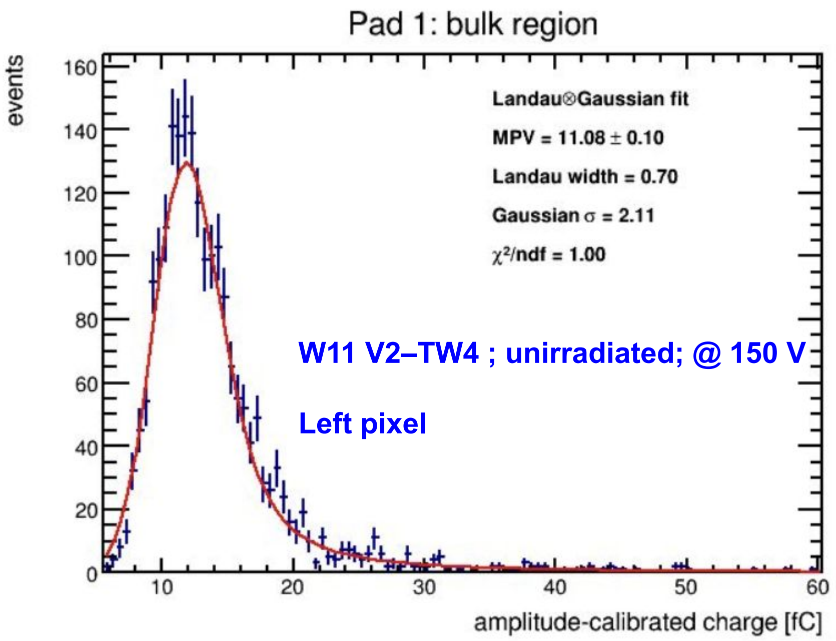

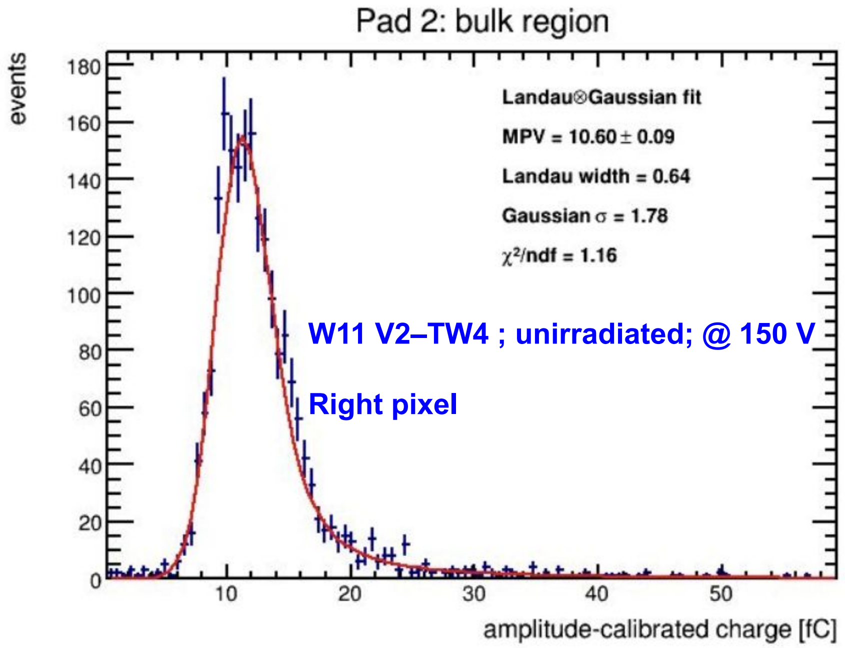

Errors on MPV are statistical; common amplitude-calibration uncertainty: 4.4%

One Charge Sample, Both Pads

- For every accepted track, BOTH neighboring pads are integrated — not just the one geometrically “hit.”
- This is what makes charge sharing measurable: an event near the gap leaves a real, comparable signal on each side, not a single ambiguous number.
- Each sample stores: q0 / q1 in fC, the raw noise-only integral (a floor estimate), peak amplitude, and the short/long-window variants.

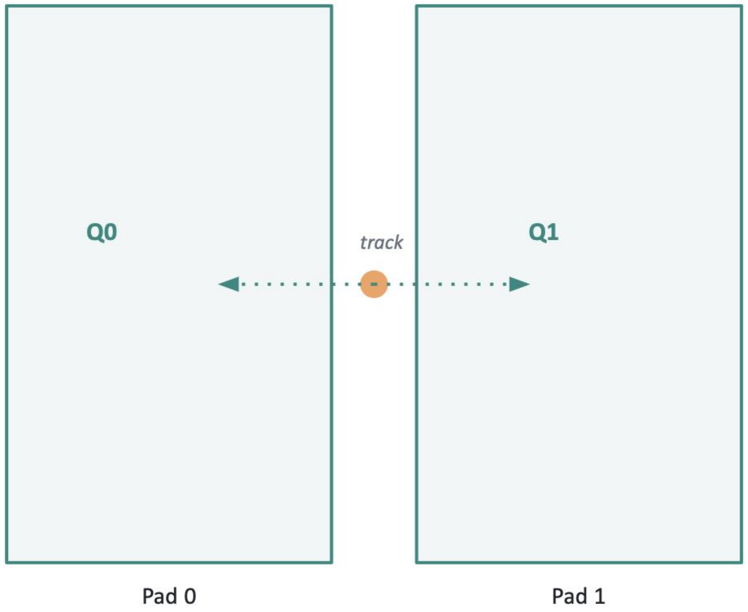

both integrated for every track, regardless of which pad it geometrically hit

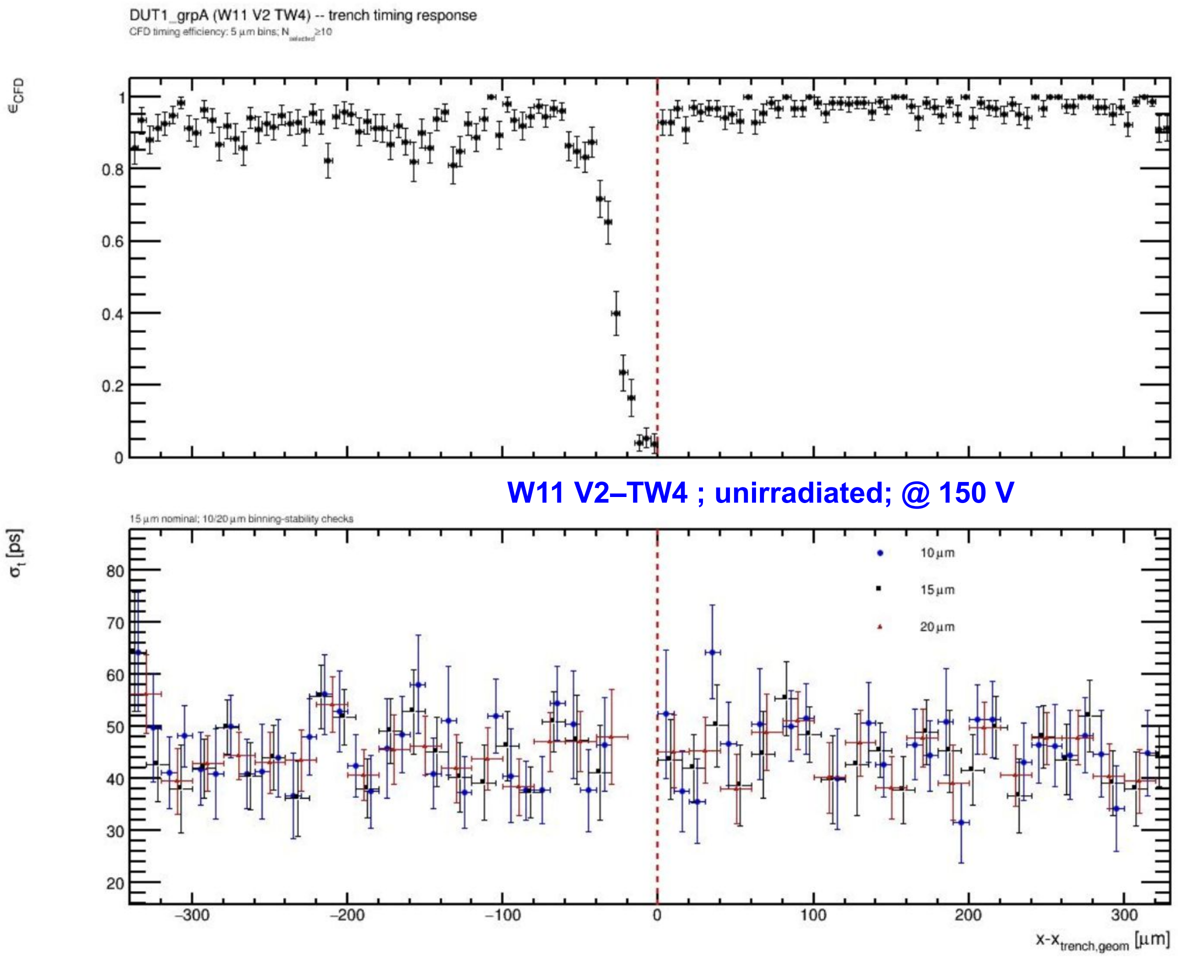

Bulk Reference & Normalized Observables

- "Bulk" = a fixed band, 0.10–0.23 mm from the pad's own inner edge, central-y only — safely away from both the gap and the outer edge, the same logic as the timing bulk-ROI.
- Each pad's own bulk Landau⊗Gaussian MPV becomes that pad's normalization scale.
- Every event's charge is then expressed as a fraction of its own pad's bulk MPV: $a = Q_0 / MPV_0$, $c = Q_1 / MPV_1$.

Normalization per pad

- Removes pad-to-pad gain/response differences before comparing charge sharing
- Makes Q_sum,norm and η insensitive to the absolute-fC calibration
- Every downstream spatial product (profile, 2D maps) is built on TOP of these two normalized numbers, plus the plain assigned charge in fC.

$$Q\_sum,norm = a + c$$

≈1 in the bulk; drops if charge is truly lost crossing the gap

$$\eta = (c - a) / (a + c)$$

−1 = all charge in pad 0, +1 = all charge in pad 1

Fine 2D Maps: Local Neighborhoods

- Query grid: 5 µm (x) × 10 µm (y)
- Each query point's value = KDE mode of its 60 nearest same-side charge samples — not a fixed spatial bin.

- Two safeguards against over-interpreting sparse regions: the 60th-nearest neighbor must be within 50 µm, AND the closest real track within 12 µm.

- Every physical map (local MPV, Q_sum,norm, η) ships with its own explicit validity mask, plus a separate "effective radius" map — so a genuinely low charge near zero is never confused with "no data here."

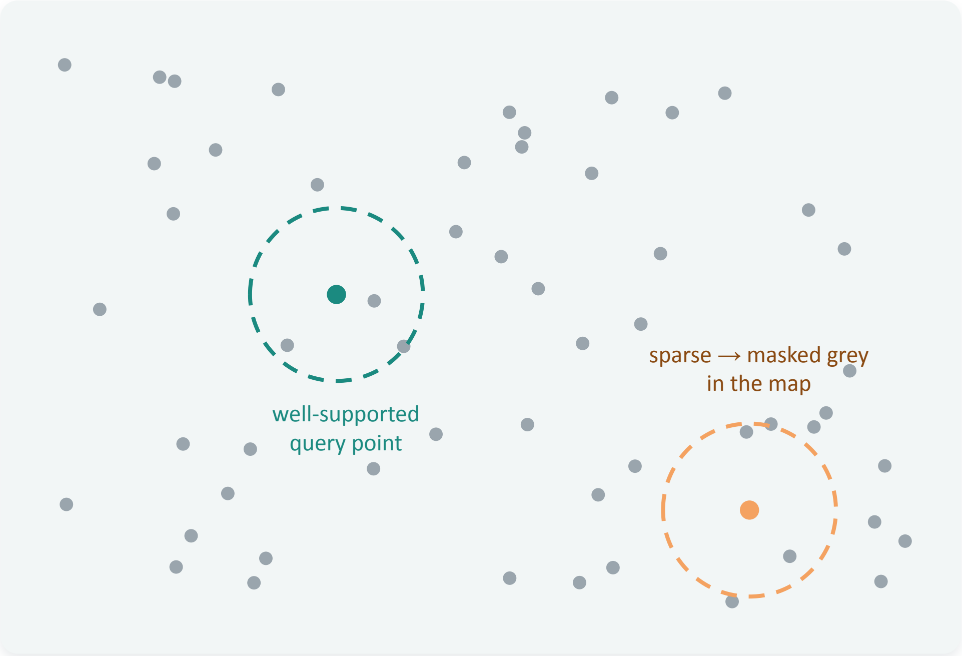

2D charge response and charge sharing — W11 V2–TW4, 150 V

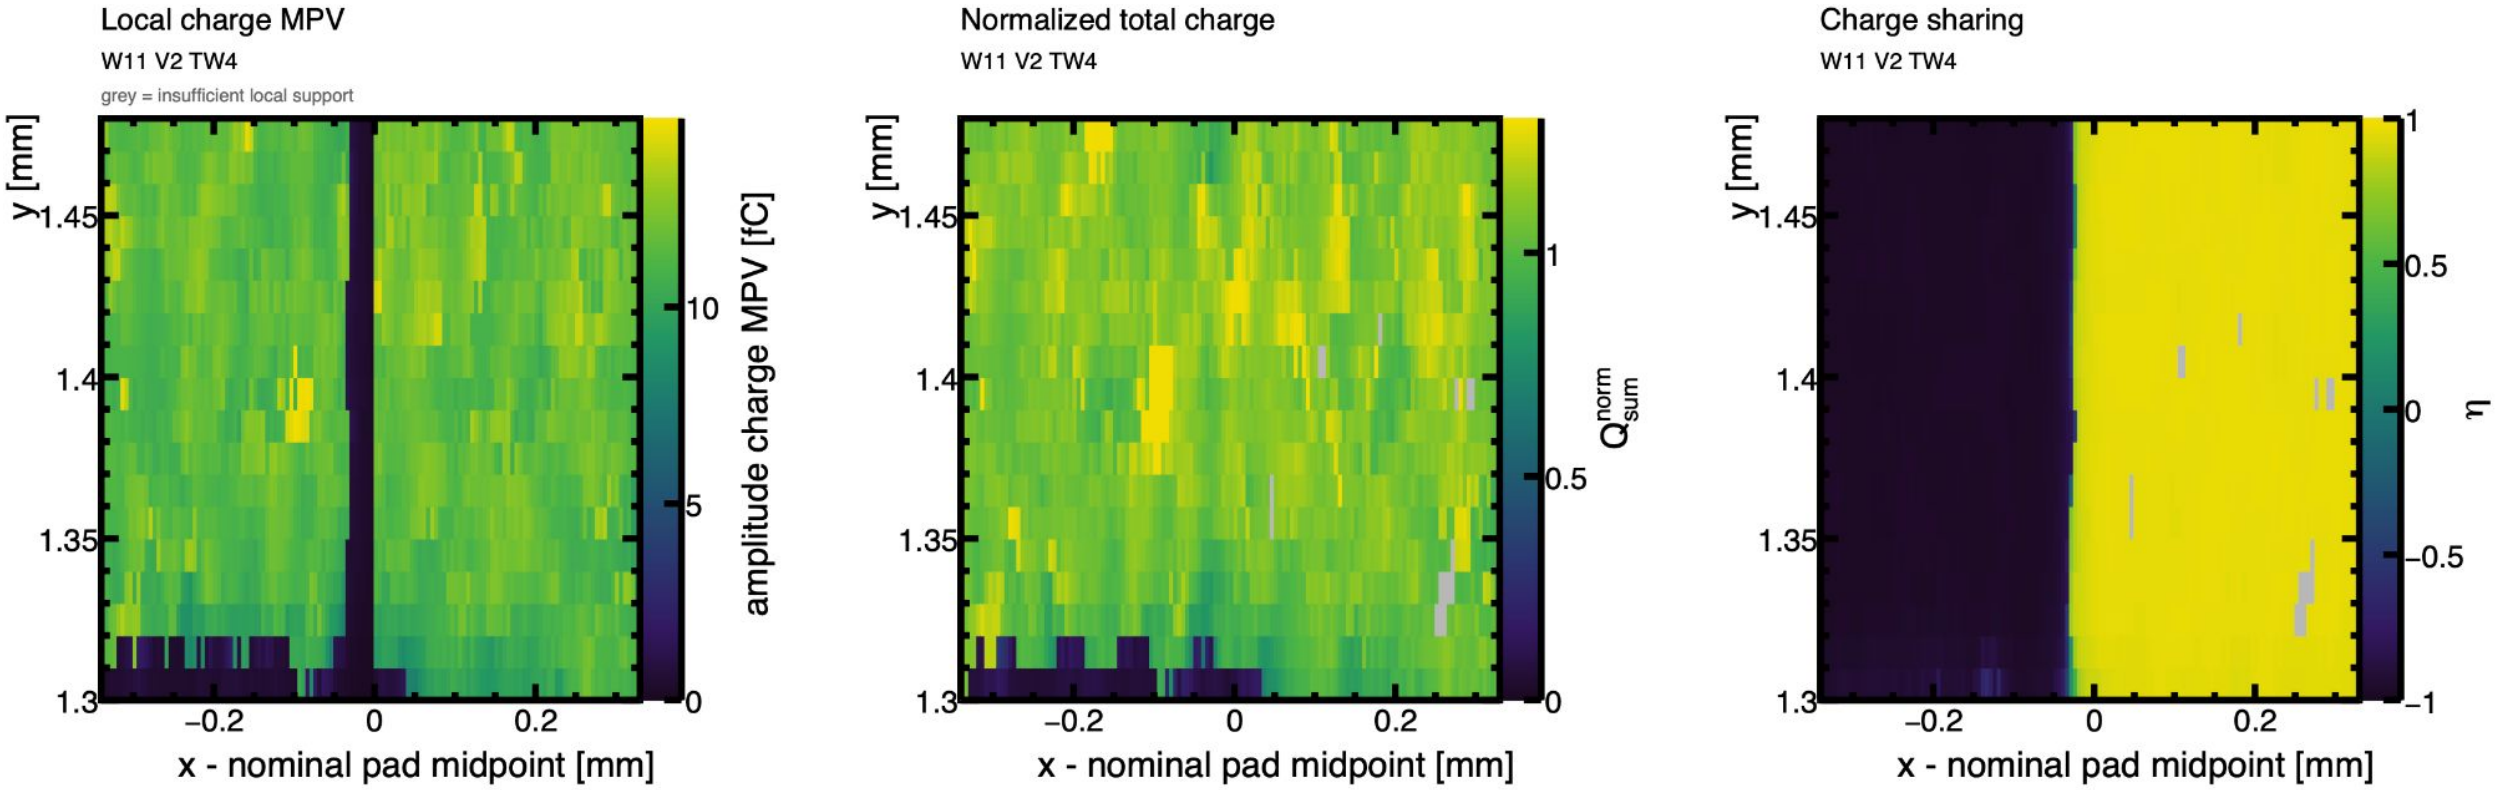

- **The inter-pixel structure is localized in 2D.**
- **The summed response and sharing observable show the same spatial transition.**
- **The 2D result is consistent with the 1D charge-sharing profile.**

