2026/9/4

TOF General Meeting

Temperature dependence of the time response of strip AC-LGADs for the ePIC BTOF

All numbers are WORK-IN-PROGRESS

Hiroshima Univ. , HU WPI-SKCM2

Kanato Matsutani

Contents

This beam test targets the temperature dependence; the measured quantities are σ_{total} and the noise jitter

1 Background

EIC-ePIC and the Barrel TOF: requirement 90 ps, stretch goal 40 ps

2 Objective

Resolution components; this study measures σ_{total} and σ_{jitter} vs temperature

3 Method

KEK PF-AR 2 GeV e−, MCP-PMT reference, 21.3–38.5 ° C

4 Analysis I — σ_{total}

CFD50, correction chain, inverse-variance two-strip combination

5 Analysis II — jitter

Definition $\sigma_{noise}/|dV/dt|$ and its ingredients vs temperature

6 Results & outlook

Temperature coefficients; decomposition with the laser as the next step

EIC-ePIC and the Barrel TOF

Strip AC-LGADs give the ePIC Barrel TOF its timing for π/K/p separation

ePIC Barrel TOF — π/K/p separation by time of flight **requirement 90 ps · stretch goal 40 ps**

EIC / ePIC

Electron-ion collider at BNL; ePIC is the general-purpose detector. PID in the barrel relies on TOF + Cherenkov.

BTOF sensor

Strip AC-LGAD: 500 μm pitch, 50 μm thick, one metal strip reads many; position from charge sharing.

Why temperature

Impact ionization falls as the sensor warms → gain and amplitude drop, which directly costs timing.

The detector will not sit at one fixed temperature. The temperature coefficient of the time response must be measured — that is the main purpose of this beam test.

What we measure: σ_{total} and the noise jitter

The time resolution is a quadrature sum of components; this beam test measures two of them directly

$$\sigma_{total}^2 = \sigma_{jitter}^2 \oplus \sigma_{Landau}^2 \oplus \sigma_{delay}^2 \oplus \sigma_{reference}^2 \oplus \cdots$$

Measured here (1): σ_{total}

Width of the arrival-time difference vs the MCP-PMT reference, after corrections — the quantity ePIC cares about.

Measured here (2): σ_{jitter}

Per-event $\sigma_{noise}/|dV/dt|$ — the electronics-noise component, measurable event by event without a tracker.

Two axes of this talk: (1) derive σ_{total} and its temperature dependence · (2) the jitter and its ingredients vs temperature.

The full decomposition of the remaining terms (Landau, position, reference) is ongoing with the laser TCT → outlook.

KEK PF-AR beam test setup

DUT upstream, MCP-PMT downstream as the time reference; sensor temperature set by a water-cooled Cu plate

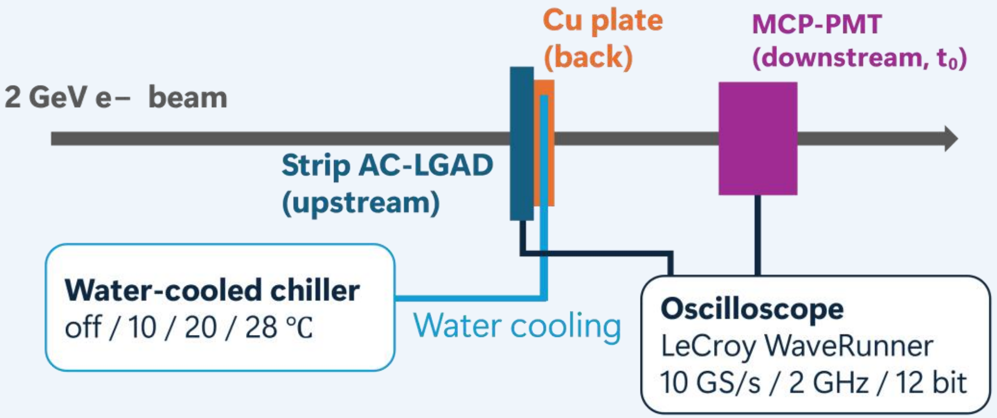

Strip AC-LGAD (DUT)

- BTOF Baseline sensor (JFY25)
- 500 µm pitch
- 50 µm x 1 mm width strip
- 64 x 2 strips per sensor

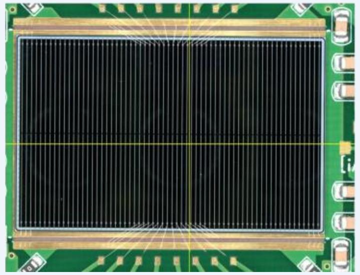

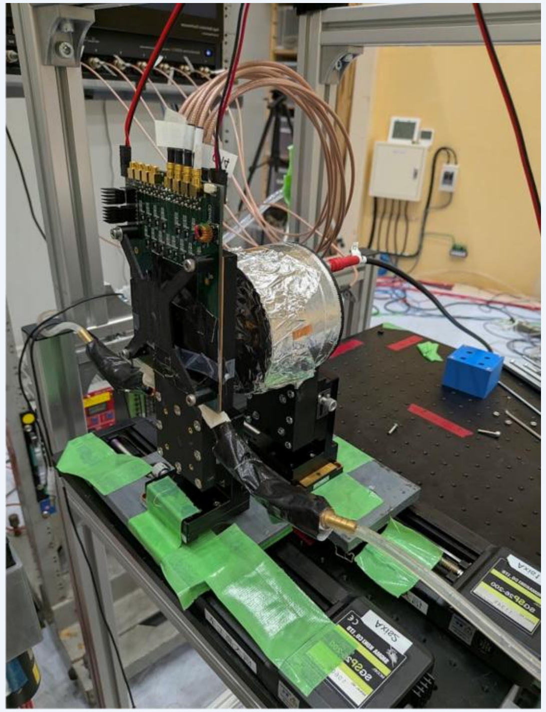

No tracker: the hit position is reconstructed from charge sharing (η_3 reconstruction).

Beam

2 GeV e−, PF-AR single bunch (1.257 µs revolution).

Reference & trigger

Photek multi-anode MCP-PMT gives t_0 and the trigger.

Readout

FNAL 16-ch amp board → LeCroy scope, 2 GHz / 10 GS/s / 12 bit; 7 strips + MCP-PMT.

Run list and temperature control

11 runs at 185 V span 21.3–38.5 ° C; the in-run drift sets the horizontal error bar of every later plot