Charge sharing across the inter-pixel region W11 V2–TW4, 150 V

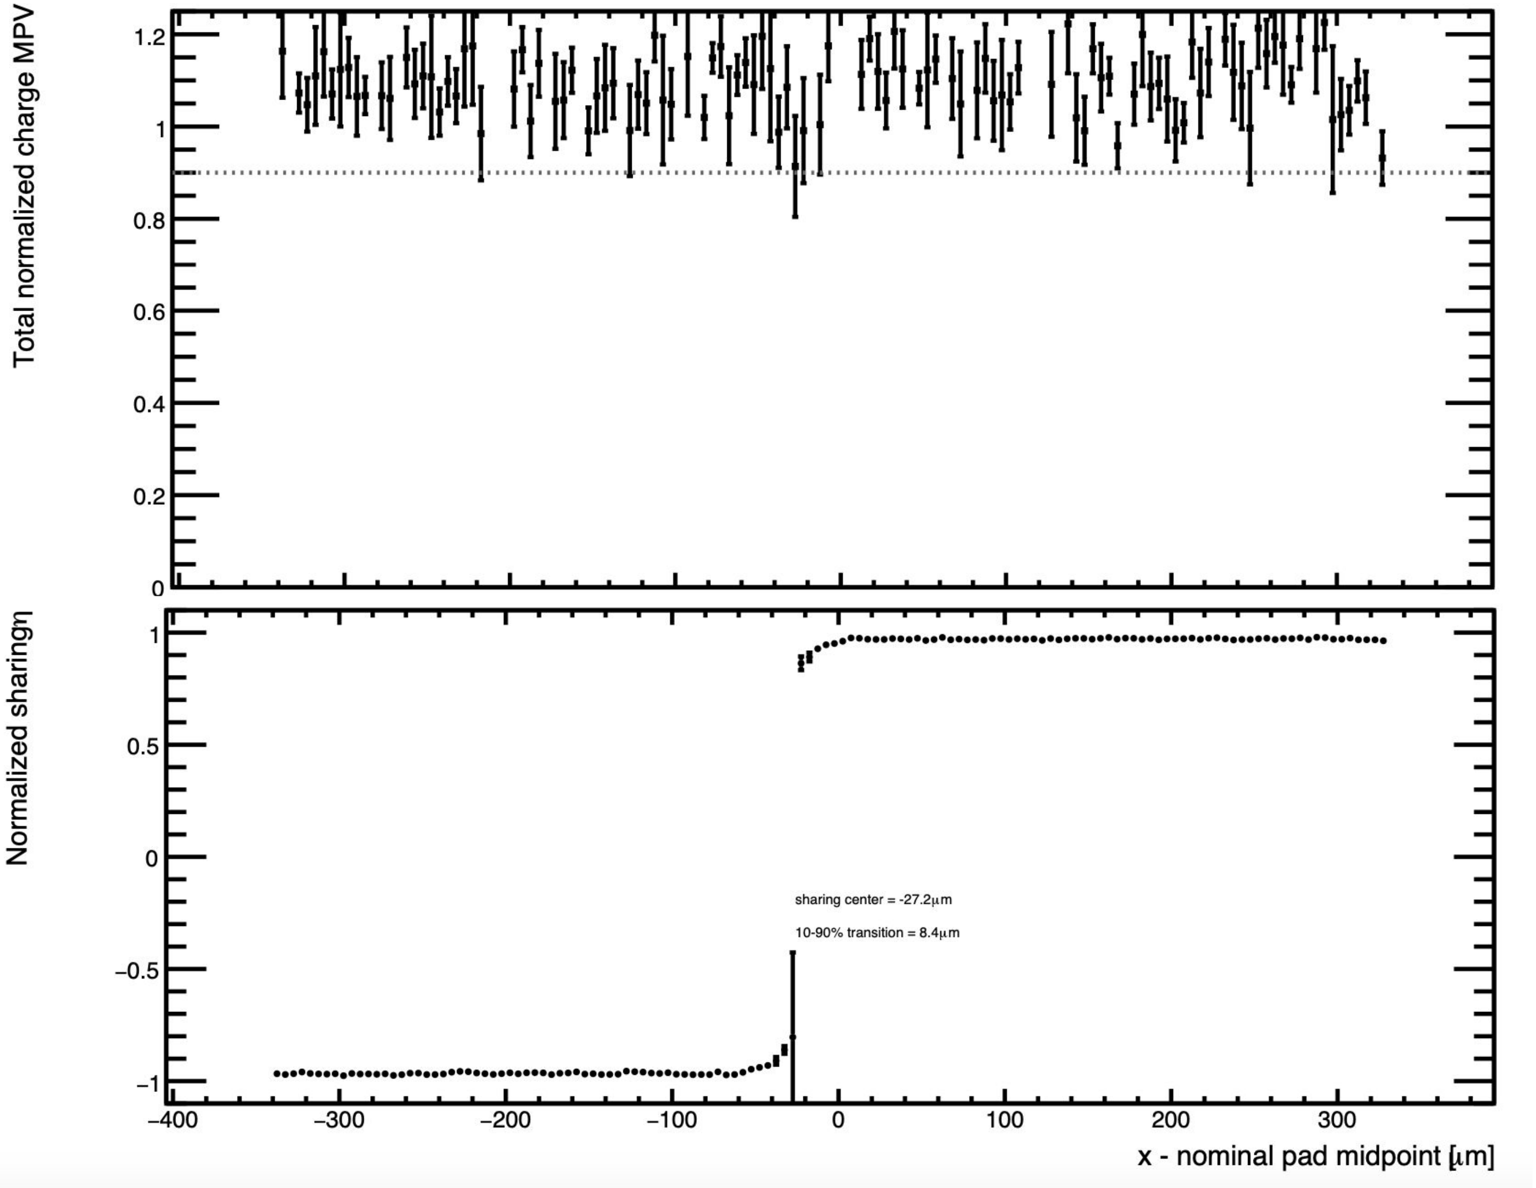

- The summed two-pad response is approximately preserved across most of the transition region.
- The signal transfers sharply from one neighboring pixel to the other.
- Charge-sharing transition width (10–90%): ~8.4 µm.

Timing and charge resolve the same inter-pixel structure — W11 V2–TW4, 150 V

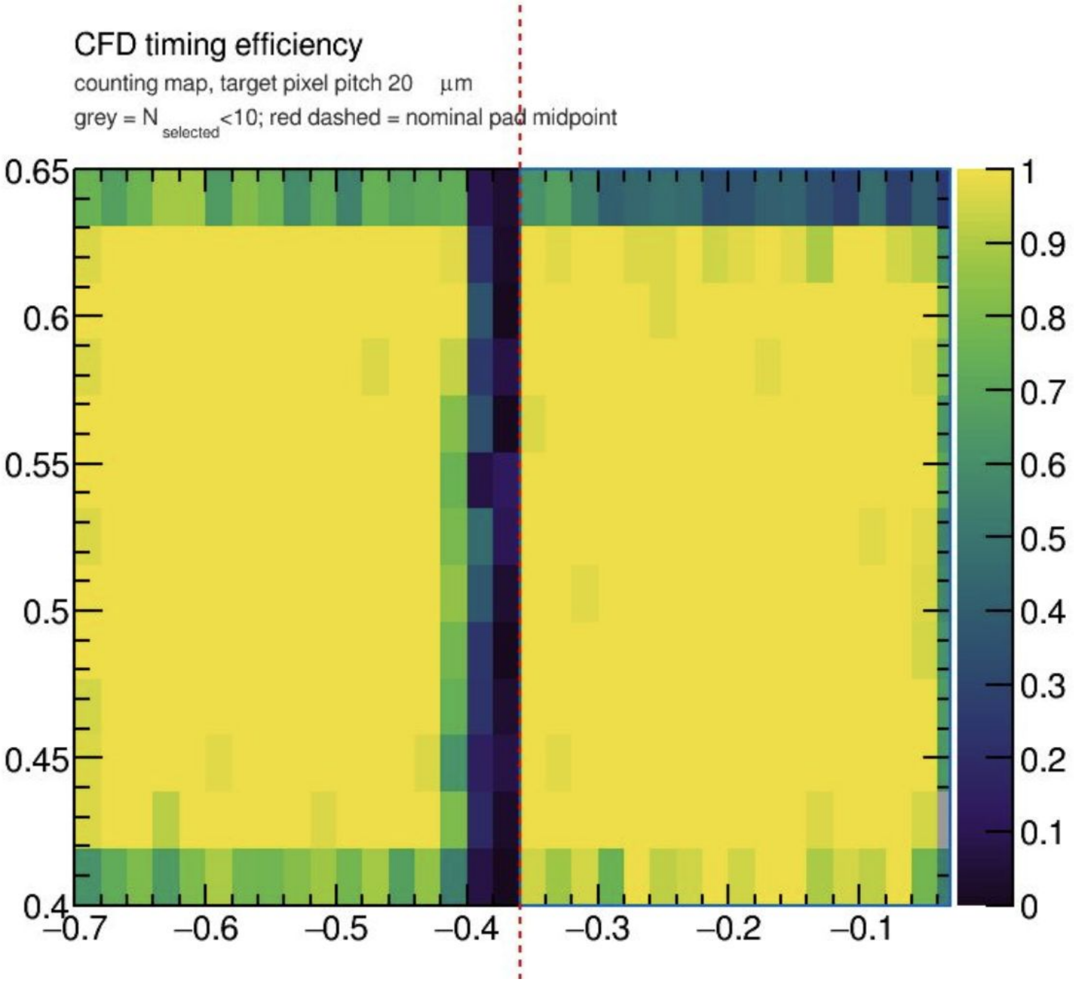

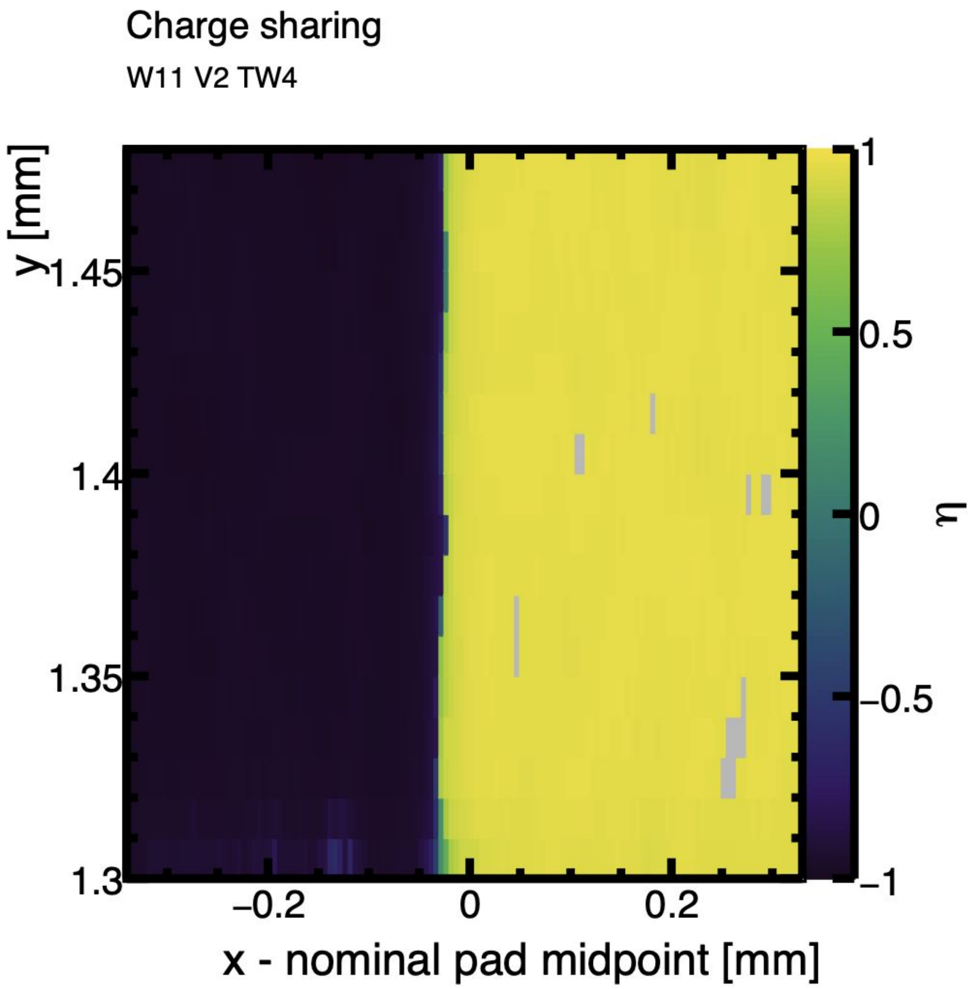

- **Bulk timing remains at ~40–42 ps away from the inter-pixel region.**
- **The probability of obtaining a CFD timestamp drops locally where the signal is shared between neighboring pixels.**
- **Timing and charge therefore provide complementary views of the same inter-pixel response.**

Spatial timing response — DUT2 W5 V2 TW 2/3, $2.5{\times}10^{15} n_{eq}/cm^2$

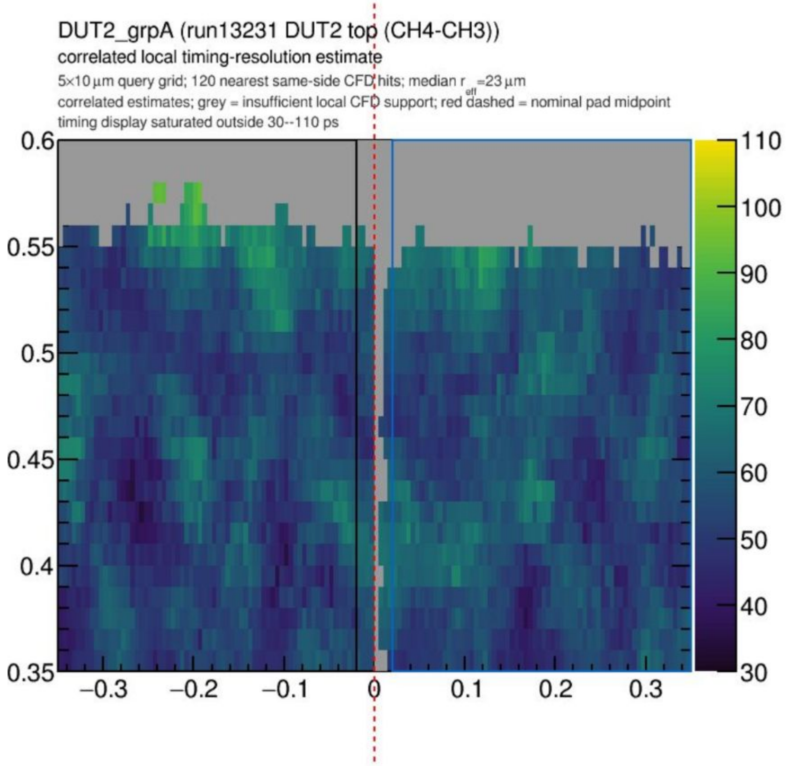

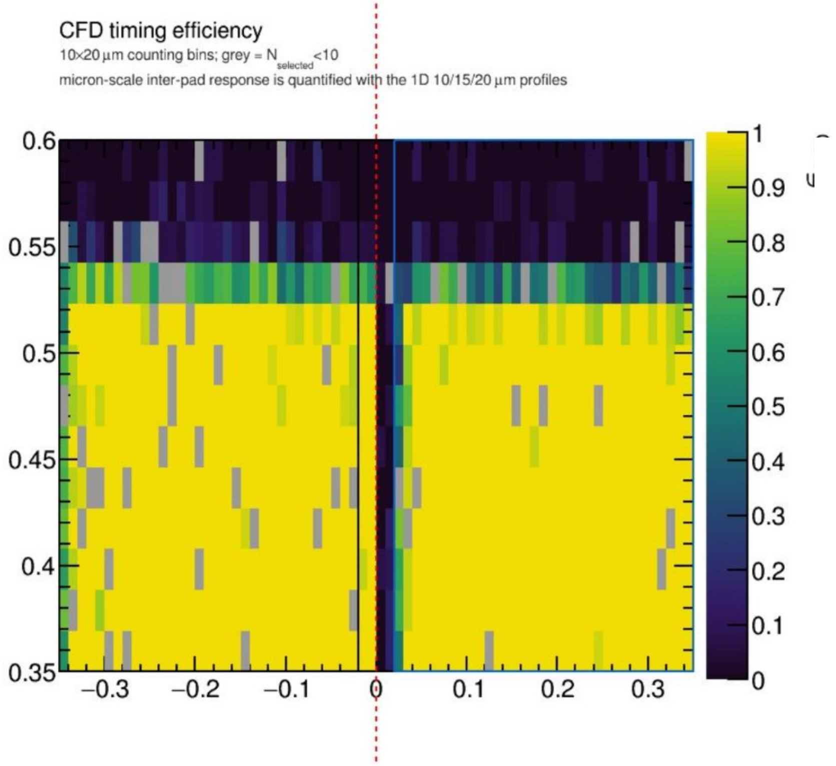

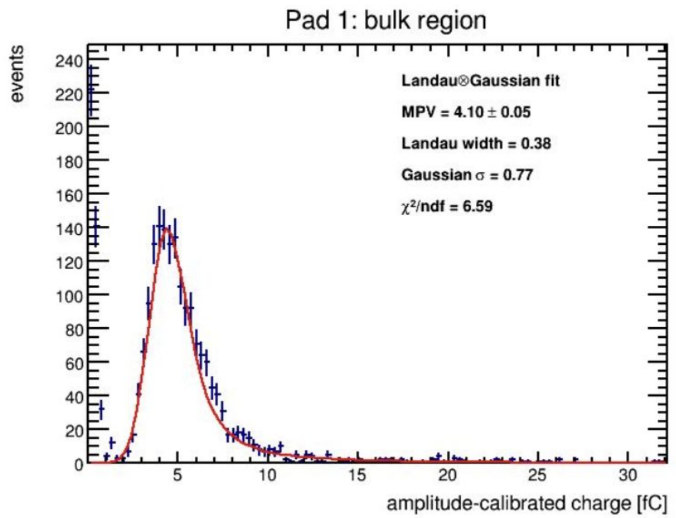

- **Bulk timing: ~55–60 ps at $2.5{\times}10^{15}$ n_{eq}/cm^2.**
- **Timing resolution degrades locally at the inter-pixel region.**
- **The degradation coincides with a localized loss of CFD timing efficiency.**

Spatial charge response after irradiation — W5 V2 TW4, $1.5\times10^{15} n_{eq}/cm^2$

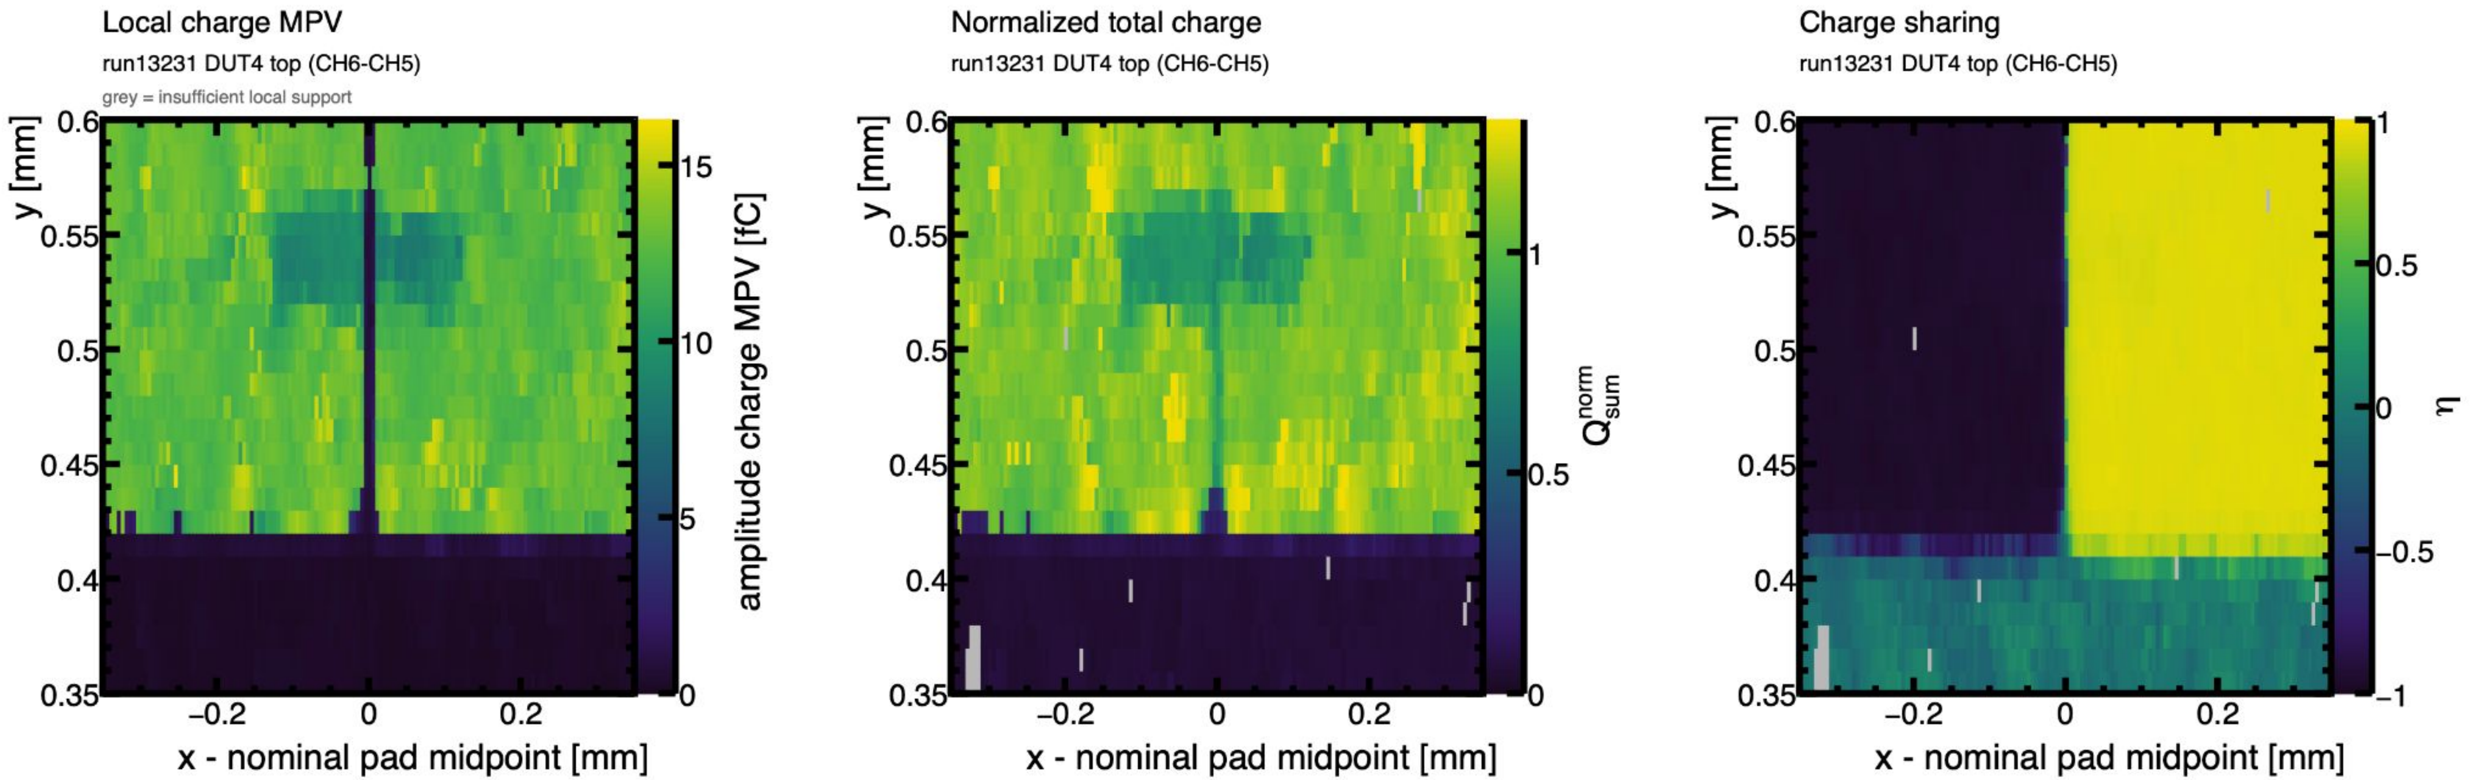

- **A clear charge-sharing transition remains visible after irradiation.**
- **The signal transfers between the two neighboring pads across the inter-pixel region.**
- **The normalized total response and sharing observable remain spatially localized around the transition.**

Inter-pixel timing and charge sharing in W11 and W5 TI-LGADs

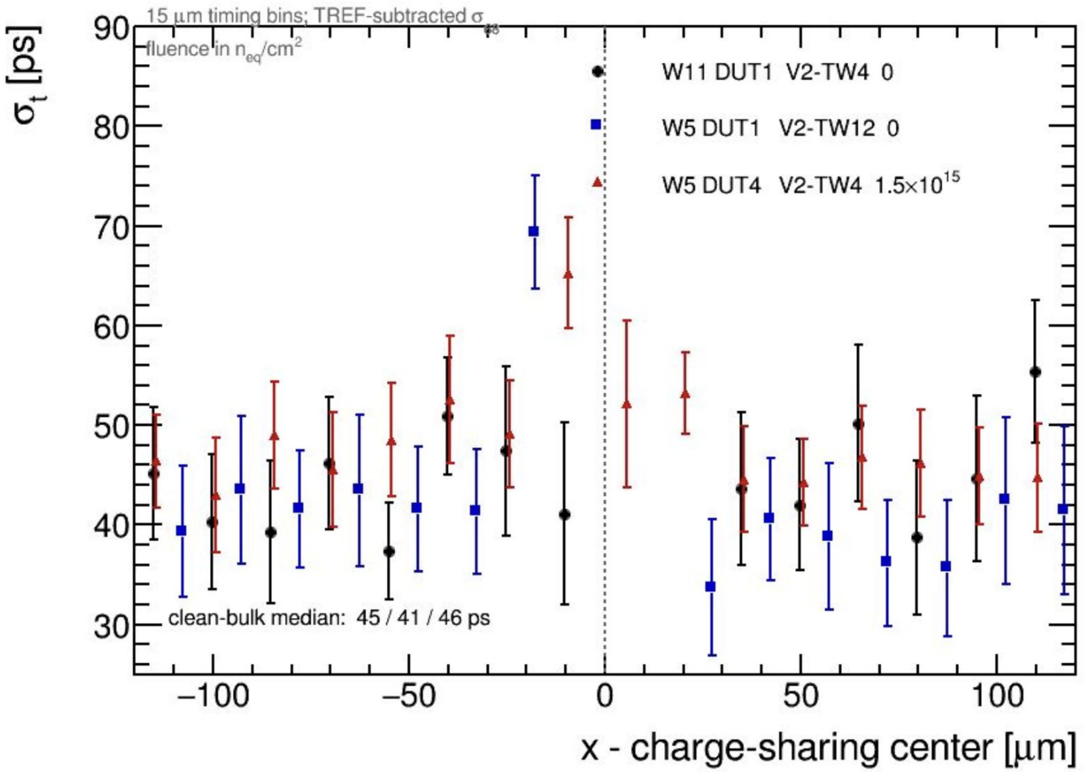

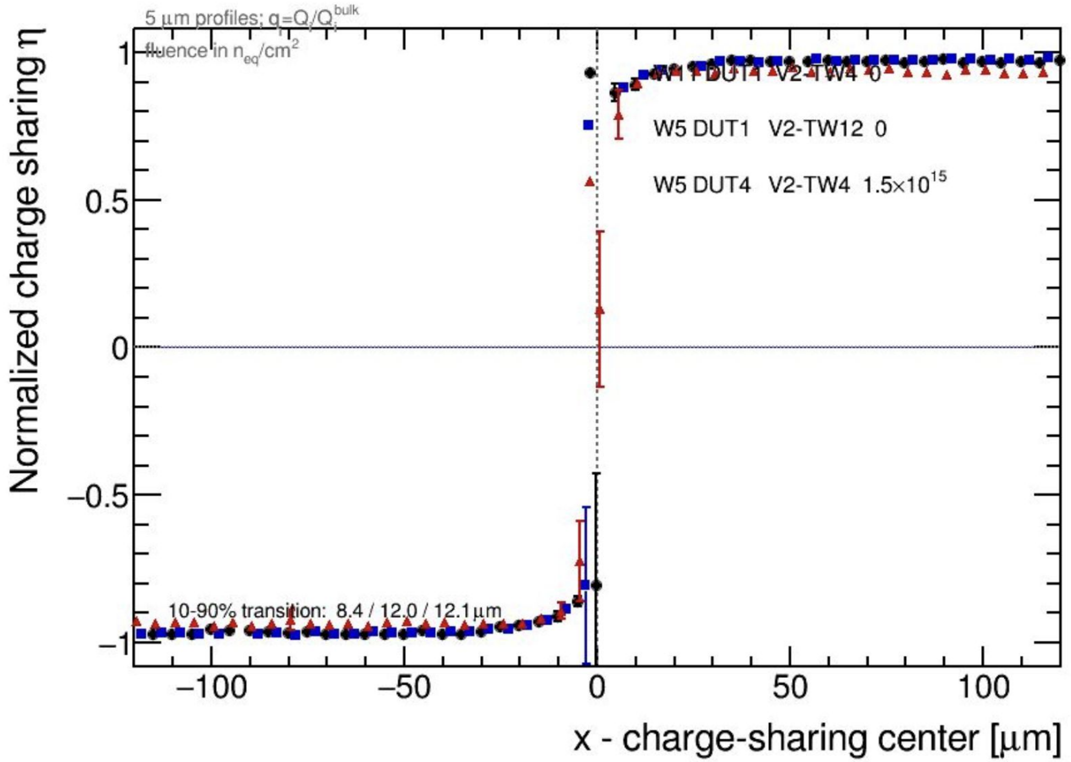

- **A narrow charge-sharing transition is observed in all three devices:** 8.4 μm for W11 and ~12 μm for both W5 devices.
- **The timing degradation is localized around the charge-sharing region; away from it, the bulk timing is recovered.**
- **The irradiated W5 DUT4 retains a sharp inter-pixel transition and ~46 ps bulk timing at 1.5×10^{15} n_{eq}/cm^2.**

BACKUP SLIDES

Converting to Absolute Charge (fC)

$$Q\ [fC] = V\_peak\ [mV]\ /\ 5.67$$

5.67 ± 0.25 mV/fC; peak found in a ±0.6 ns window anchored to TREF's predicted pulse center

- Charge is read from the pulse's peak amplitude, not the integrated area — searched only in a narrow ±3-sample (±0.6 ns) window centered on the position TREF predicts, so a noise spike elsewhere on the trace can't be mistaken for the pulse.
- The 5.67 ± 0.25 mV/fC scale is the same amplitude-based transimpedance constant independently published for this exact Chubut board (5.67×10^{12} V/C)
- This is still the only place an external, absolute calibration enters — and it now carries a quantified ±4.4% (0.25/5.67) common systematic on every absolute fC number, not folded into the per-fit statistical errors.