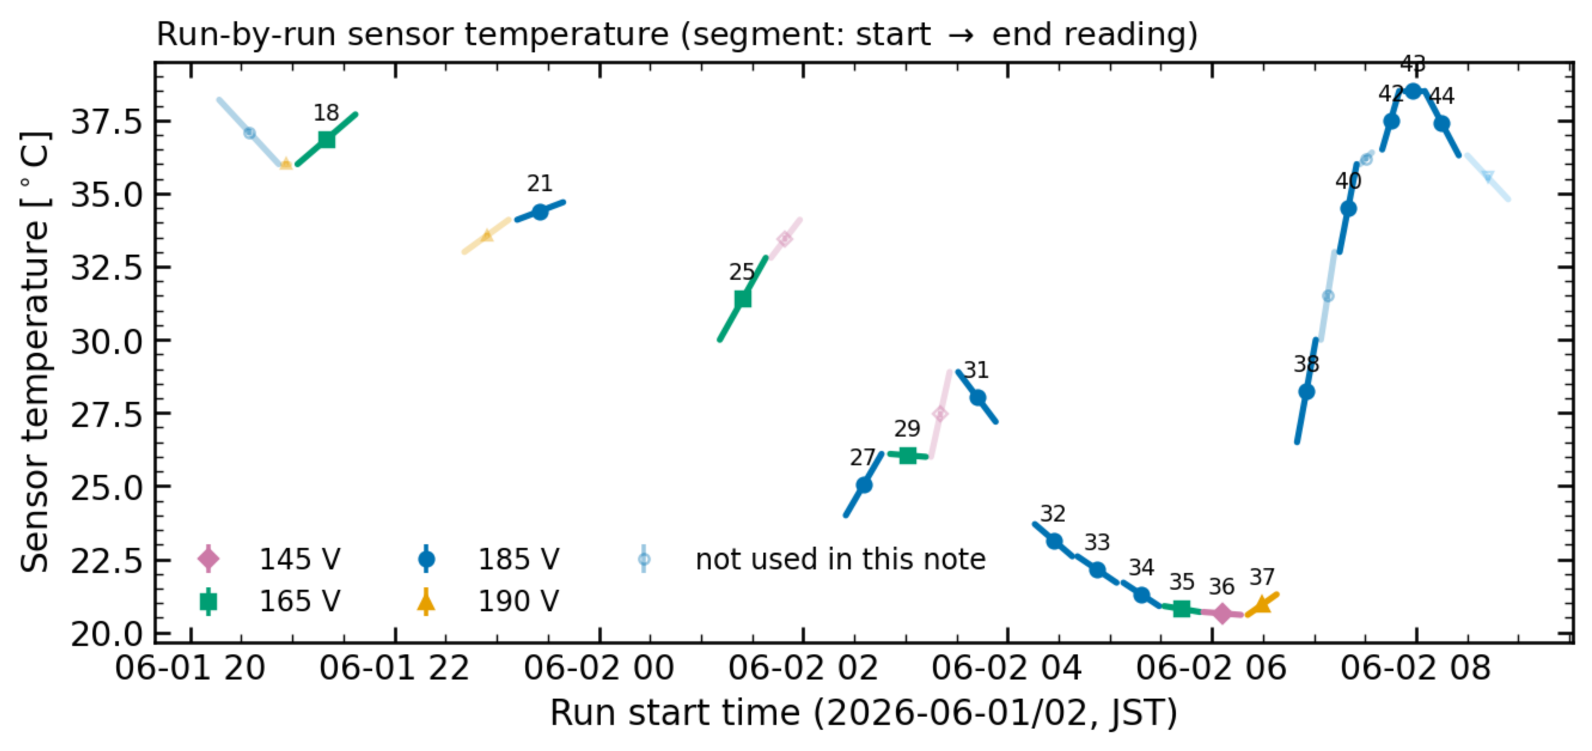

Temperature read by hand at run start / end (0.1 ℃ thermometer); mean in-run drift 1.4 ± 1.1 ℃.

Main series — 185 V

11 runs, T = 21.3–38.5 ℃: population of every temperature fit.

Sub series — 165 V

4 runs: low-gain control of the same effect.

HV series — T ≈ 21 ℃

145–190 V: temperature ↔ gain equivalence.

Reference run — run 34

185 V, 21.3 ℃, n = 5280 clean events.

Timing resolution Analysis

- Event selection
- Position reconstruction
- Arrival time definition (CFD50%)
- Correction of timing resolution

Event selection and hit position vs amplitude

One cut set for every run; charge sharing maps the hit position and modulates the leading amplitude

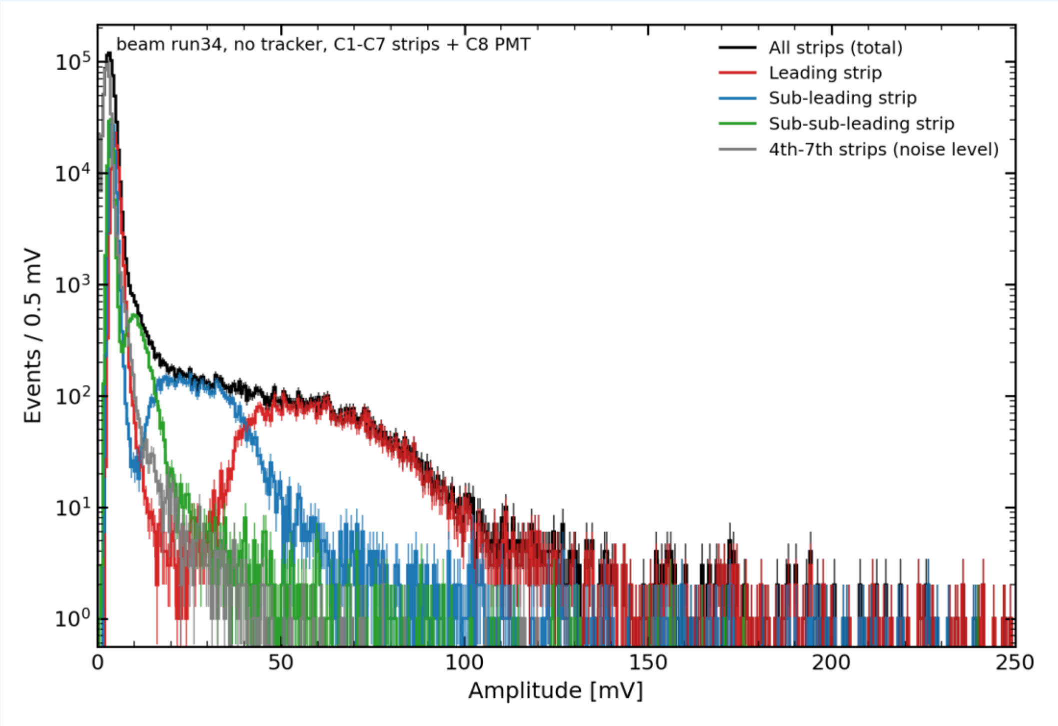

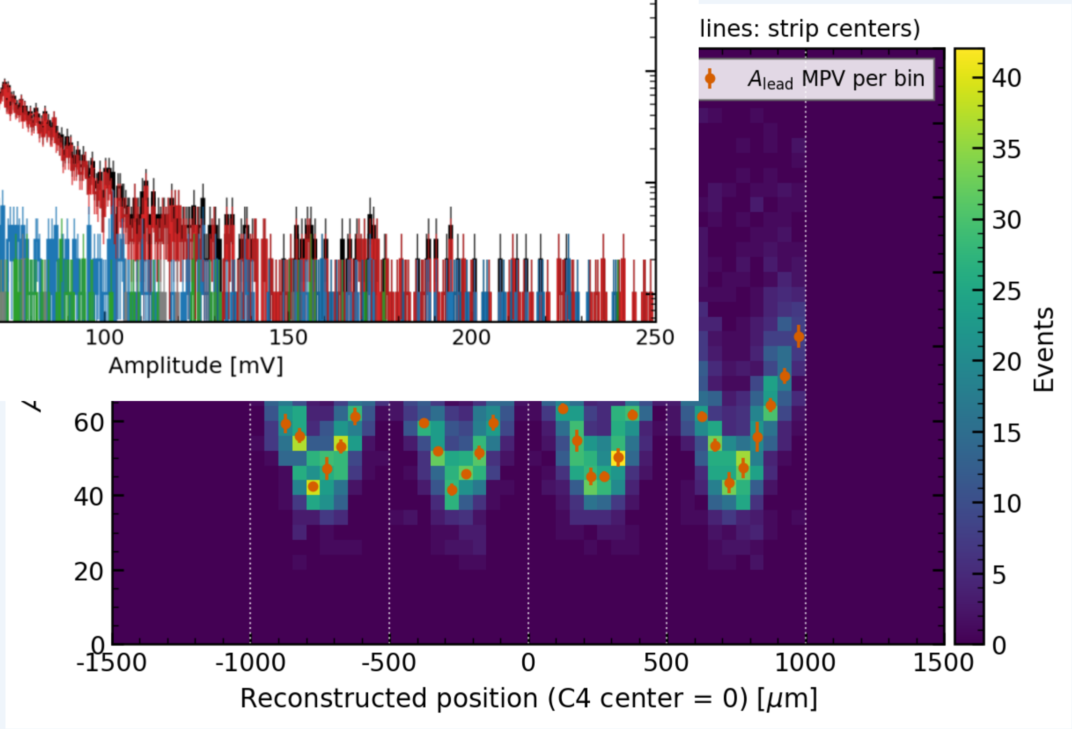

Cuts (identical for all runs)

MCP ≥ 500 mV · leading strip in C2–C6 with 20 ≤ A ≤ 150 mV · |Δt| < 1 ns · C1/C7 veto.

Position variable

$\eta = A_{right}/(A_{left} + A_{right})$, scaled by the laser caribration.

What the map shows

A_{lead} is largest on-strip and drops between strips, where the charge is shared — the basis of the η position.