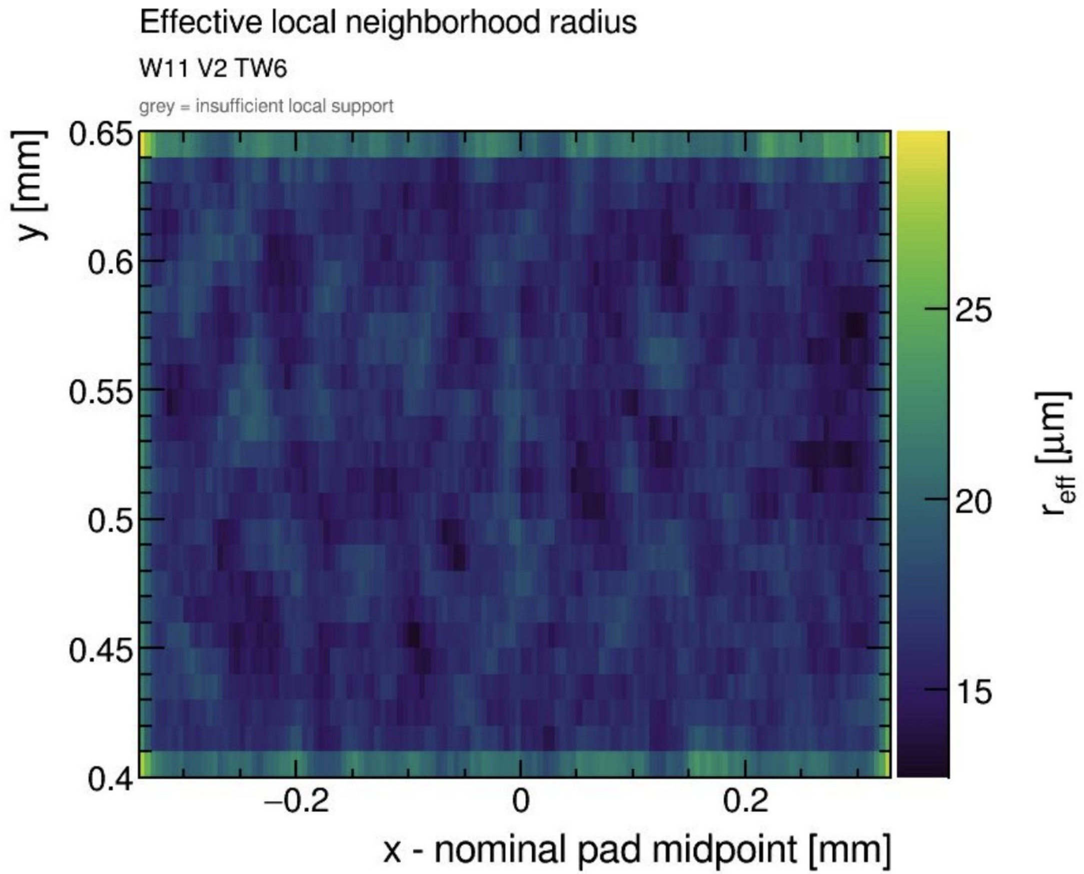

W11 V 3/4 TW5, unirradiated @120 V

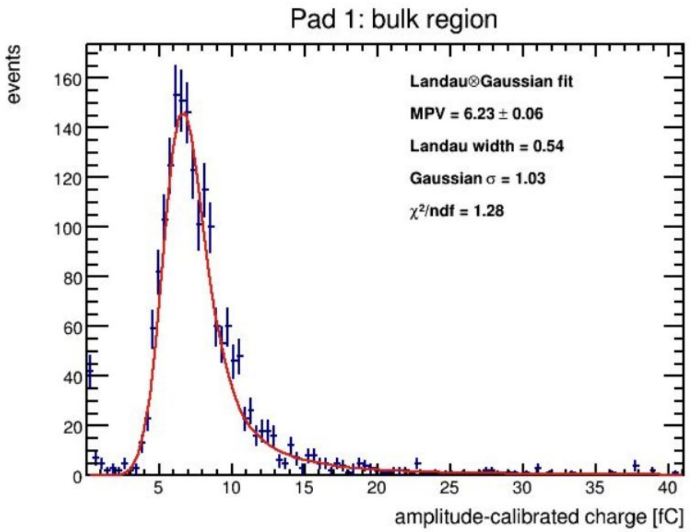

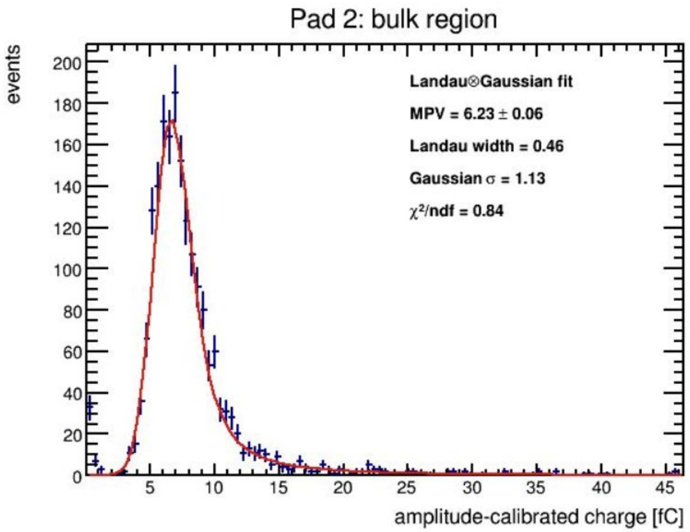

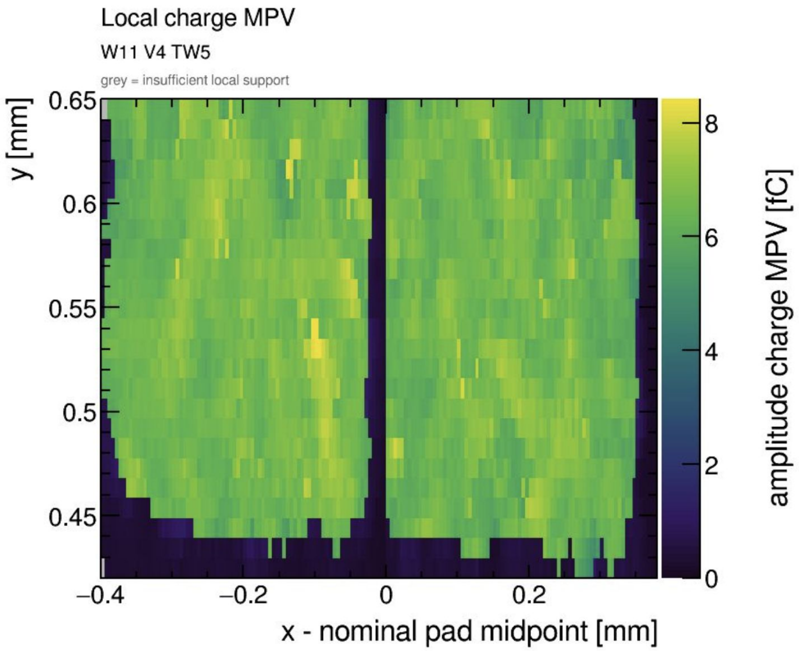

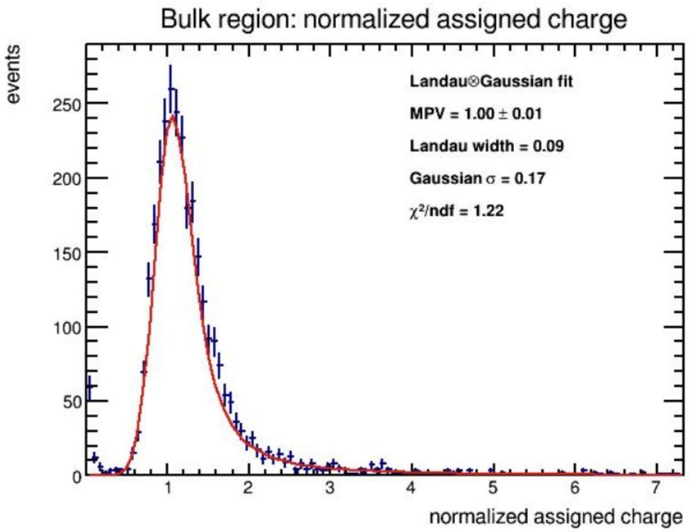

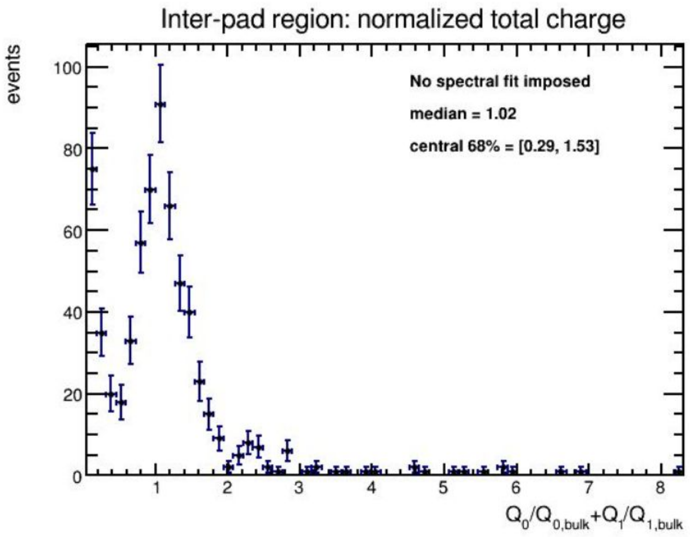

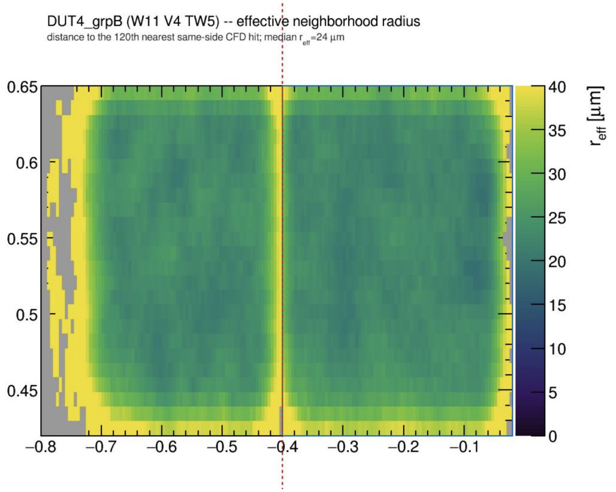

W11 V 3/4 TW5, unirradiated @120 V

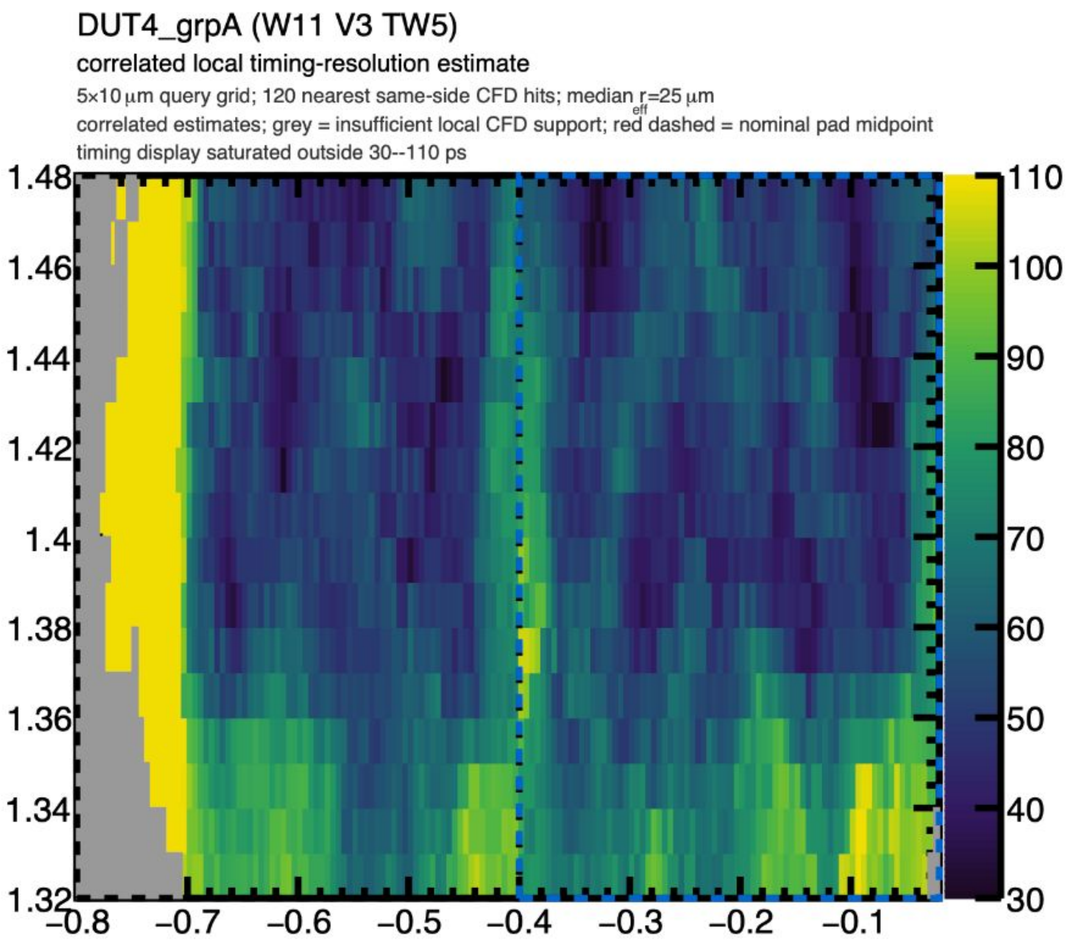

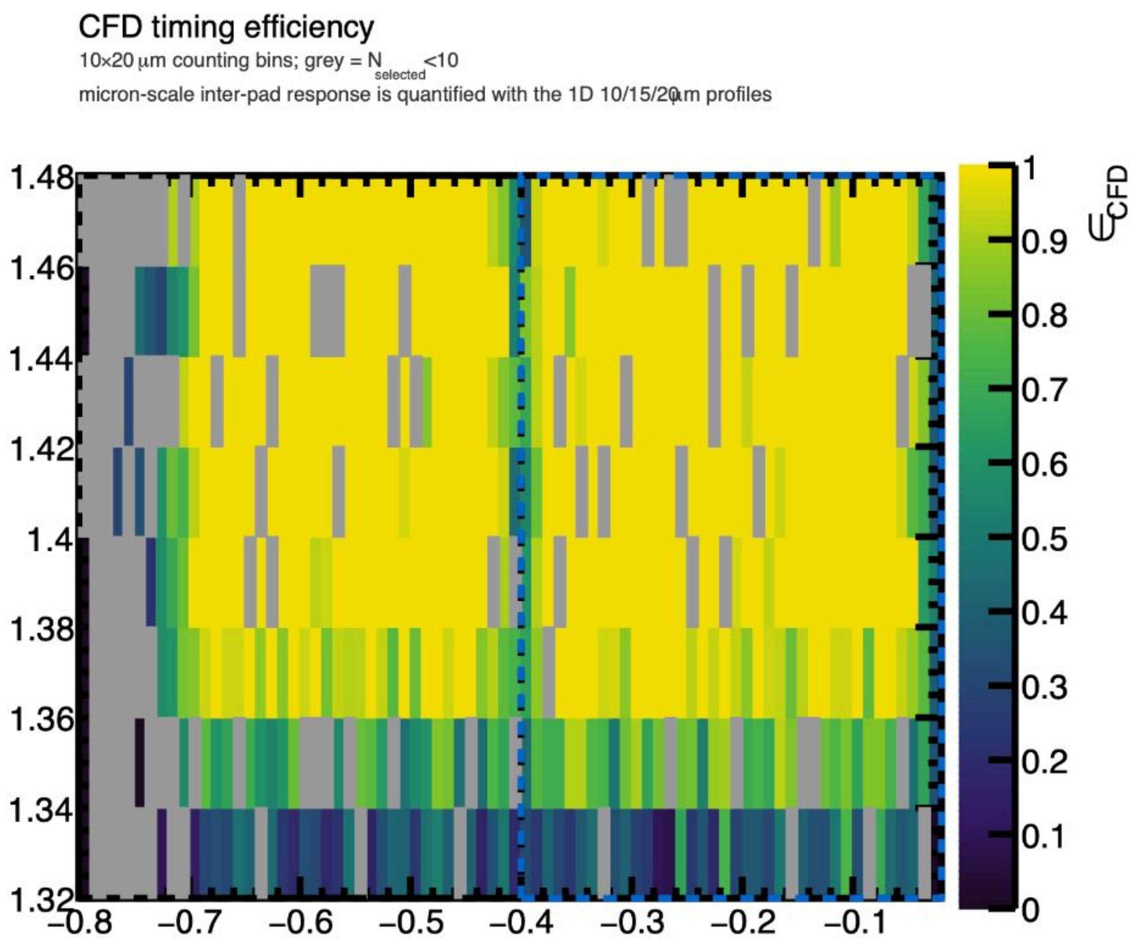

Spatial timing response and track coverage — W11 V2–TW6, 150 V

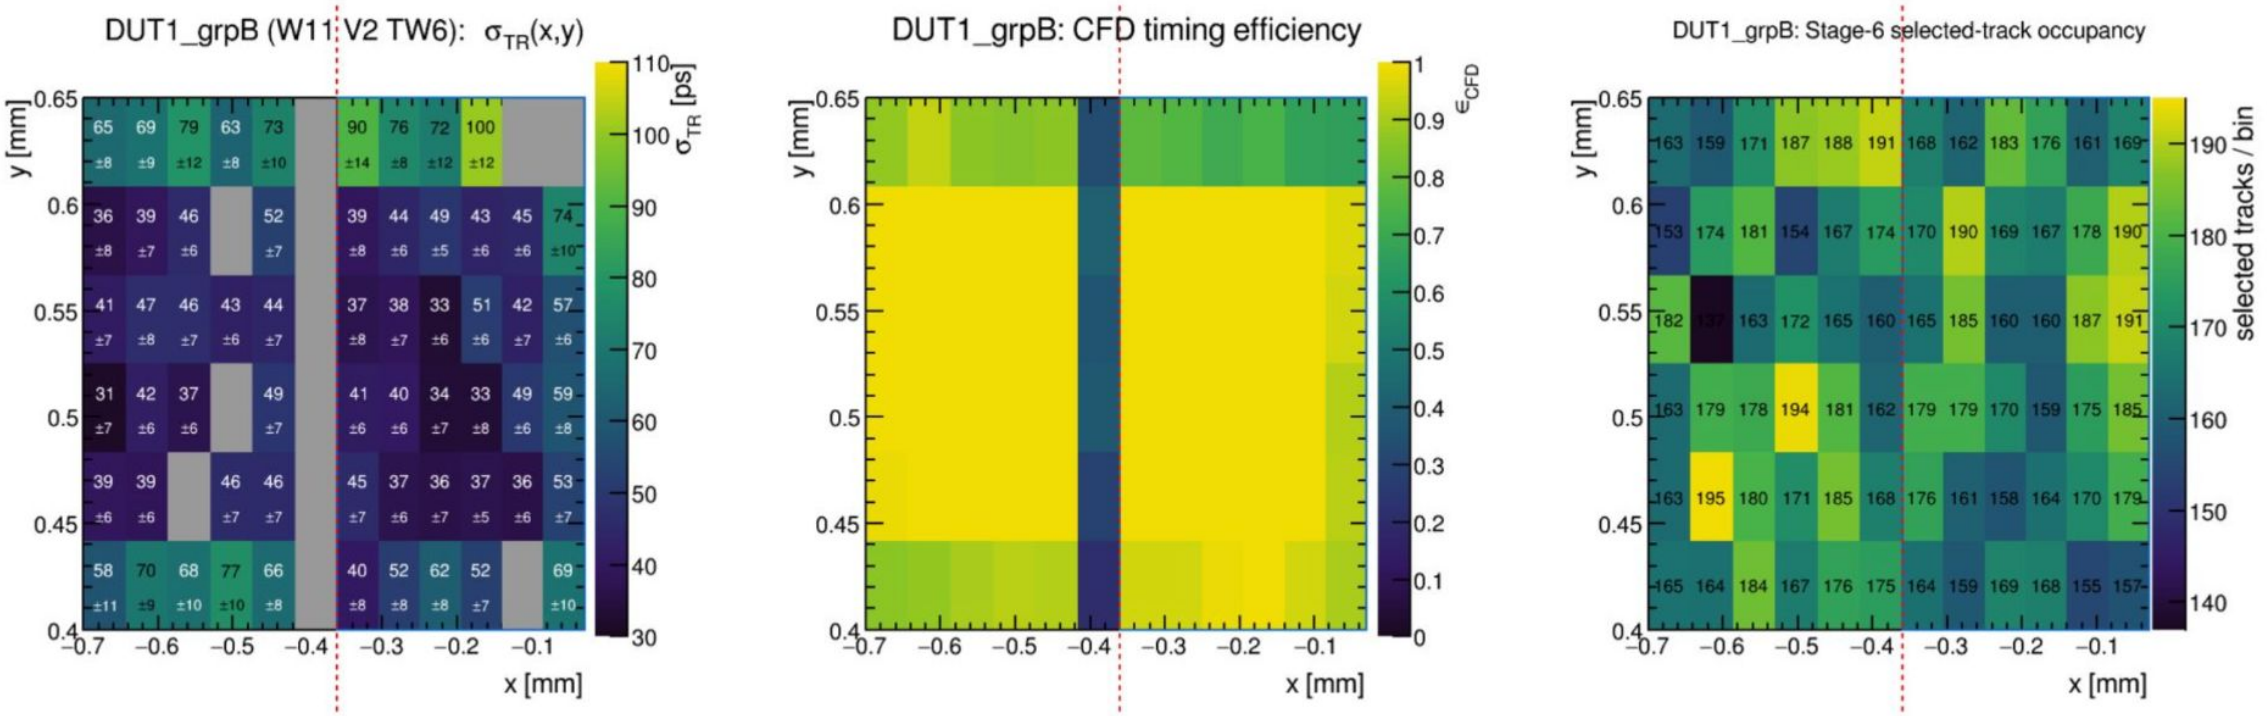

Electronic-jitter diagnostic vs CFD fraction — W11 V2–TW6

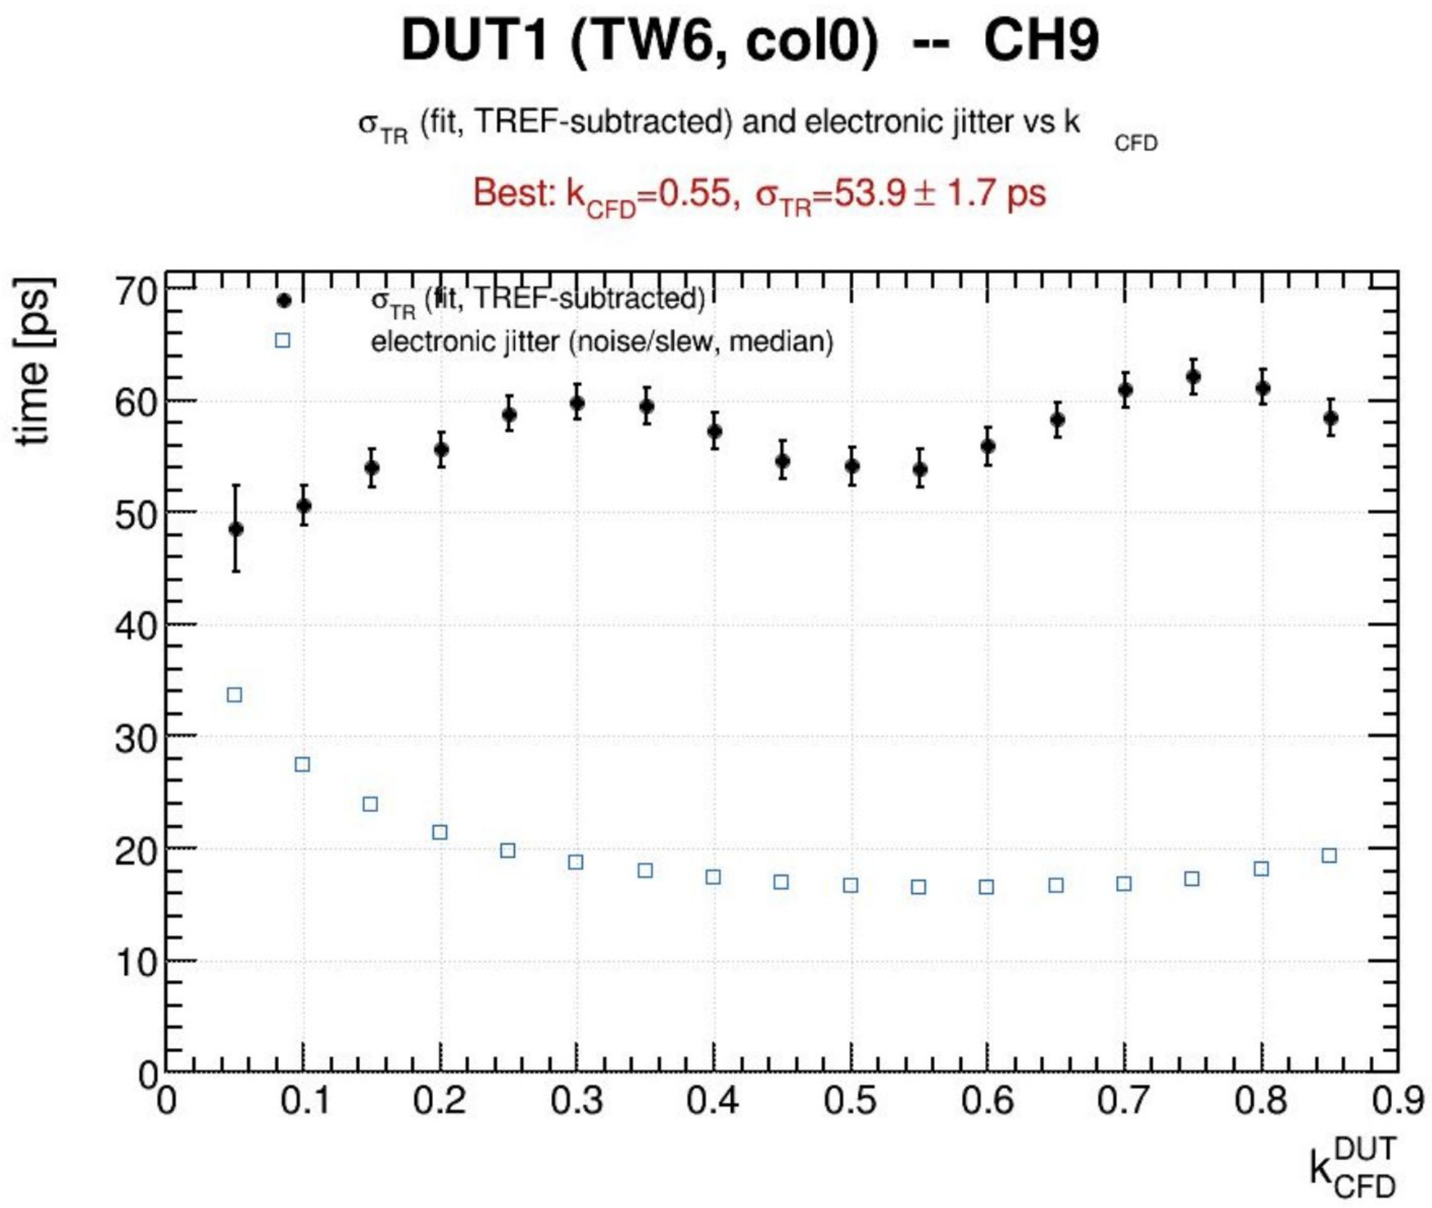

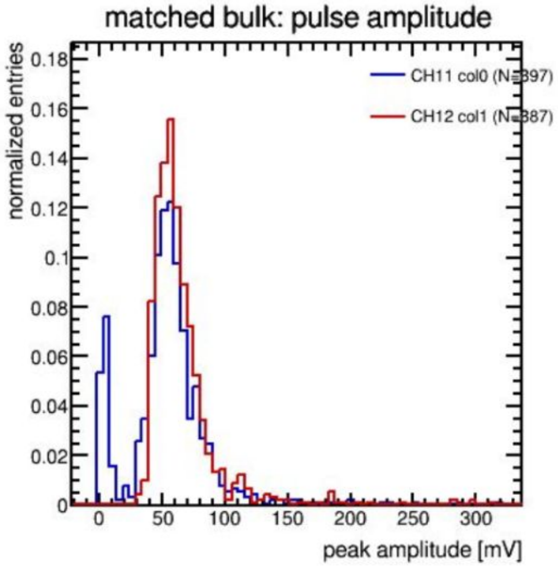

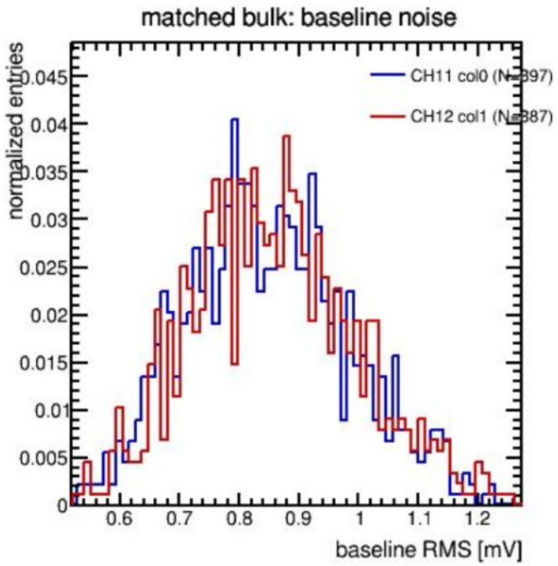

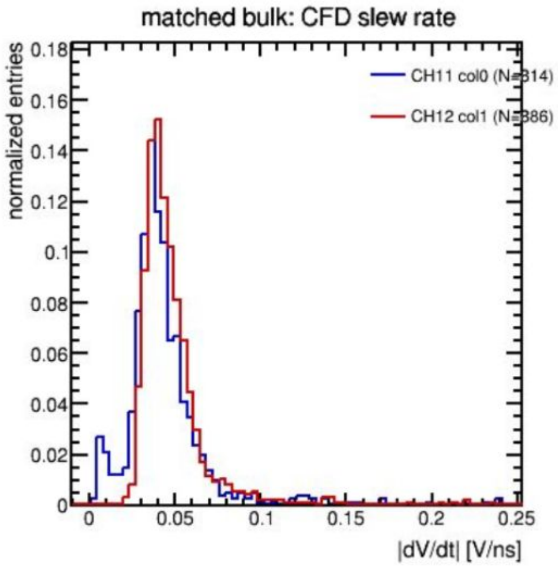

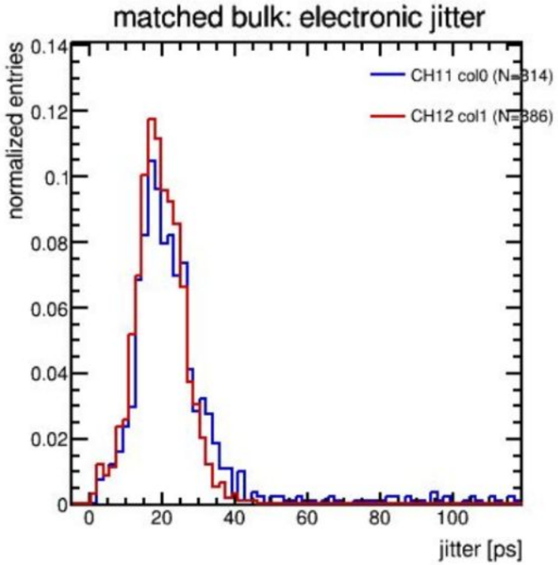

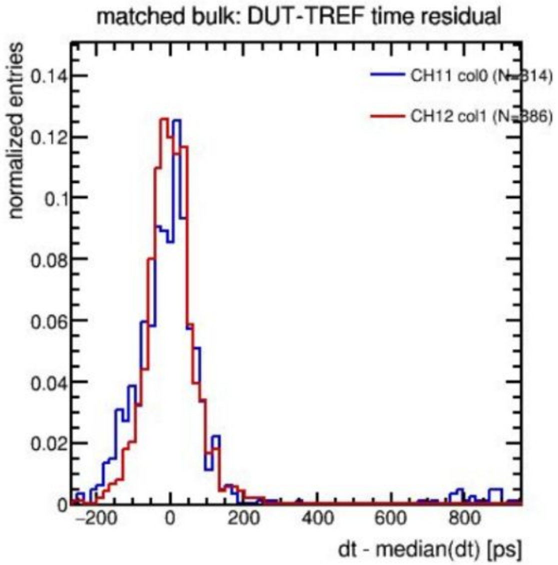

DUT2 group B = W11 TW2_V2

matched $|x-x_{trench}|$ = 100-170 μm

common y = 0.490-0.660 mm

CH11 col0: Nsel=897, NCFD=814, ε=0.907

median A/noise/slew/jitter = 54.28 mV / 0.845 mV / 0.040 V/ns / 21.0 ps

σ^{68}_{DUT} = 64.57 ± 4.01 ps

CH12 col1: Nsel=887, NCFD=886, ε=0.999

median A/noise/slew/jitter = 59.03 mV / 0.854 mV / 0.044 V/ns / 19.1 ps

σ^{68}_{DUT} = 39.14 ± 3.27 ps

DIAGNOSTIC ONLY: same frozen k_{CFD}, same Stage-1 TREF, same matched geometry

Use this to identify whether the CH11/CH12 timing asymmetry is pulse/electronics driver

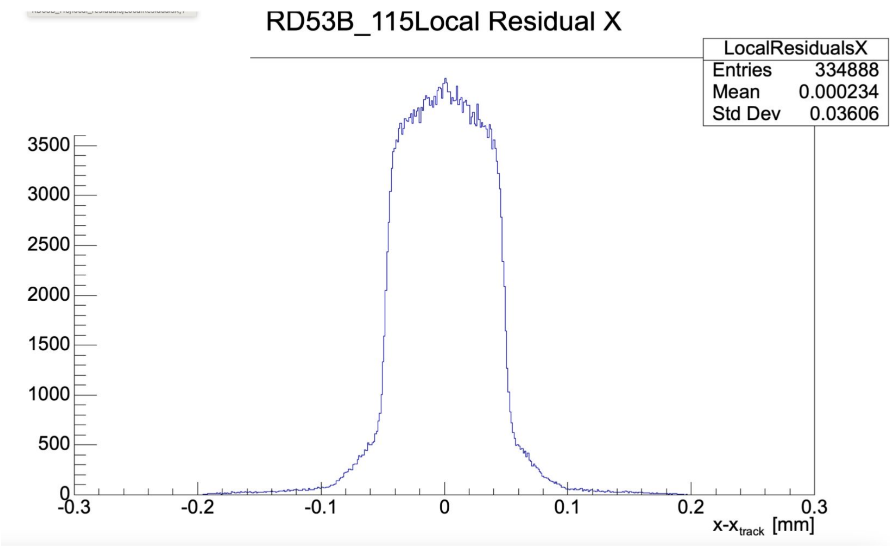
RD53B_115Local Residual X
LocalResidualsX
Entries 334888
Mean 0.000234
Std Dev 0.03606
x-x_track [mm]

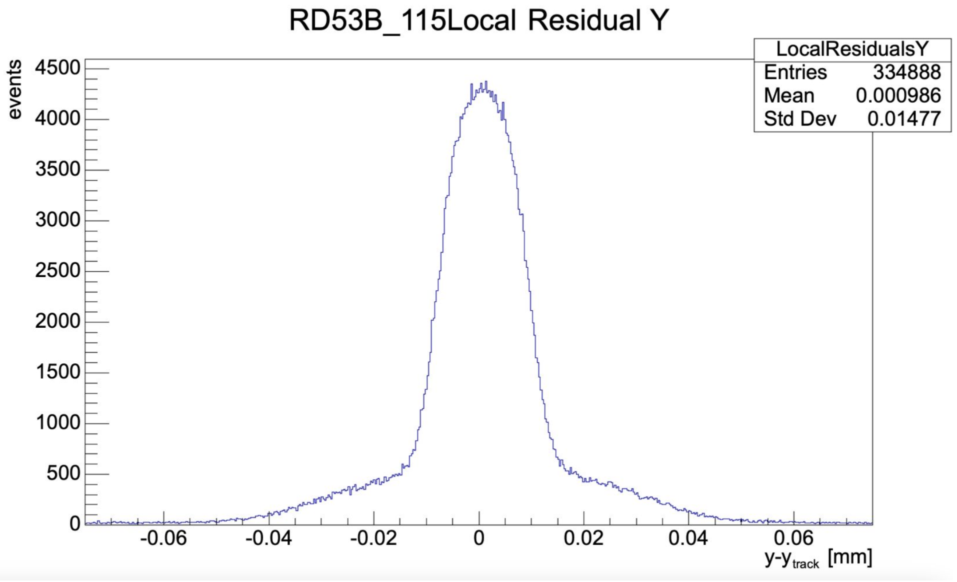
RD53B_115Local Residual Y
LocalResidualsY
Entries 334888
Mean 0.000986
Std Dev 0.01477
events
y-y_track [mm]

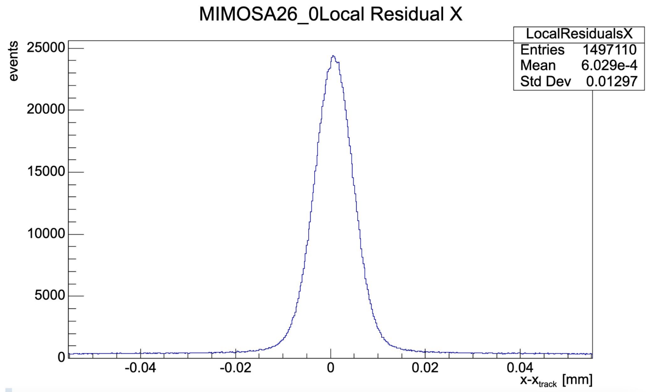
MIMOSA26_0Local Residual X
LocalResidualsX
Entries 1497110
Mean 6.029e-4
Std Dev 0.01297
events
x-x_track [mm]

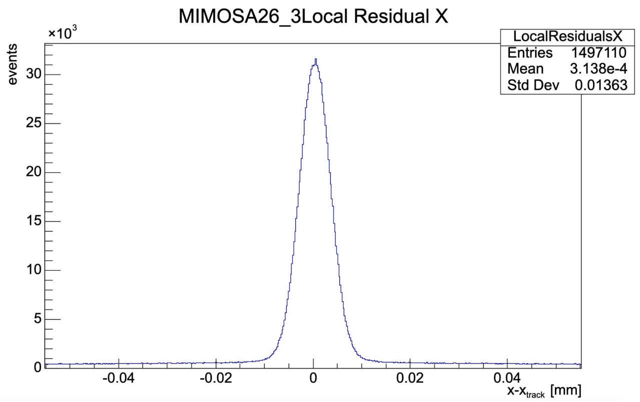
MIMOSA26_3Local Residual X
LocalResidualsX
Entries 1497110
Mean 3.138e-4
Std Dev 0.01363
events
×10^3
x-x_track [mm]

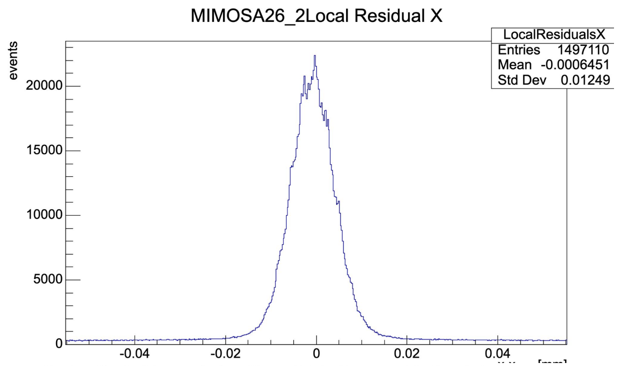
MIMOSA26_2Local Residual X
LocalResidualsX
Entries 1497110
Mean -0.0006451
Std Dev 0.01249
events

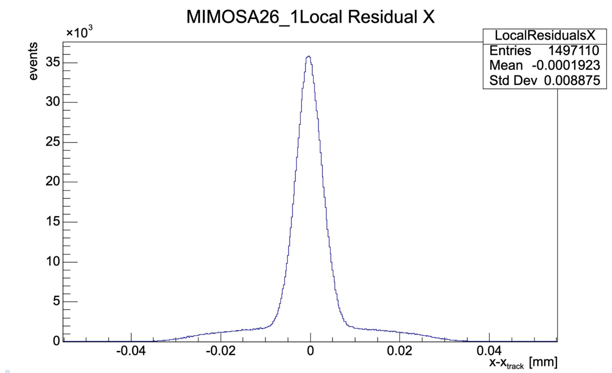
MIMOSA26_1Local Residual X
LocalResidualsX
Entries 1497110
Mean -0.0001923
Std Dev 0.008875
events
×10^3
x-x_track [mm]

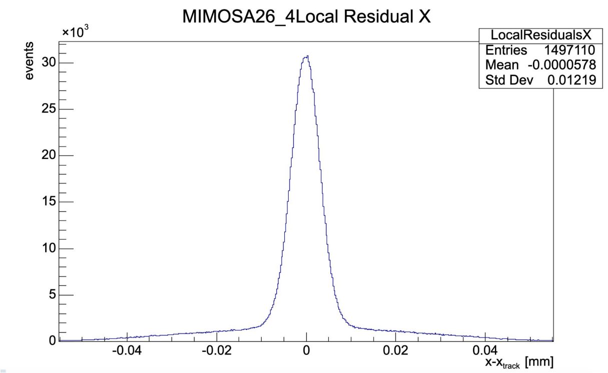
MIMOSA26_4Local Residual X
LocalResidualsX
Entries 1497110
Mean -0.0000578
Std Dev 0.01219
events
×10^3
x-x_track [mm]

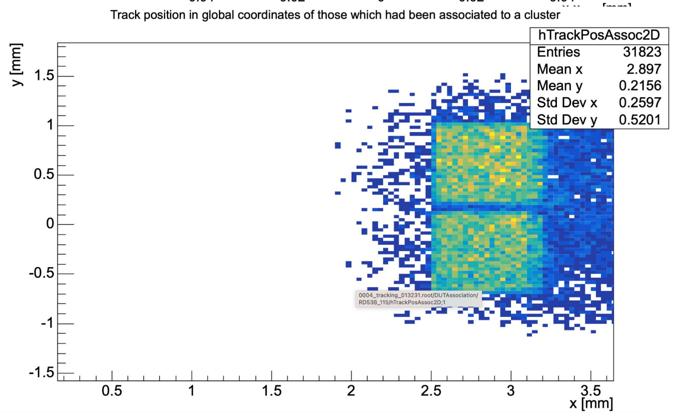
Track position in global coordinates of those which had been associated to a cluster
hTrackPosAssoc2D
Entries 31823
Mean x 2.897
Mean y 0.2156
Std Dev x 0.2597
Std Dev y 0.5201
y [mm]
x [mm]