Event selection and hit position vs amplitude

One cut set for every run; charge sharing maps the hit position and modulates the leading amplitude

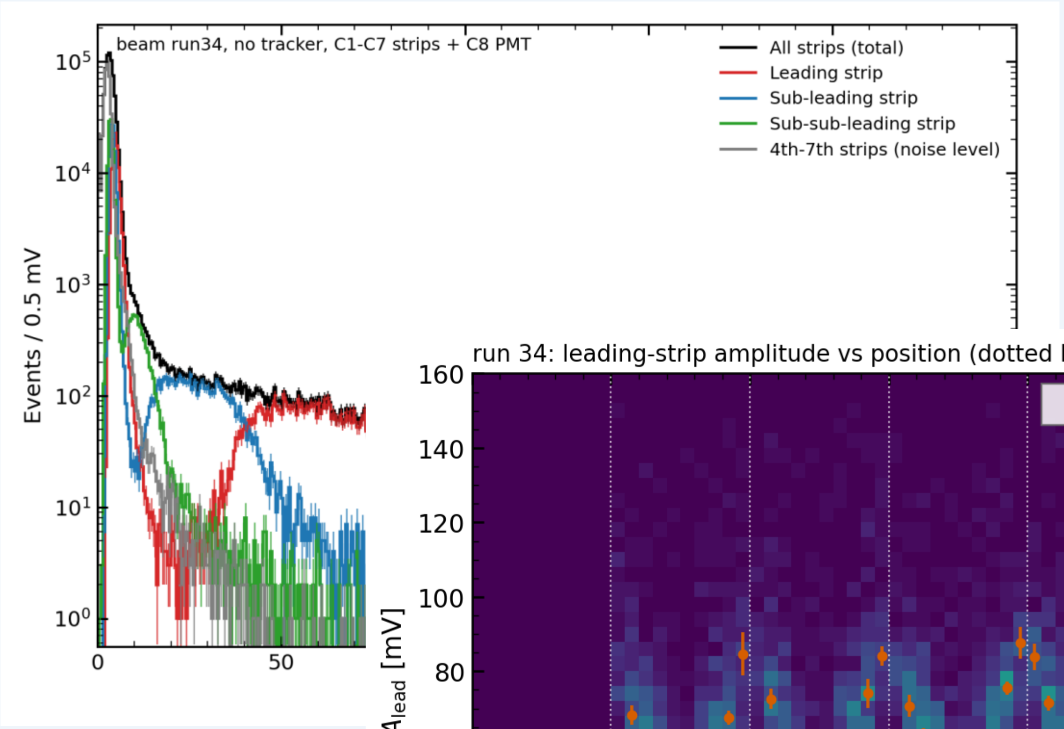

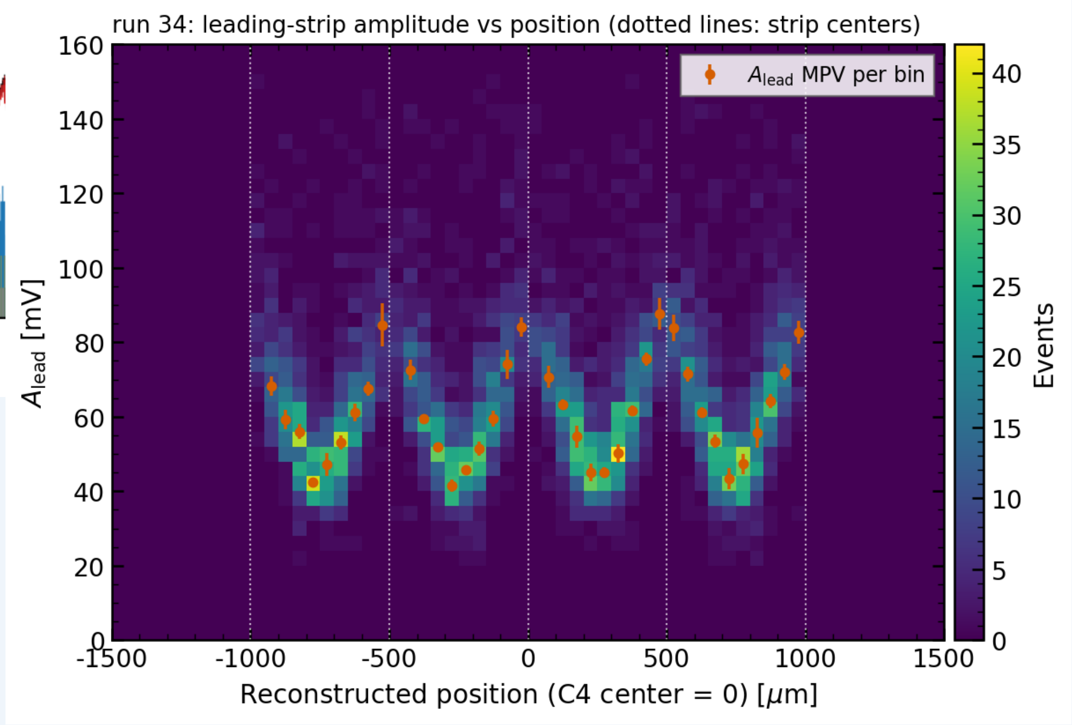

Cuts (identical for all runs)

MCP ≥ 500 mV · leading strip in C2–C6 with 20 ≤ A ≤ 150 mV · |Δt| < 1 ns · C1/C7 veto.

Position variable

$\eta = A_{right}/(A_{left} + A_{right})$, scaled by the laser caribration.

What the map shows

A_{lead} is largest on-strip and drops between strips, where the charge is shared — the basis of the η position.

Arrival time definition: CFD50

The arrival time is the 50%-of-peak crossing on the leading edge, interpolated between waveform samples

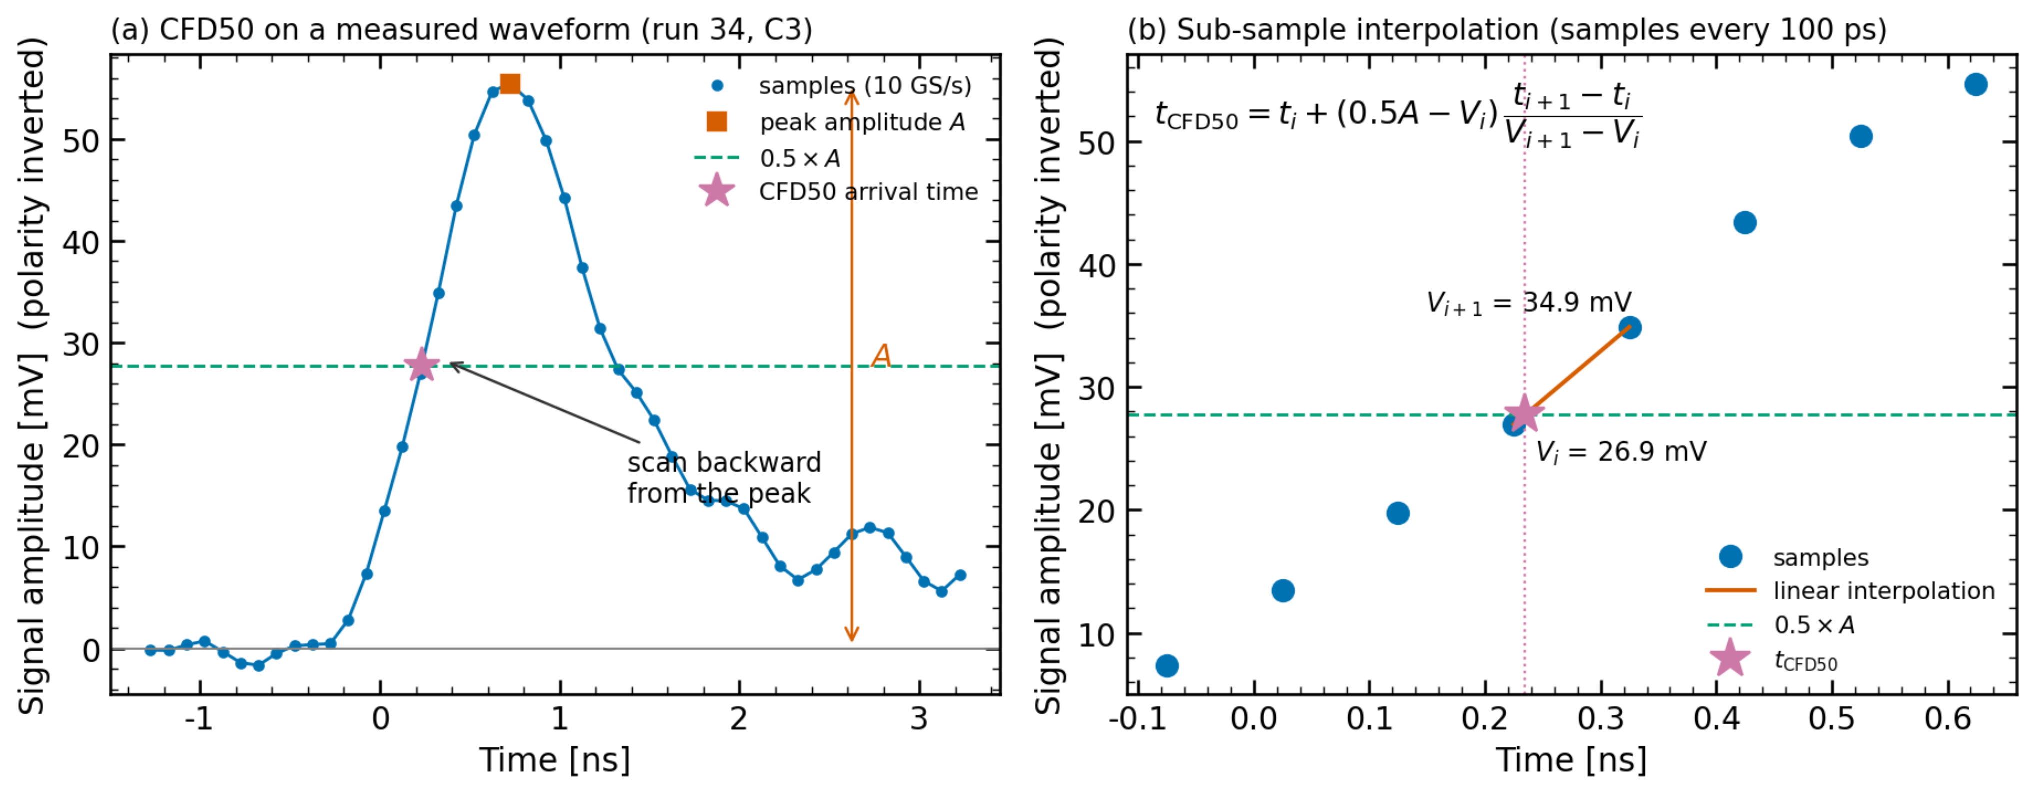

1 · peak amplitude

A = waveform maximum (pedestal subtracted).

2 · scan from the peak

Walk backward to the samples bracketing 0.5 × A.

3 · linear interpolation

Sub-sample precision from the two 100 ps samples.

Correction chain

Three corrections, iterated three times — and every resolution is quoted twice

Step 1
TW(A_{lead})
Time-walk vs. the leading-strip amplitude

→

Step 2
TW(A_{sub})
Residual walk vs. the sub-leading amplitude

→

Step 3
Propagation(y_η)
Transverse signal propagation vs. hit position

→

LUTs are equal-population-bin medians · the whole loop runs ×3 · CFD 50 % arrival time throughout

Arrival time defference

$$\Delta t_{total} = t_{sensor} - t_{MCP}$$

t_{sensor} : the combined arrival-time leading strip with sub-leading strip used weight

Time resolution decomposition

$$\sigma_{total}^2 = \underbrace{\boxed{\sigma_{jitter}^2 + \sigma_{Landau}^2}}_{\sigma_{sensor}} + \sigma_{MCP-PMT}^2 + \cancel{\sigma_{TW}^2} + \sigma_{delay}^2 + \sigma_{unknown}^2$$

measured

$$\cancel{\sigma_{trans}^2} + \sigma_{long}^2$$

The chain converges at reference run (21℃, -185V)

run 34: σ(Δt_sensor) 90.4 ps → 48.6 ps, and the residual dependence on every correction variable flattens

Before

After

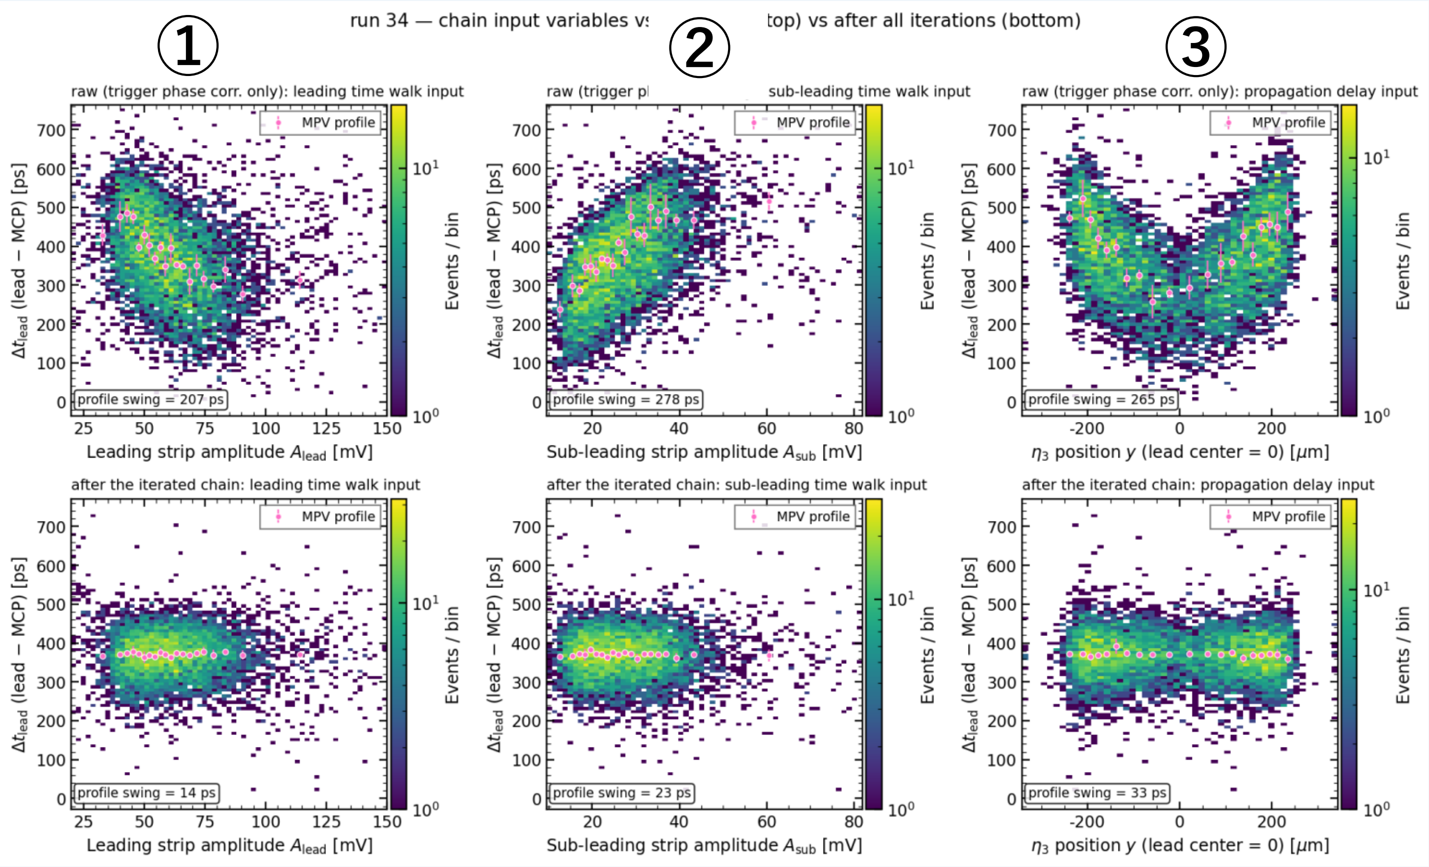

① : Leading amplitude vs Δt

② : sub-leading amplitude vs Δt

③ : η reconstructed position vs Δt

σ_{total}	raw	after iteration 1	after iteration 3
	90.4 ± 3.0 ps	51.2 ± 1.9 ps	48.6 ± 1.3 ps

Per-strip σ from the 3-detector method

Three pairwise time differences between t_{lead}, t_{sub} and t_{MCP} separate the per-strip terms from the common term