Pedestal Noise Comparison Before and After Irradiation

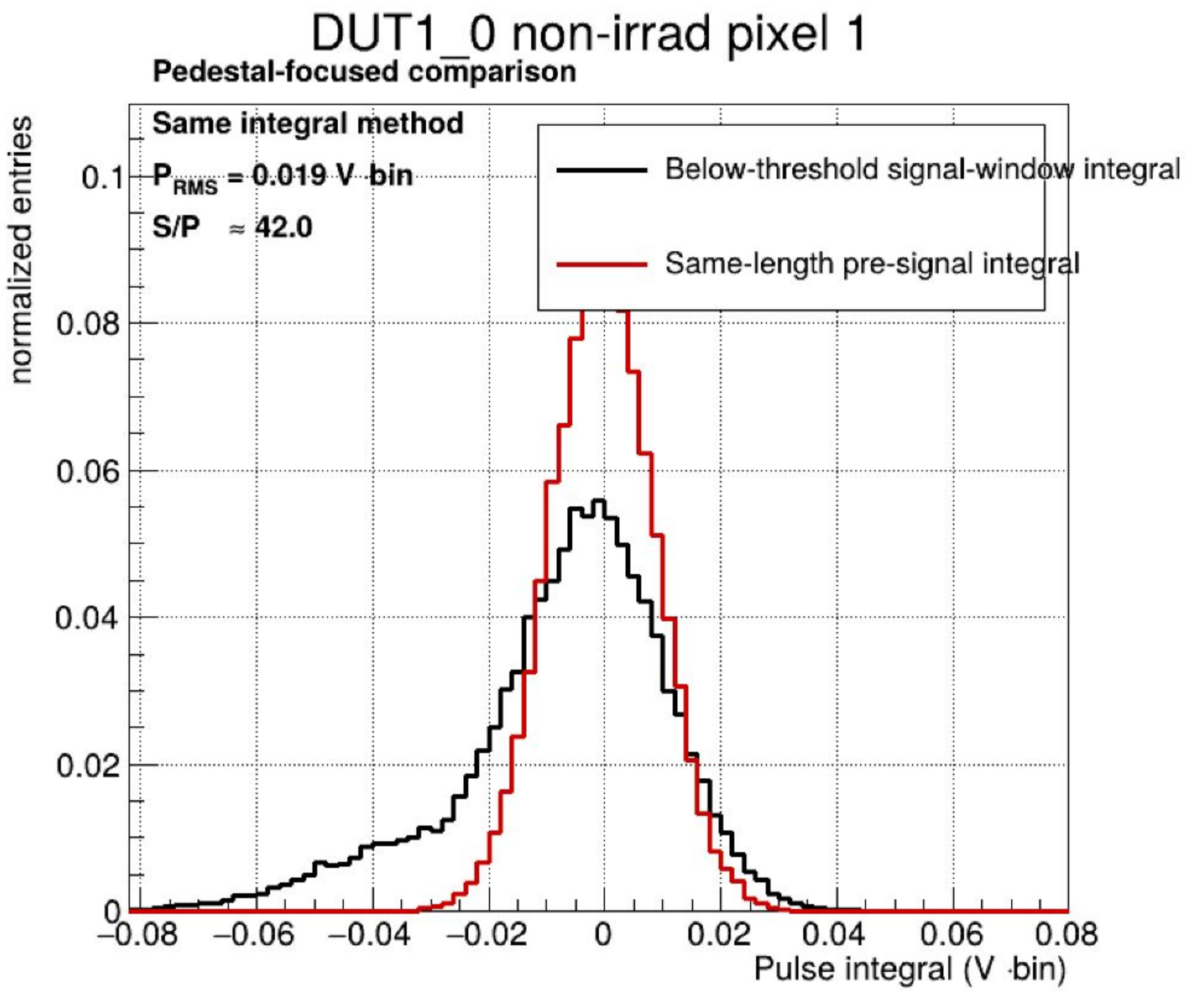

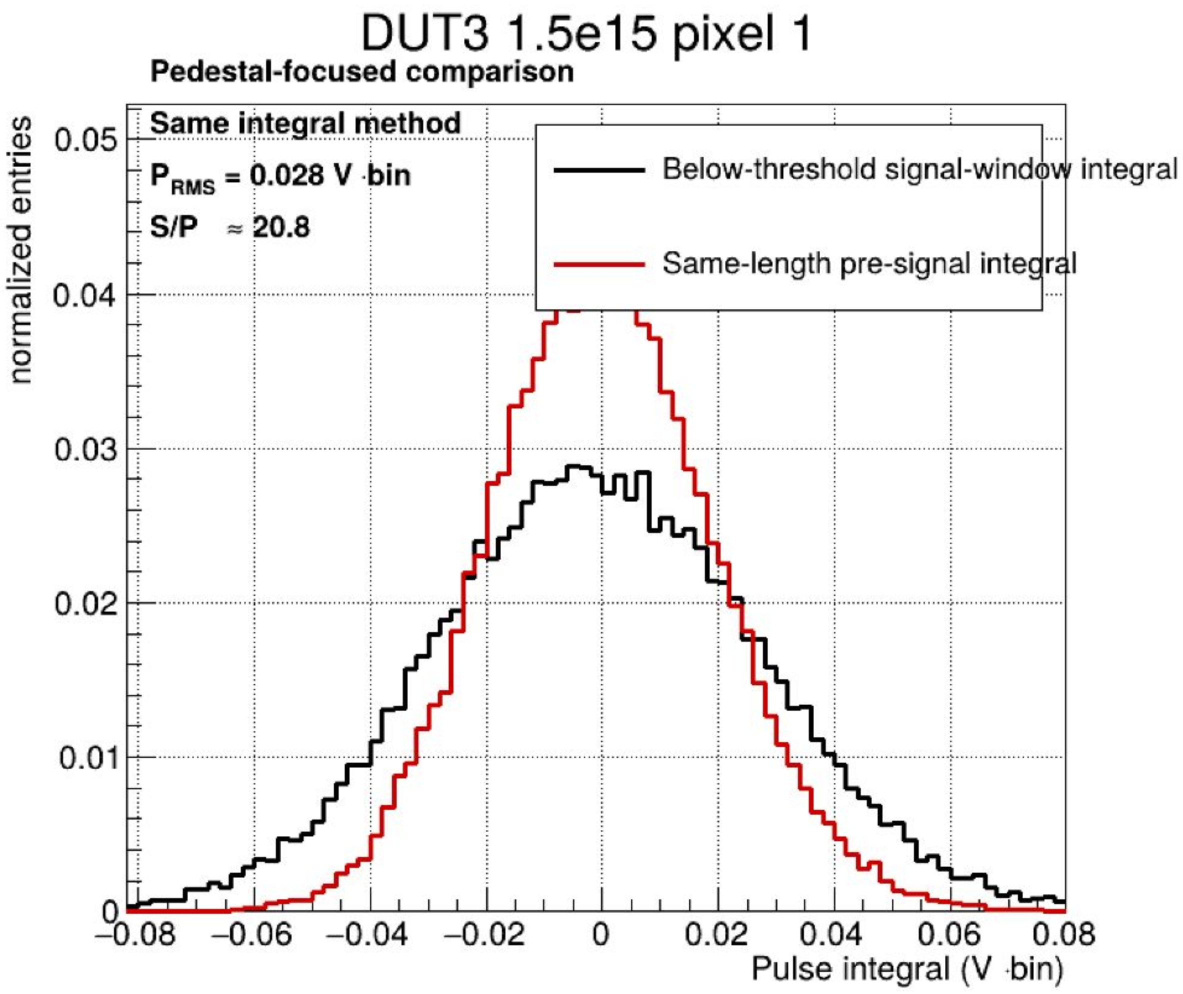

Analysis of tench and off-trench regions with tracks required

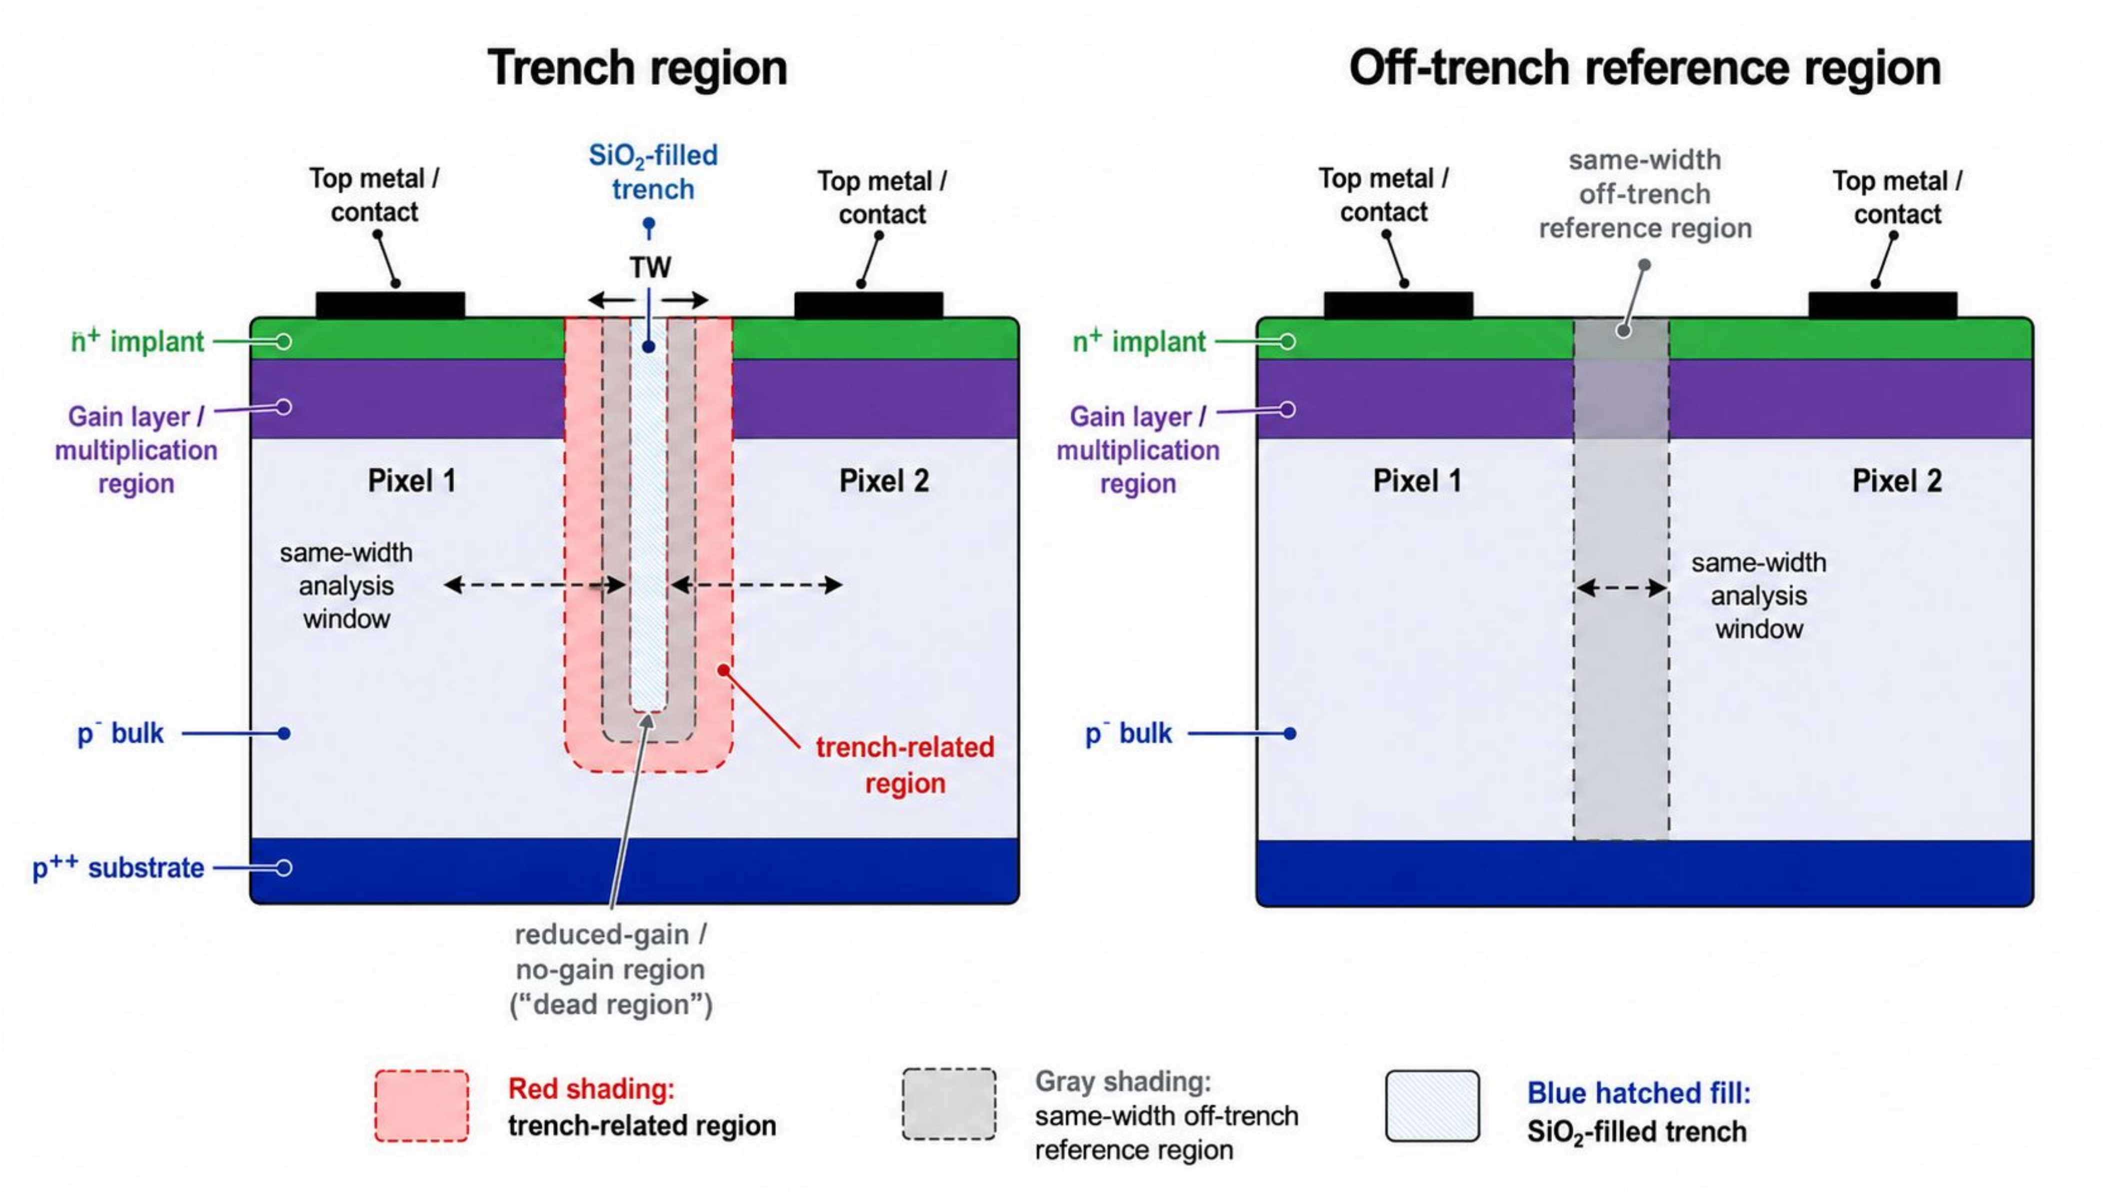

How does the trench-region signal spectrum compare to a nearby non-trench region?

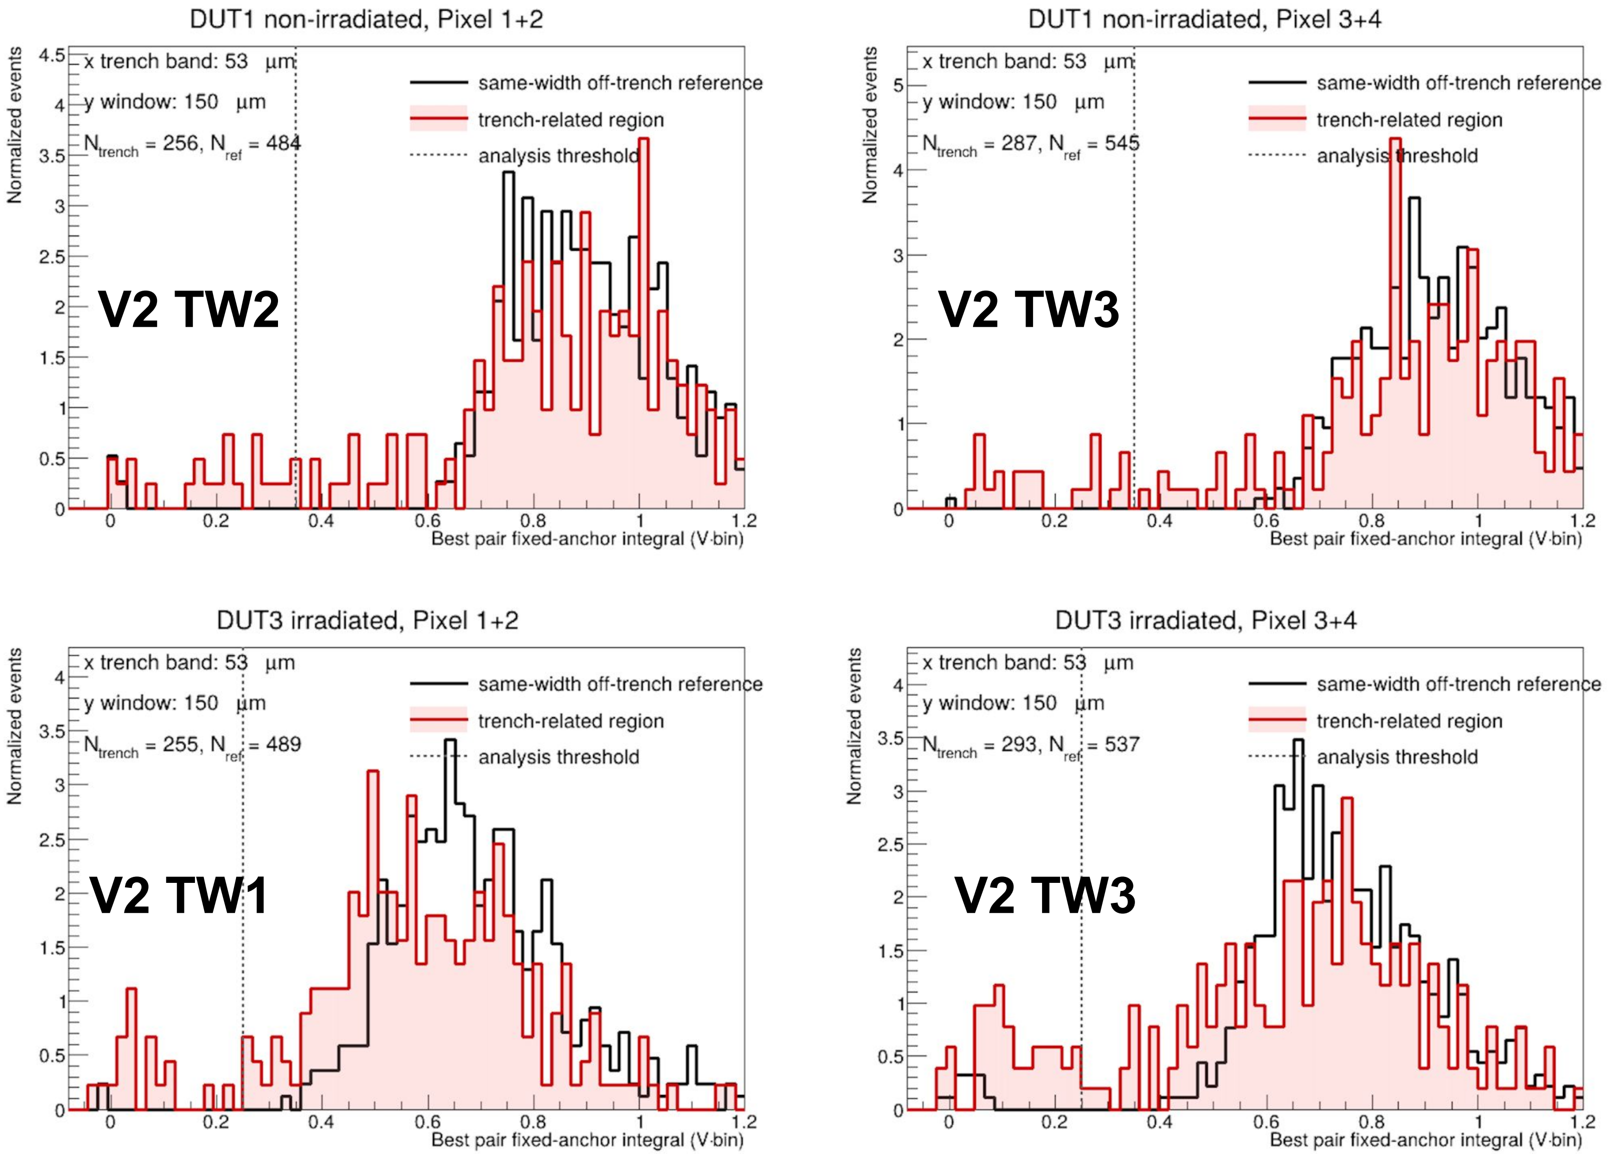

The trench-related region shows a lower-charge population than the off-trench reference region, which is consistent with the local efficiency loss.

x trench band: 53 µm

y trench band: 150 µm

Tracking-Gated[1] Timing Results Explanation

We measure three time differences per event, gated by the telescope track

For each event, the track is interpolated across the three sensors, and the time differences between their signals are calculated.

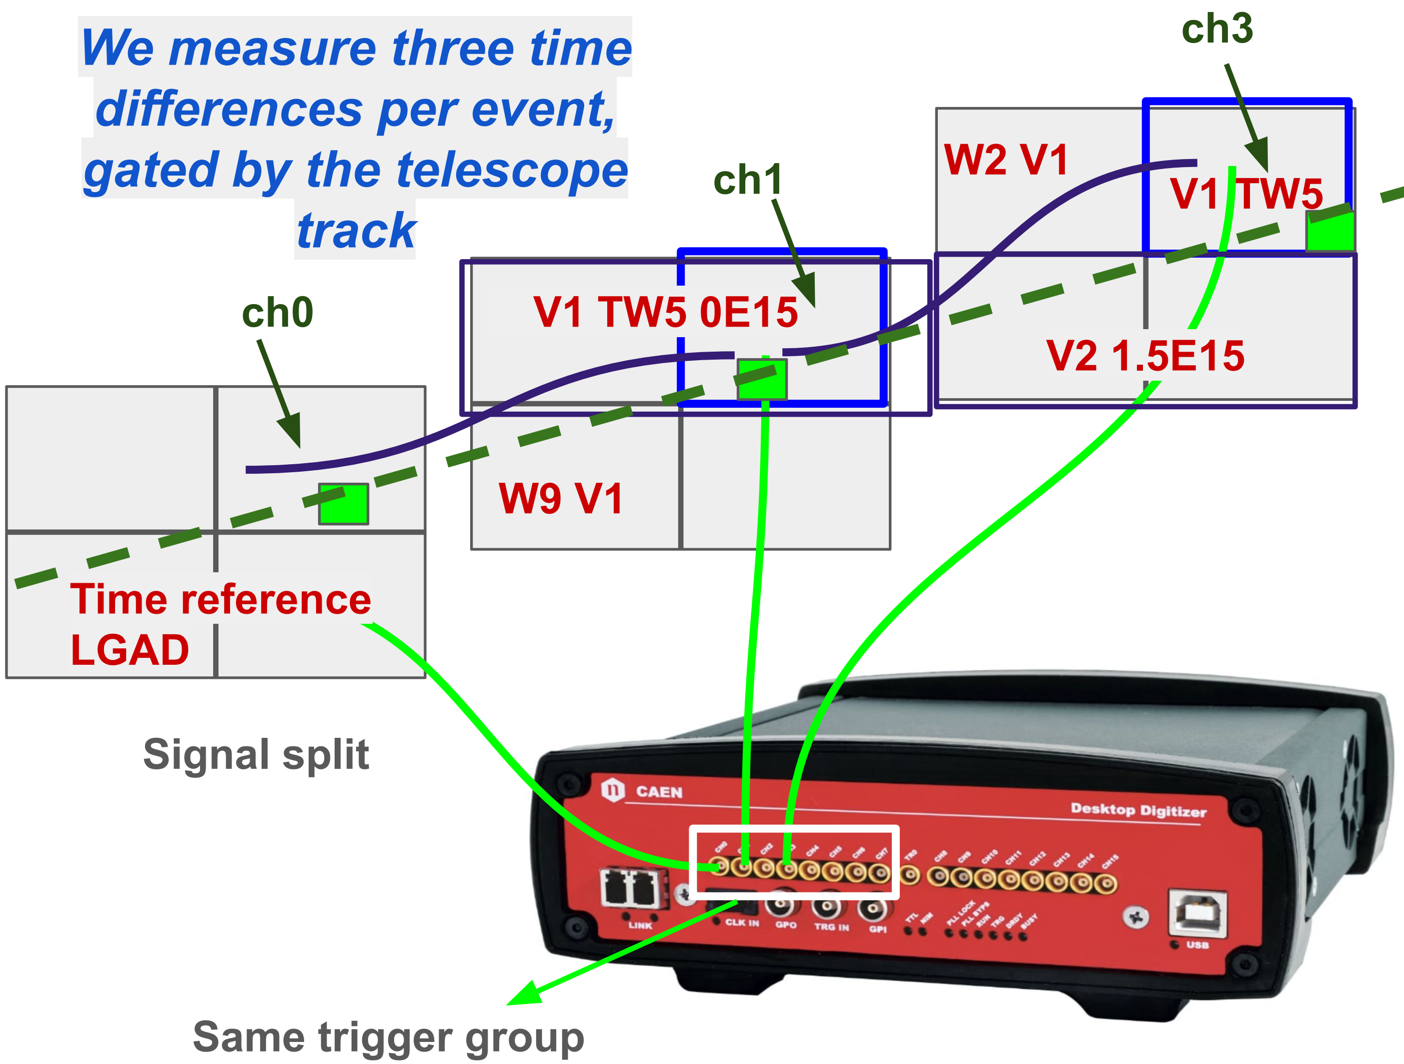

- **A** = t(ch1) – t(ch0) (non-irradiated DUT − REF)
- **B** = t(ch3) – t(ch0) (irradiated DUT − REF)
- **C** = t(ch1) – t(ch3) (DUT − DUT, REF-independent)

Triplet system equations:

$$A^2 = \sigma_1^2 + \sigma_R^2, \qquad B^2 = \sigma_3^2 + \sigma_R^2, \qquad C^2 = \sigma_1^2 + \sigma_3^2$$

[1] Track gating is the process of keeping only data associated with valid particle tracks and rejecting unrelated or noisy events