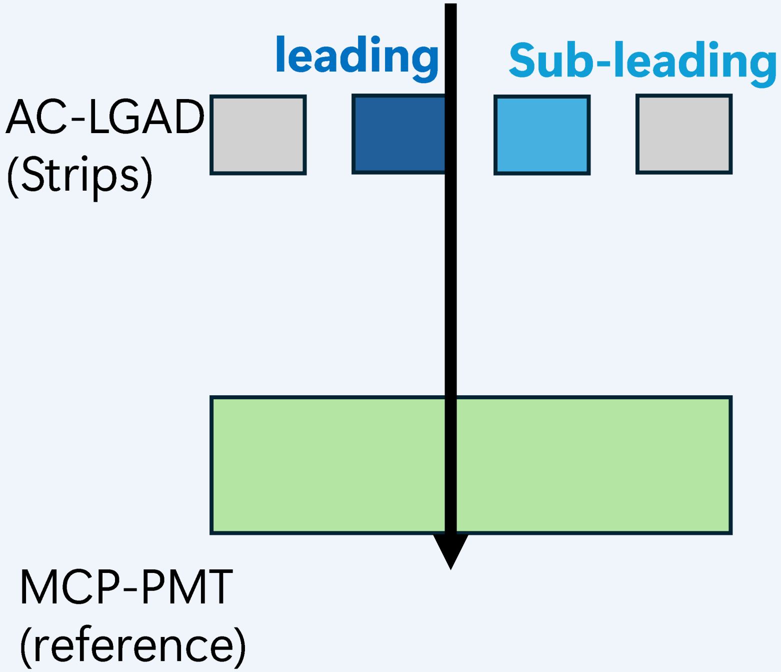

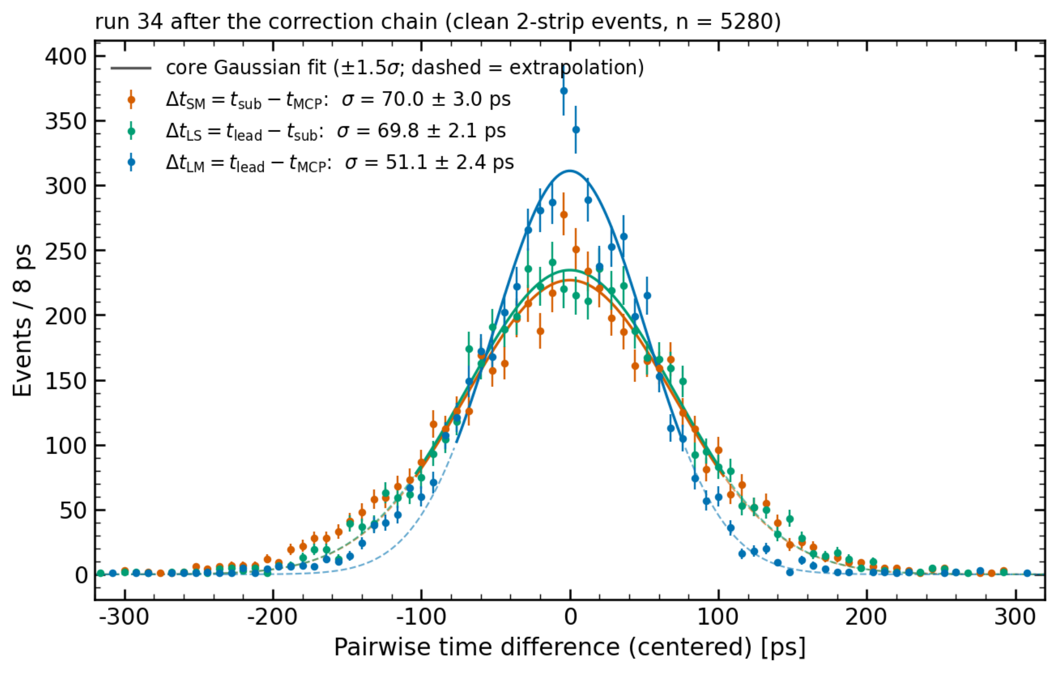

$$\sigma_{LM}^2 = \sigma_{lead}^2 + bucket^2$$
$$\boldsymbol{\sigma_{LM} = 51.1\ \mathrm{ps}}$$
$$\sigma_{SM}^2 = \sigma_{sub}^2 + bucket^2$$
$$\boldsymbol{\sigma_{SM} = 70.0\ \mathrm{ps}}$$
$$\sigma_{LS}^2 = \sigma_{lead}^2 + \sigma_{sub}^2$$
$$\boldsymbol{\sigma_{SM} = 69.8\ \mathrm{ps}}$$

"bucket": common to both strips (MCP⊕ beam bucket ⊕ common landau⊕…)

Result (reference run): σ_{lead}= 35.9 ± 4.0 ps, σ_{sub}= 59.8 ± 2.4 ps. Here they are used only to fix the combination weight on the next slide; the full decomposition is deferred to the outlook.

σ_{total}: inverse-variance combination of the two strips

$t_{sensor} = w * \cdot t_{lead} + (1 - w *) \cdot t_{sub}$; the combined arrival-time histogram is measured directly

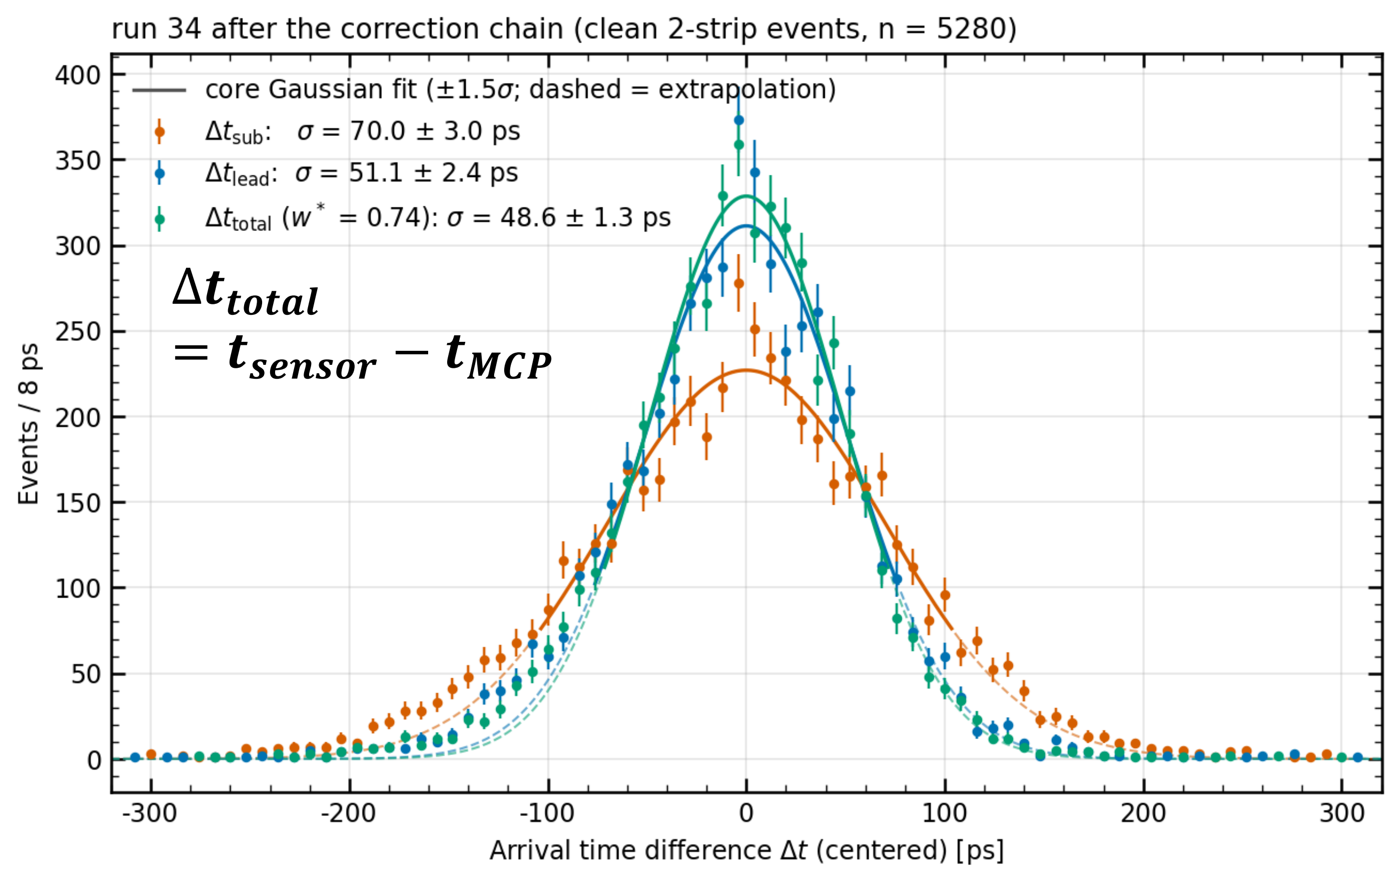

Cross-checks: amplitude weight 48.4 ps, plain mean 50.7 ps — consistent.

$$w * = (1/\sigma_{lead}^2) / (1/\sigma_{lead}^2 + 1/\sigma_{sub}^2) = 0.74$$

$\sigma_{lead}, \sigma_{sub}$ from the 3-detector method (previous slide).

Why 1/σ² weights?

The two strips see the same common term but independent per-strip fluctuations → 1/σ² is the minimum-variance combination. Fixed channel delays are aligned first.

Reference run (-185 V, 21 ℃) result

$$\sigma_{total} = 48.6 \pm 1.3 \text{ ps}$$

σ_{total} vs sensor temperature

σ_{total} stays in the 45–61 ps band over 21.3–38.5 ℃ — far below the 90 ps requirement

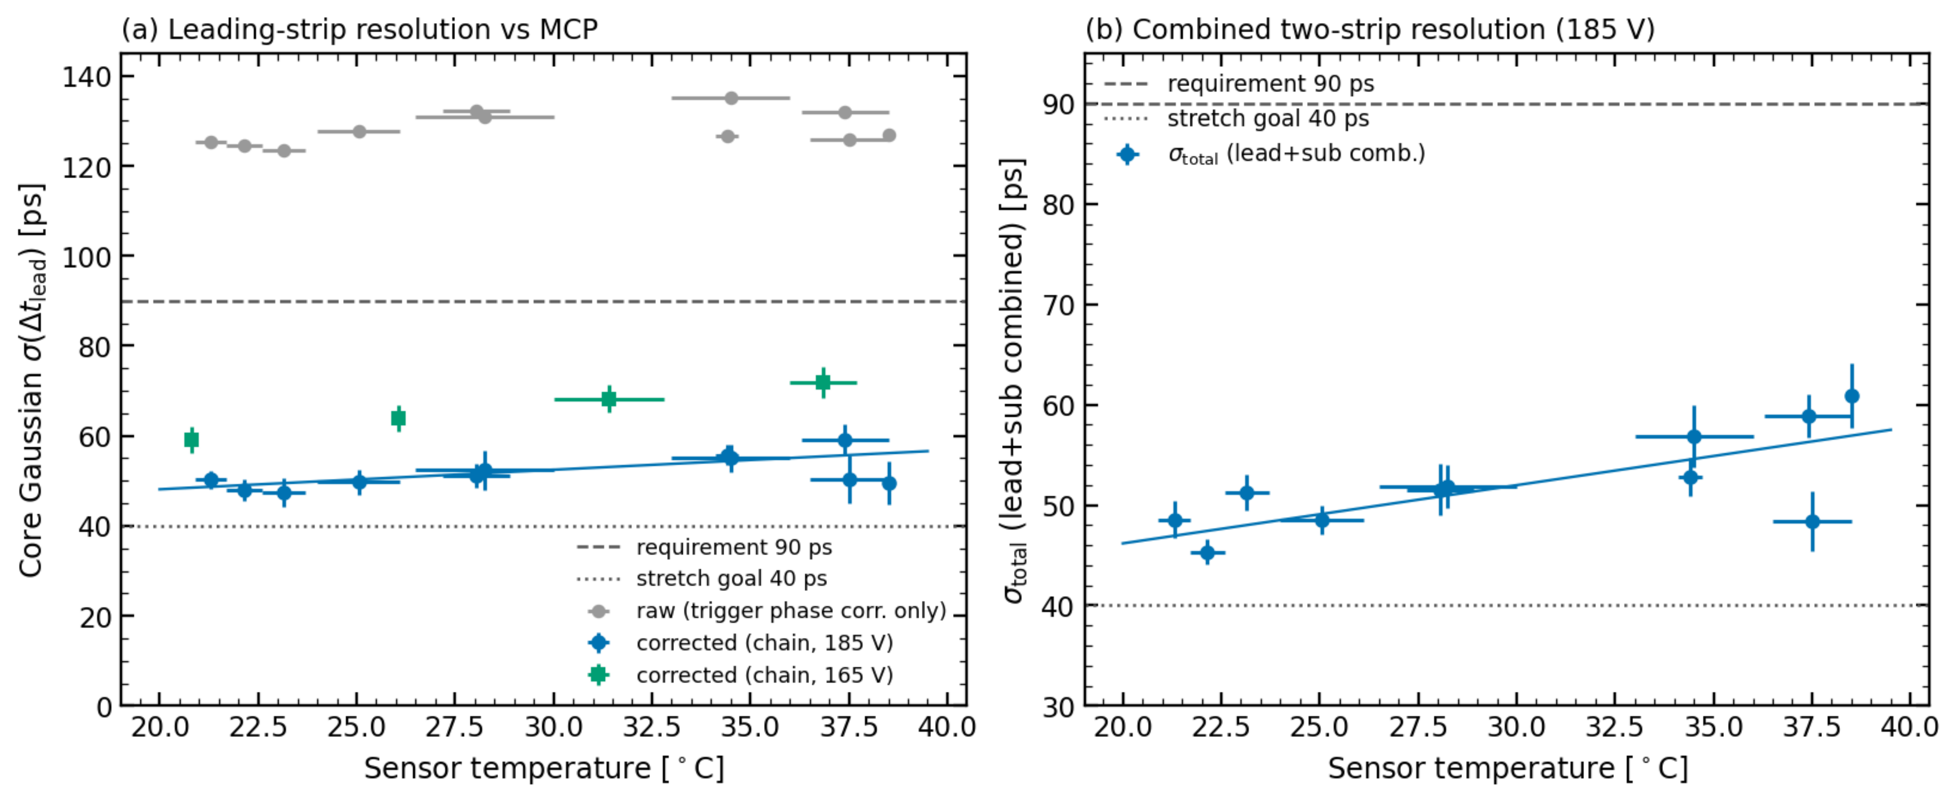

48.6 ± 1.3 ps

σ_{total} at 21.3 ℃

+0.58 ± 0.14 ps/℃

slope, in-situ

≪ 90 ps

requirement met at every point

Jitter Analysis

- Definition
- Temperature dependence

Noise jitter: definition

$\sigma_{jitter} = \sigma_{noise} / |dV/dt|$, **evaluated event by event; the run value is the MPV of the distribution**