Timing tracking-anchored results

Batch2 / run509

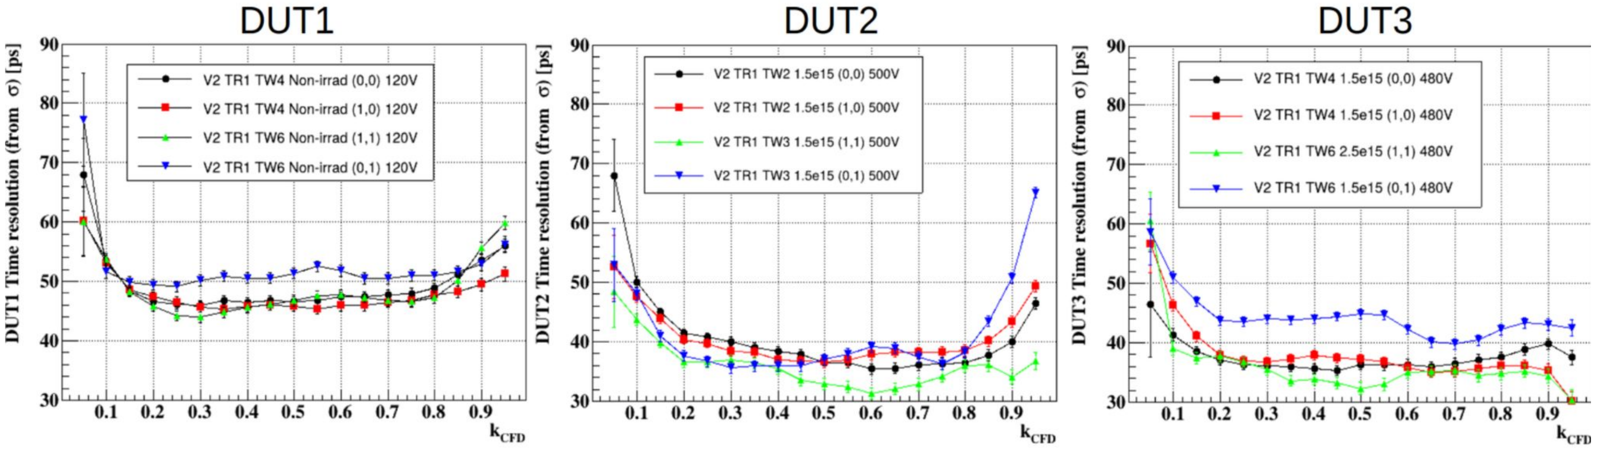

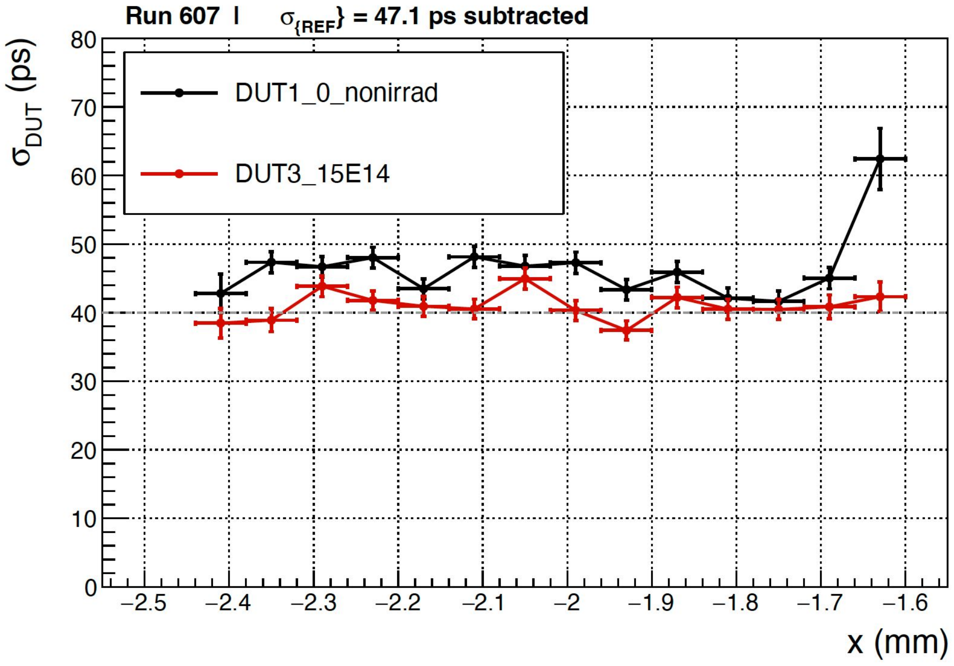

Cuts:
NPeaks=1
RiseTime>0.6
(to get rid of digitizer fast noise)

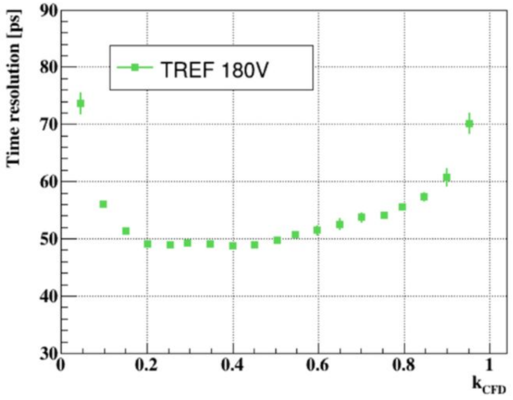

0,1	1,1
0,0	1,0

Comparison with Marcos' results

	Batch 2 run509	Batch 3 run607
DUT1 non-irrad	~47 ps @ 120V	**43 ps @ 165V**
DUT3 1.5×10^{14}	~38 ps @ 500V	**38 ps @ 495V**
TREF 180V	~49 ps	47 ps ✓

Batch 3: σ_{REF}= 47.1 ps (triplet method)

Threshold and CFD scans with track-based TREF-anchor selection

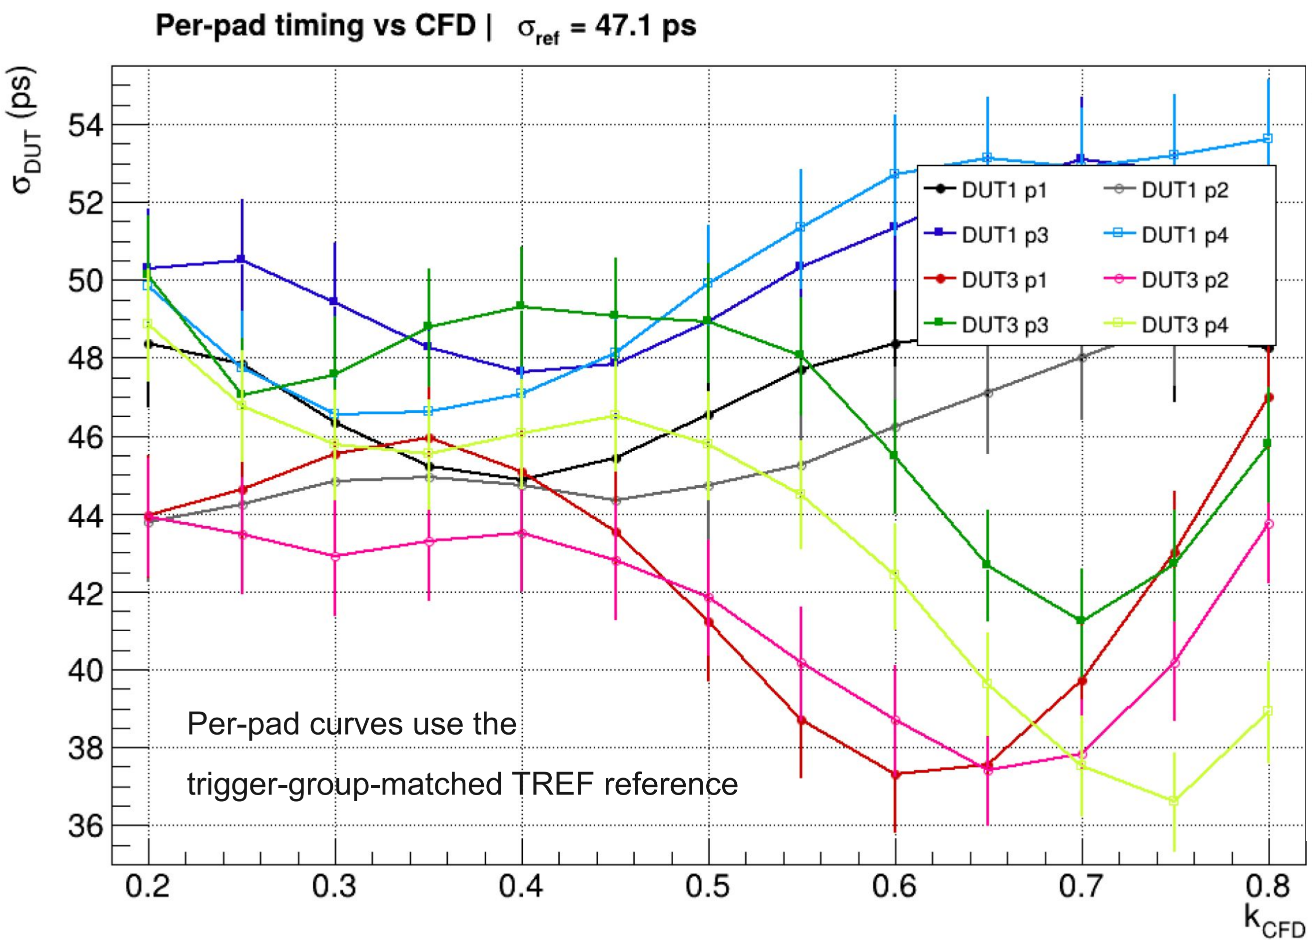

Detector-averaged timing shown as a function of CFD fraction

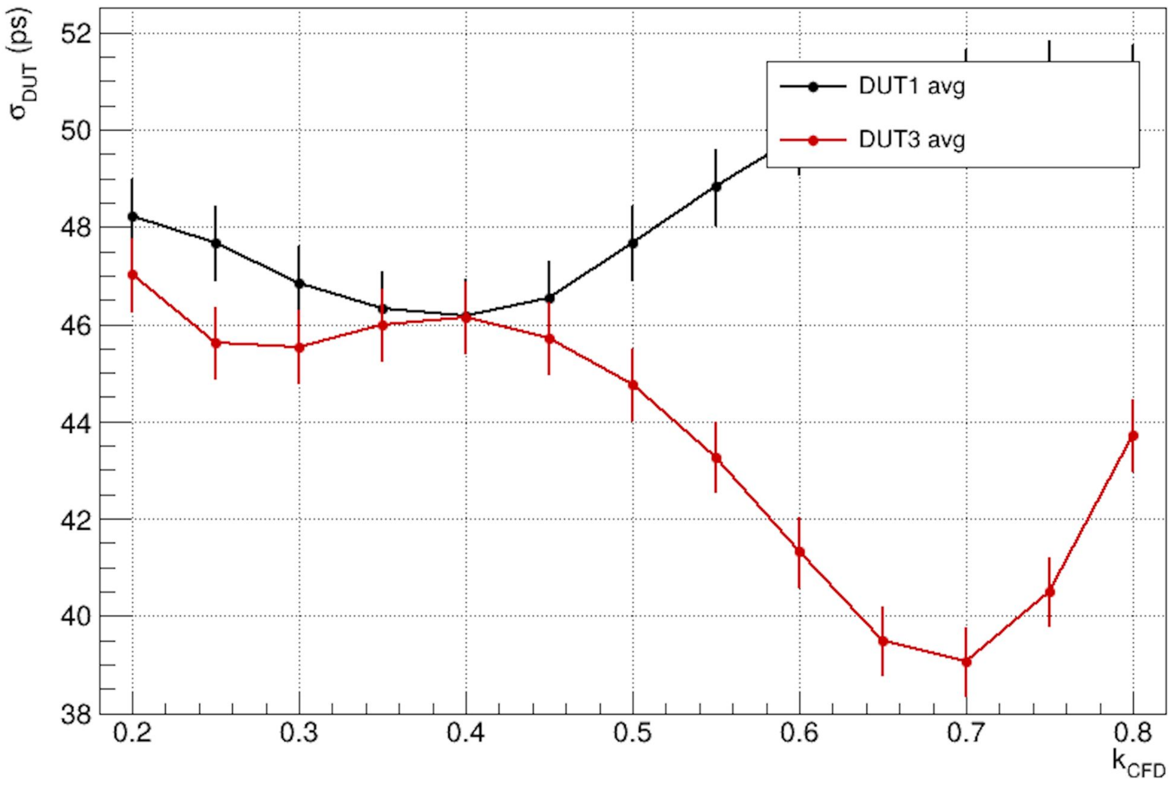

DUT 1: W2 V2 TW2/3 pixel 1/2 (p1/2) @ 165 V

DUT3: W2 V2 TW1/3 pixel 1/2 (p1/2) @ 550 V

Pulse-integral distributions for TI-LGAD pixels before and after irradiation

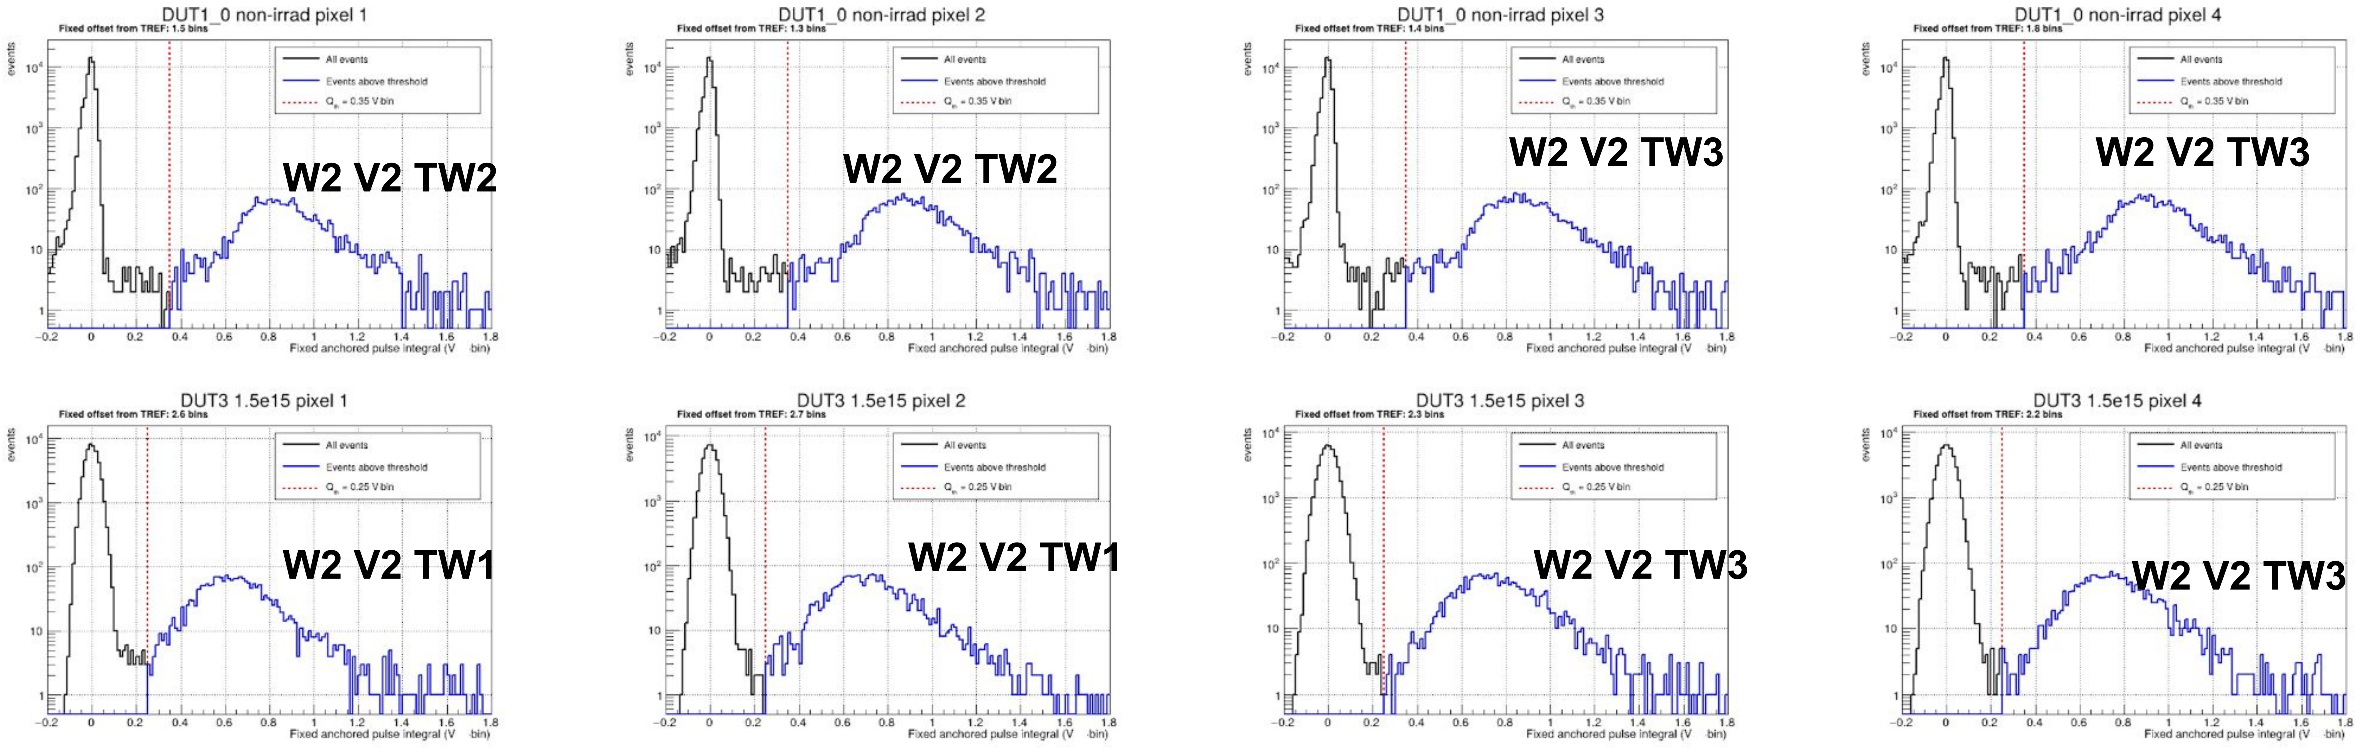

Top: non-irradiated DUT
Bottom: DUT irradiated to $1.5{\times}10^{15}\ n_{eq}/cm^2$

Black distributions contain all events, blue distributions show events passing the pulse-integral threshold

W5 V2 TW12 0e15 inclined at 10°

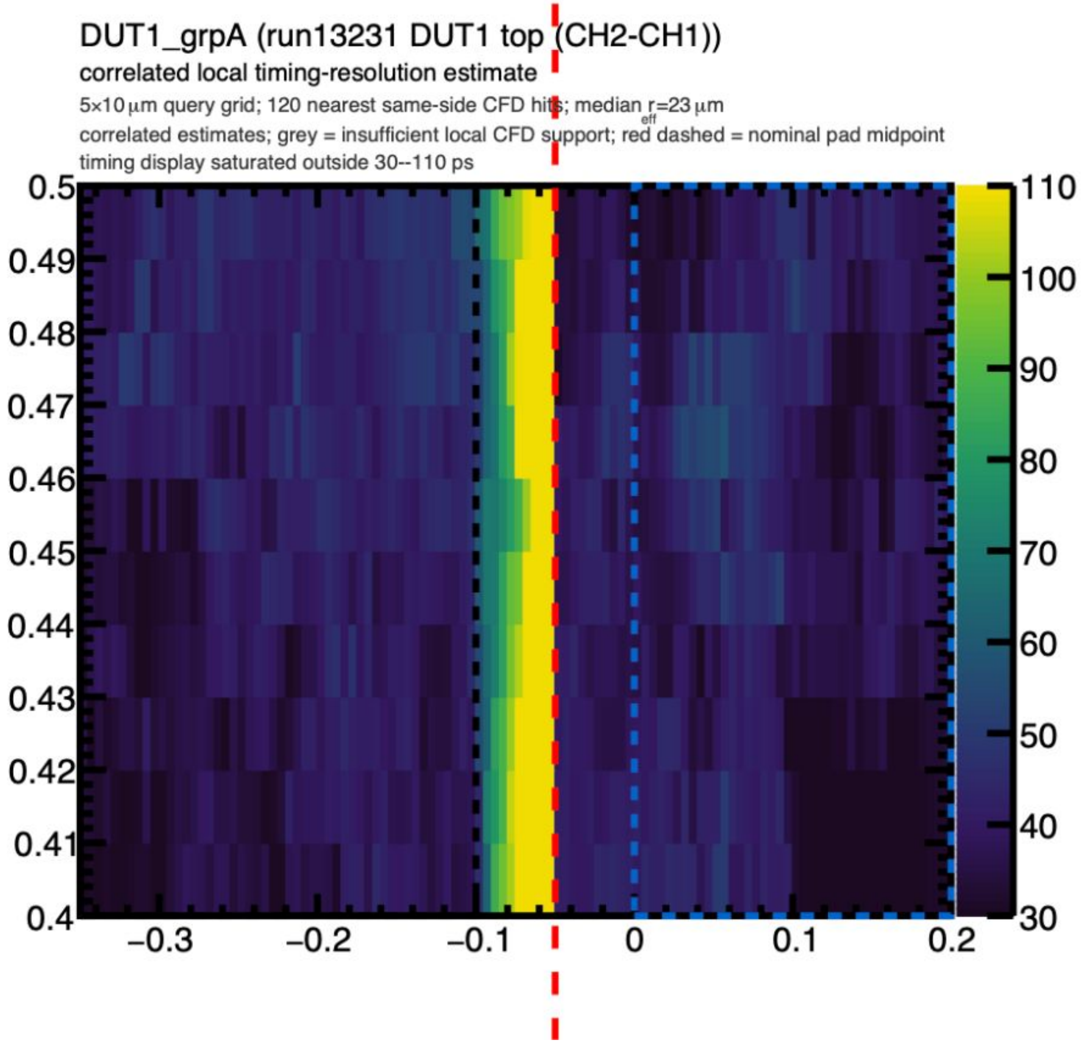

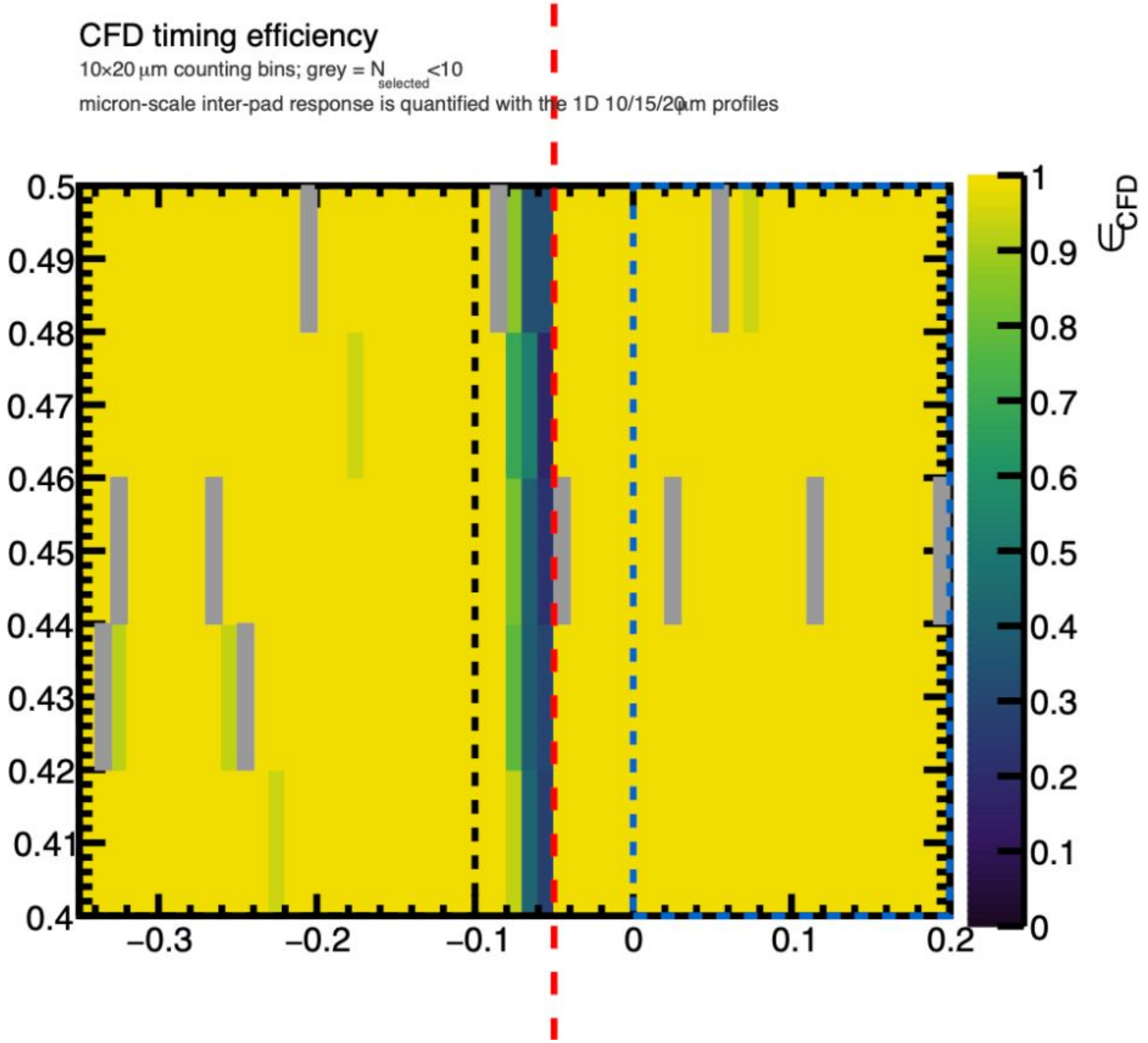

Clean-bulk timing resolution: ~38 ps

Spatial timing for DUT4 W5 V2 TW4, $1.5\times10^{15}\ n_{eq}/cm^2$

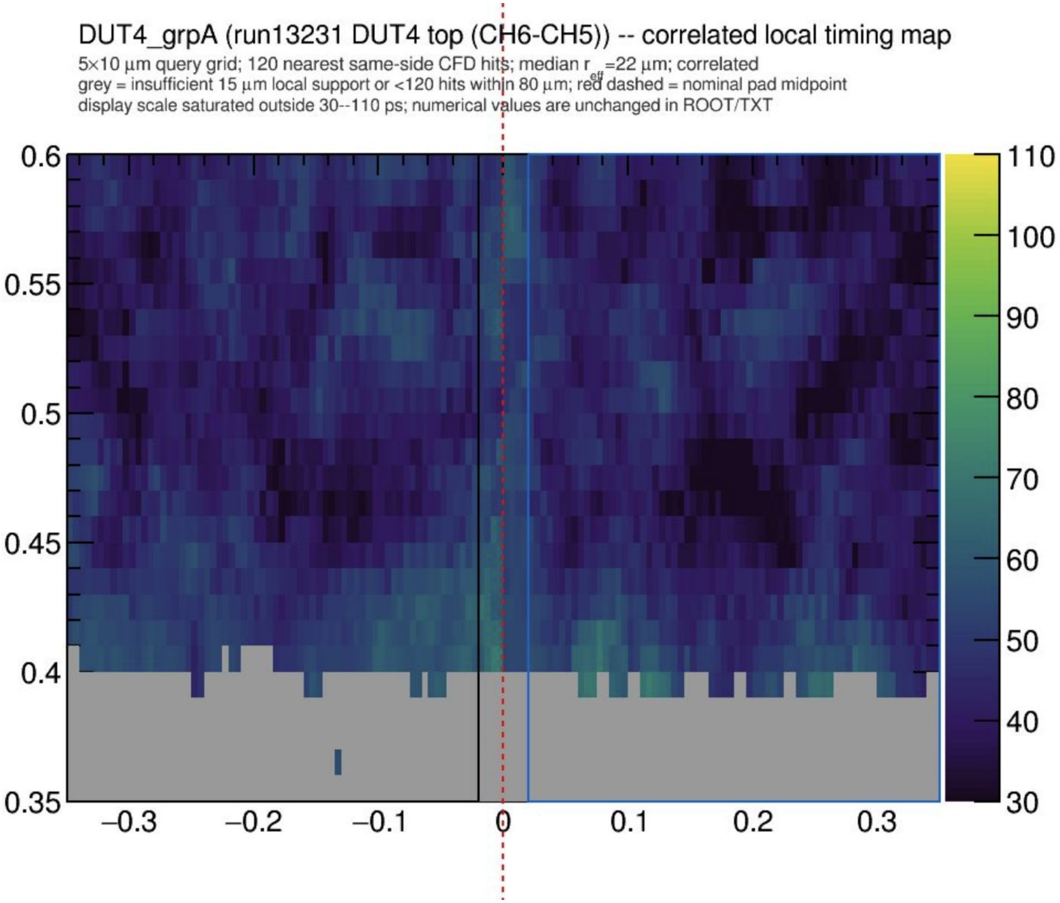

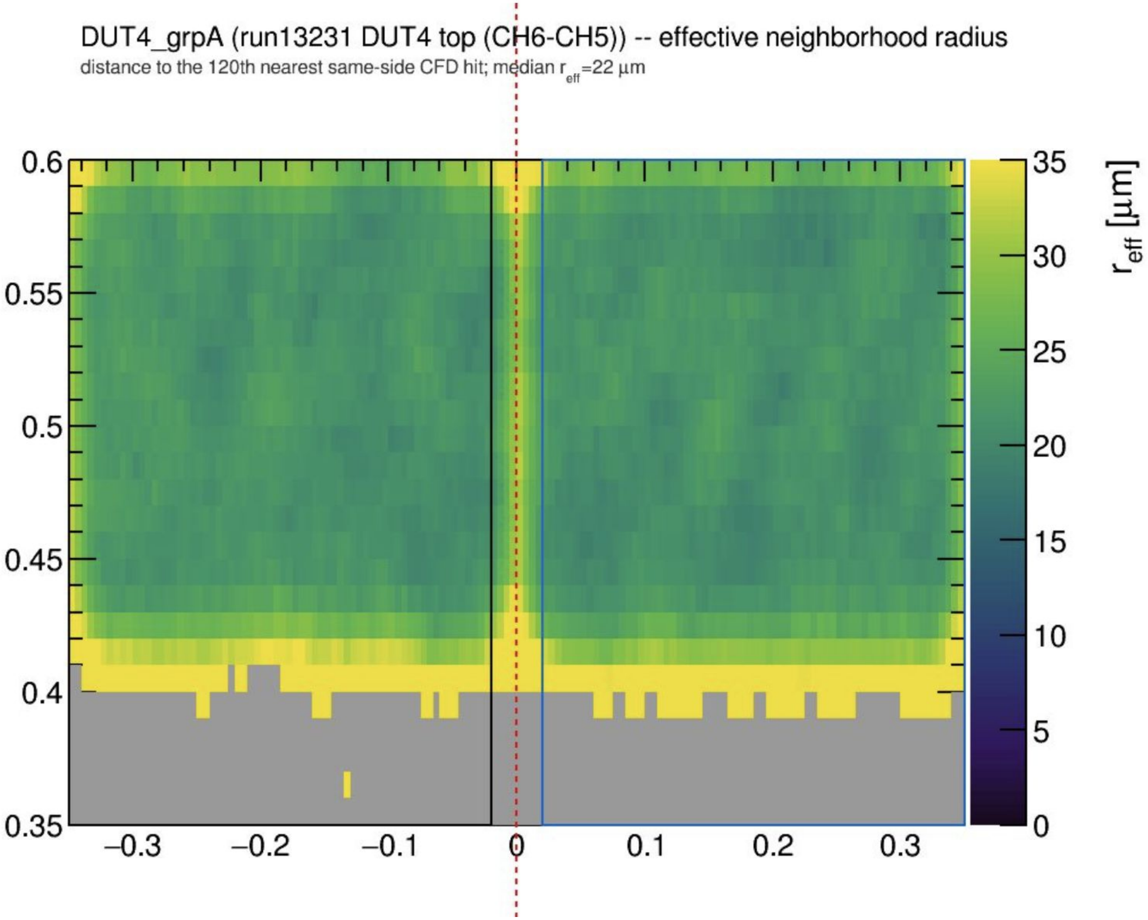

Bulk timing: ~42–46 ps.
Local estimates use 120 same-side CFD hits, with median reff≈22 µm.
The relatively uniform reff confirms that the timing map is not driven by strongly varying local statistics.

Wafer 5 analysis

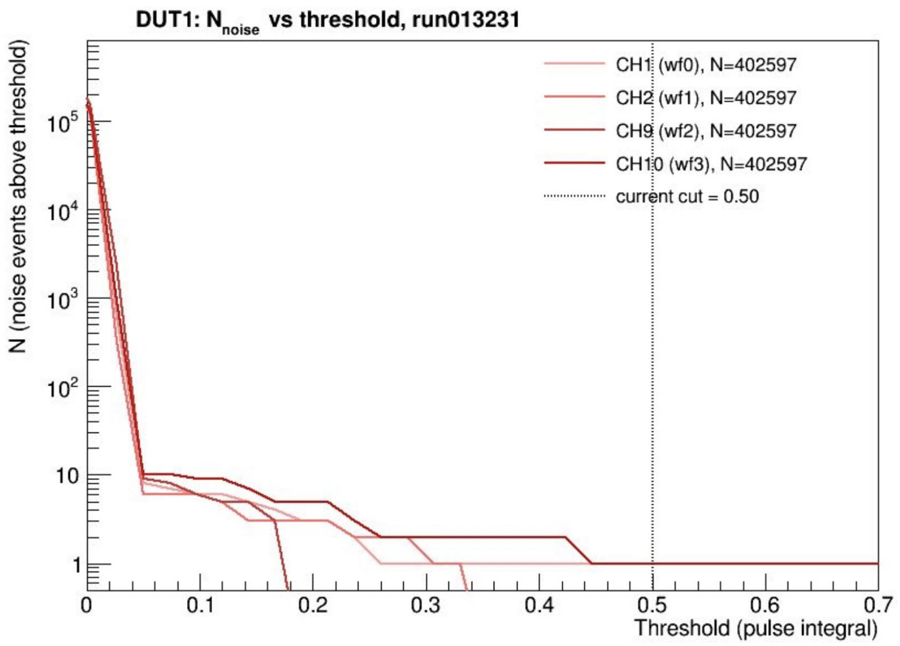

W5, V2, TW1-2, 0e14 n_{eq}/cm^2, angled at 10°

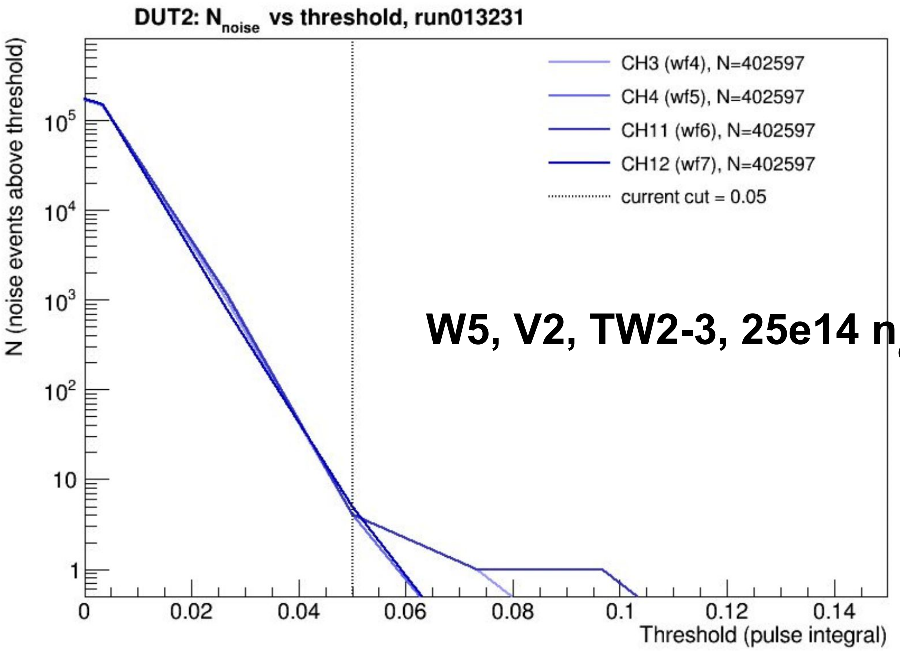

W5, V2, TW2-3, 25e14 n_{eq}/cm^2

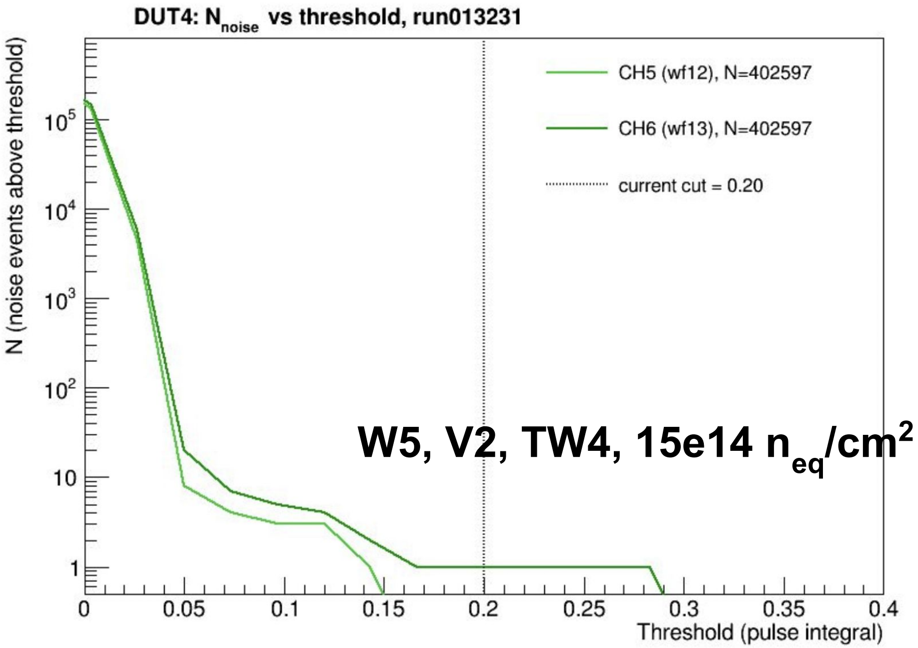

W5, V2, TW4, 15e14 n_{eq}/cm^2

Anchoring the Window: DUT–TREF Peak-Offset Calibration

- The window's center can’t depend on the DUT's own CFD edge -> so it's anchored to TREF's peak, plus a fixed per-channel time offset.
- That offset is measured once per channel, using only “clean bulk” tracks (inset from the pad edges), with S/N > 8 required on both TREF and the DUT.
- Raw peak-bin difference (DUT − TREF) is collected for up to 2000 such events per channel.
- Offset = robust median of that sample; spread = 1.4826×MAD; a 5σ clip removes outliers before finalizing.
- A channel is calibrated only if ≥120 events survive — otherwise its charge samples are skipped entirely, never guessed.

Calibration gates, at a glance

Target sample	2000 events / channel
Minimum sample	120 events / channel
SNR gate	amplitude > 8× noise RMS
Outlier clip	5σ (robust MAD-based)
Bulk region	inset from pad edges

Bulk ROI: One Timing Resolution Per Pad

- Scanning in the bulk is done for every k_{CFD} (TREF and DUT) and fixed before any position-resolved analysis.
- The bulk ROI pools all clean-bulk tracks from a fixed box each pad's own geometric box, shrunk inward away from both the trench and the outer sensor edge — into one sample, and computes σ_{DUT} once with the same σ_{68}/Gaussian-cross-check machinery used everywhere else.
- The 1D trench profile uses two bin widths (10 μm for pure counting quantities like amplitude/efficiency/jitter, a coarser 15 μm for σ_t's fit stability) with edges anchored at x=0 so no bin ever straddles and mixes both sides of the trench, and that x=0 itself is refined from a validity-gated erf fit to the data's own efficiency.. An analytical per-event jitter curve is then quadrature-subtracted from σ_t to expose a non-jitter residual

This number is jitter-INCLUSIVE

- The bulk ROI's σ_{DUT} is total detector resolution — electronic jitter is NOT subtracted here.
- A separate, analytical per-event jitter estimate (noise/slew at the CFD crossing) exists and IS quadrature-subtracted