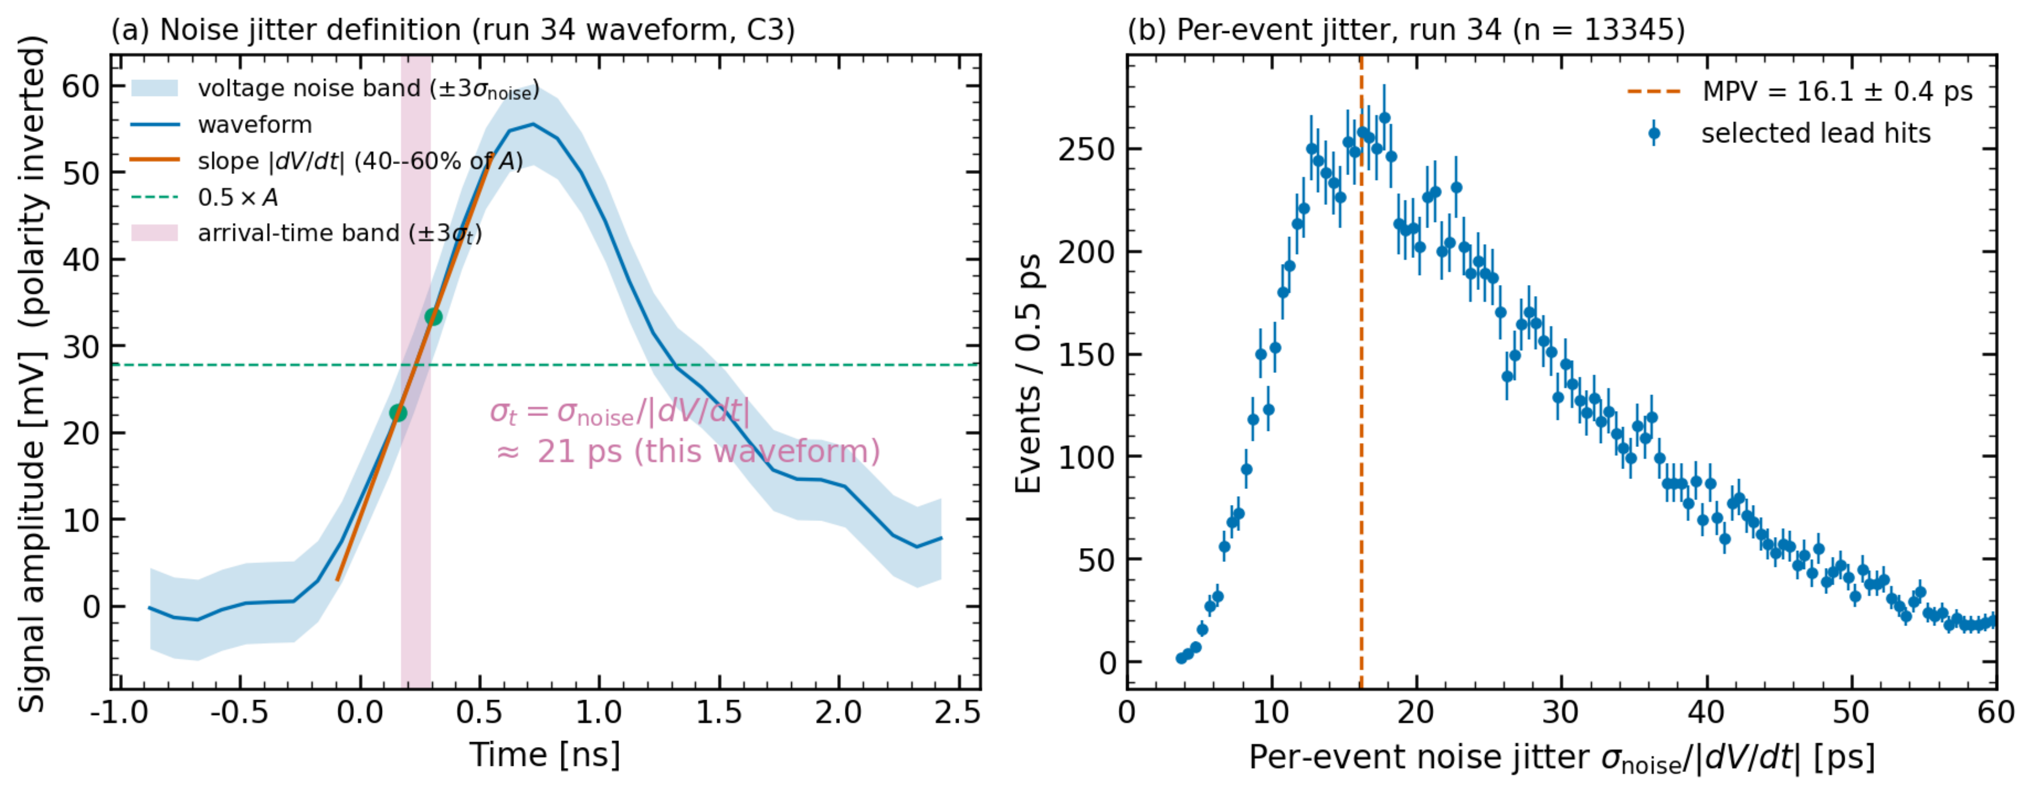

σ_{noise}

Pedestal RMS of the leading strip, per event.

$|dV/dt|$

Mean slope between the 40% and 60% crossings.

MPV, not mean

The 1/A tail is asymmetric → quote the MPV.

What changes with temperature: the gain only

The amplitude falls by −1.59 ± 0.15 %/℃ while noise and rise time stay flat

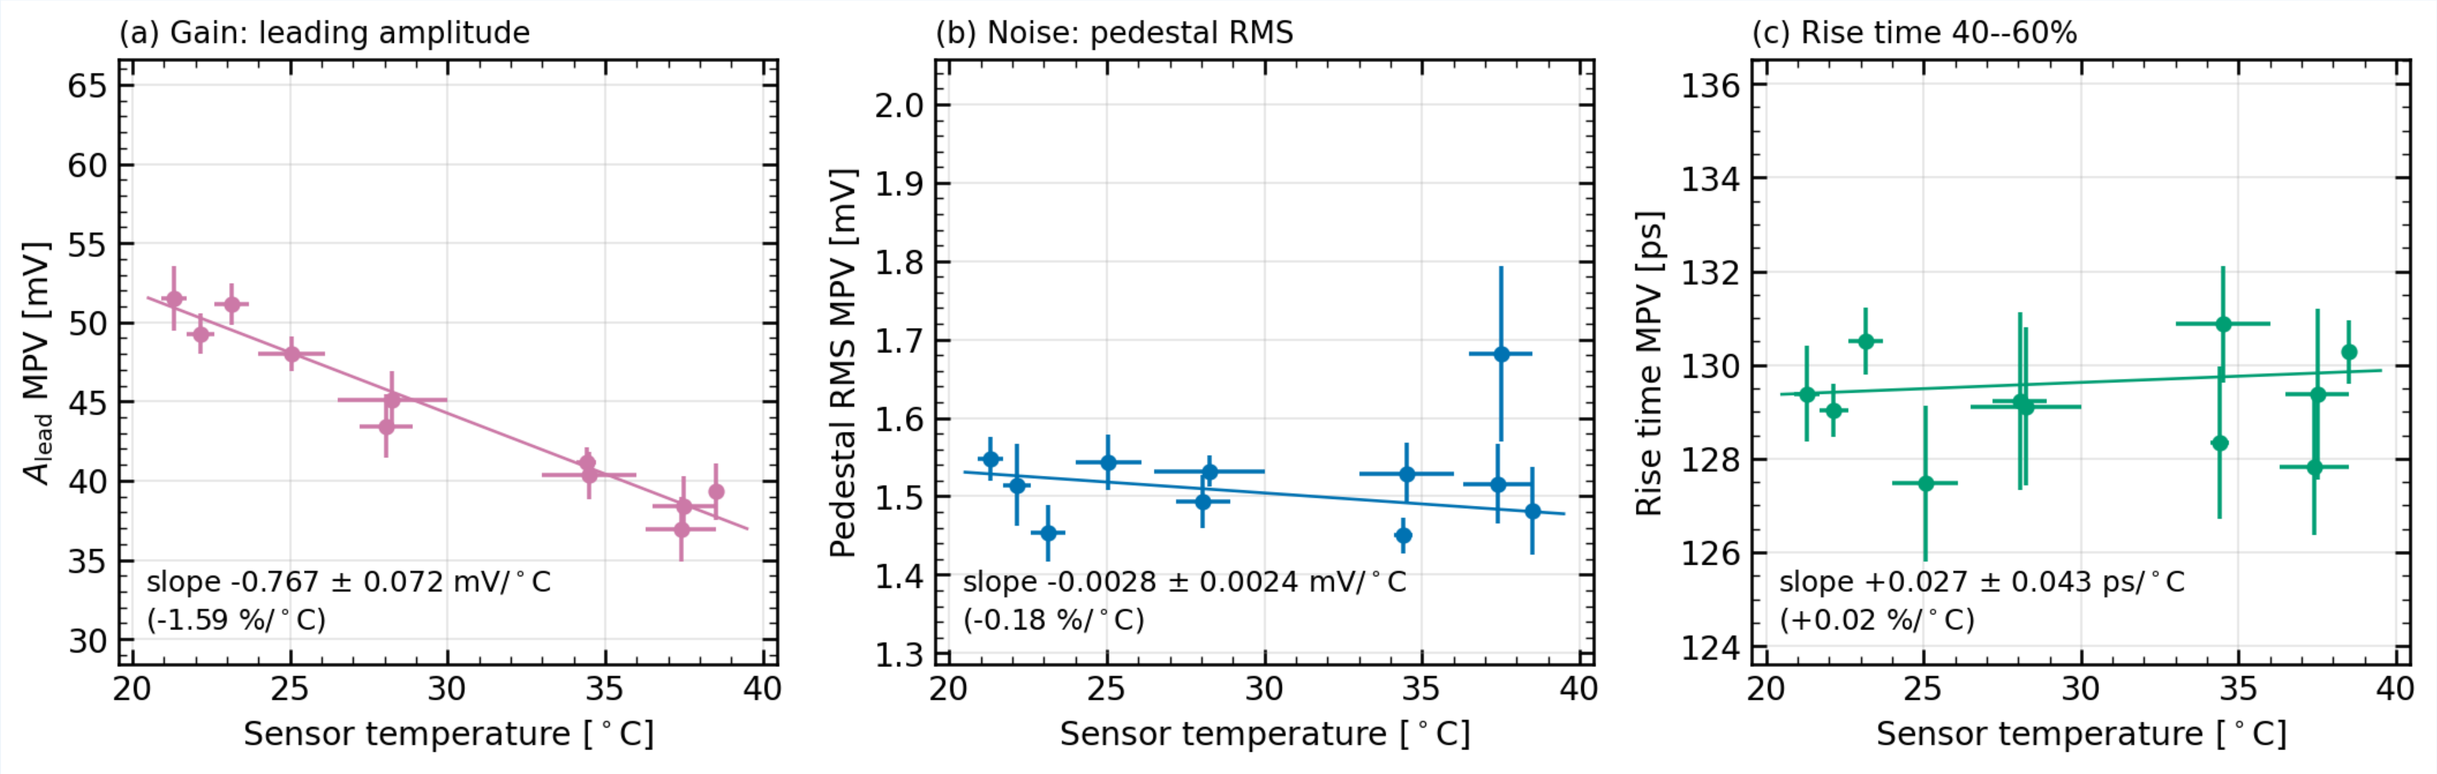

−1.59 ± 0.15 %/℃

A_{lead} MPV — the gain drop

−0.18 ± 0.16 %/℃

pedestal noise — flat

+0.02 ± 0.03 %/℃

rise time — flat

Jitter temperature dependence

σ_{jitter} grows by +0.31 ± 0.04 ps/℃ and tracks the 1/A law — the temperature effect is the gain drop

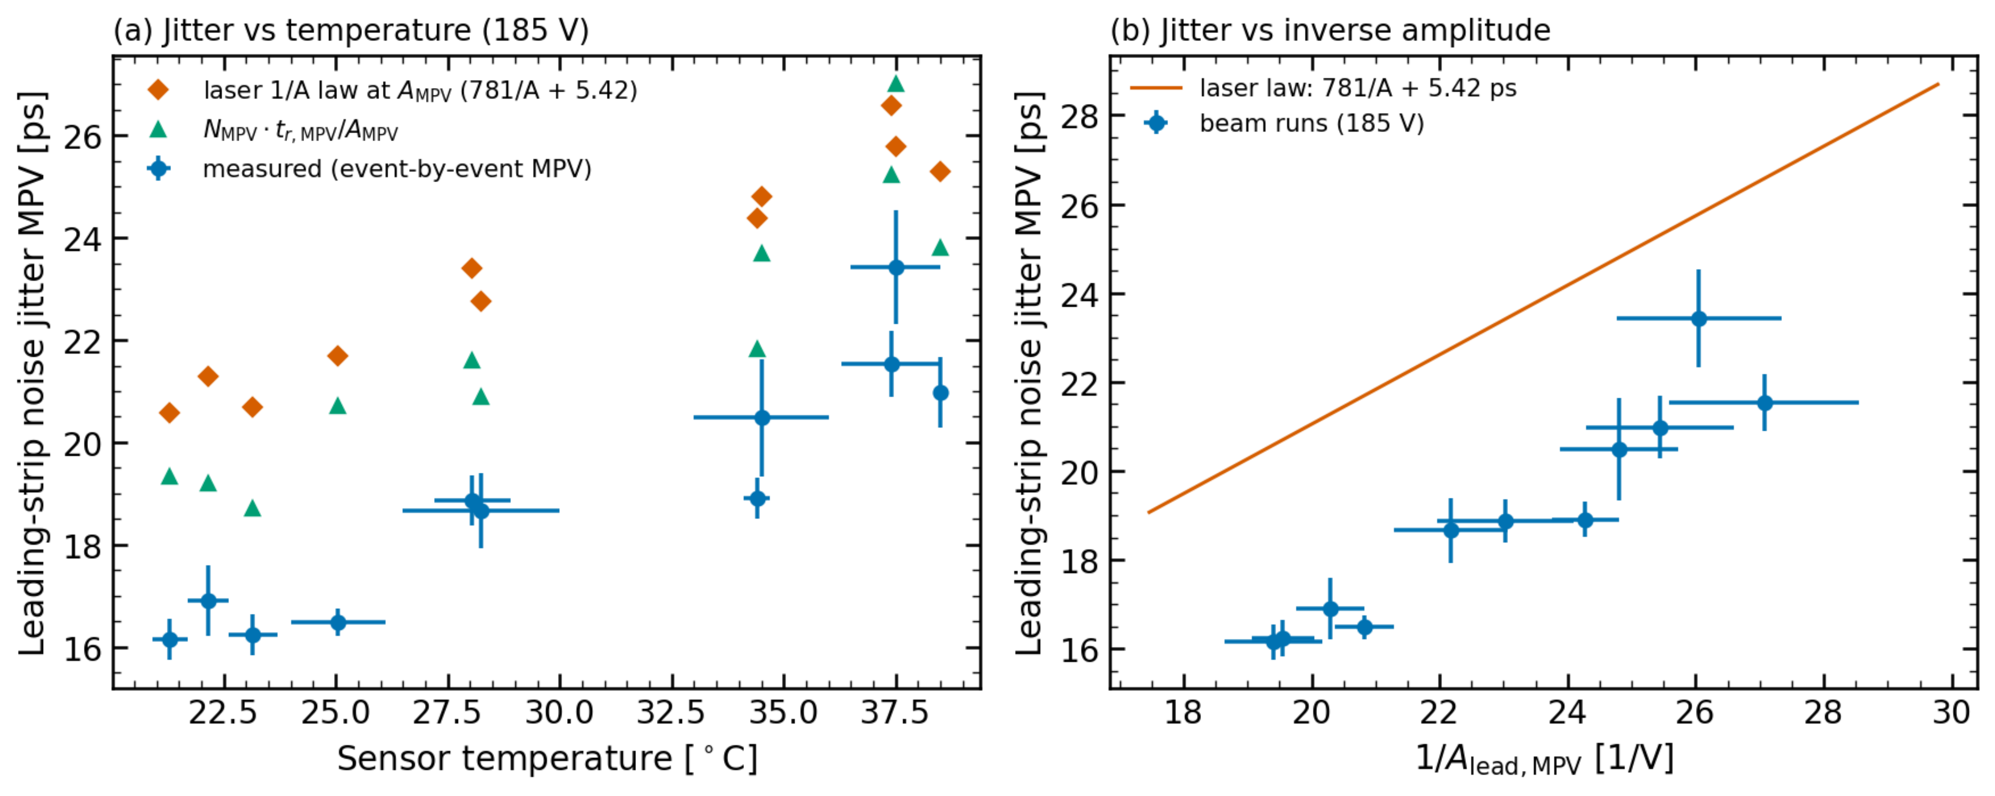

16.1 ± 0.4 ps

jitter MPV at 21 ℃ (-185 V)

+0.31 ± 0.04 ps/℃

slope of the jitter MPV

1/A law

shape follows the laser 781/A + 5.42 ps

SUMMARY

Summary

Both measured quantities are stable in temperature, and the dependence that remains traces back to the gain

$\boldsymbol{\sigma_{\mathrm{total}} = 48.6 \pm 1.3\ \mathrm{ps}}$ (reference run, -185 V, 21.3 ℃)

— from the inverse-variance two-strip combination against the MCP-PMT.

$\boldsymbol{\sigma_{total}\ slope : +0.58 \pm 0.14\ \mathrm{ps/°C}}$

over 21.3–38.5 ℃ — the 90 ps requirement is met at every temperature.

$\boldsymbol{\sigma_{jitter} = 16.1 \pm 0.4\ \mathrm{ps}}$ (at 21℃, −185V); $\boldsymbol{slope = +0.31 \pm 0.04\ \mathrm{ps/°C}}$

Noise and rise time are flat; only the amplitude falls (−1.59 ± 0.15 %/℃)

— the temperature dependence of the time response is the gain drop.

Outlook

Decompose the remaining terms with the laser TCT, and close the measurement gaps of this setup

Decomposition of σ_{total}(in progress)

Combine the 3-detector method with the laser TCT (no Landau fluctuations, known position) to separate the per-strip and common terms — preliminary and beyond the scope of this talk.

Extend below 20 ℃

All points are above room temperature; the cooled ePIC operating point is currently an extrapolation.

Bias compensation

≈1.2 V/℃ would hold the gain constant (first-order); validate with a 2-D temperature × HV map.

Temperature instrumentation

Continuous logging and a sensor-mounted thermometer, replacing two hand-read points per run.

Additional checks

Charge-sharing drift (+2.8 %/℃ in A_{sub}/A_{lead}), rate dependence, and reproducibility on a second sensor.

Backup

Why the temperature coefficient matters

Warming the sensor lowers the gain — and a lower signal amplitude directly costs timing

ePIC Barrel TOF — π/K/p separation by time of flight

requirement 90 ps · stretch goal 40 ps

Temperature ↑ → Gain ↓ → Amplitude A ↓ → $\sigma_{jit} \simeq \frac{\sigma_{noise}}{|dV/dt|}$ ↑

Impact-ionisation coefficients fall with temperature; the n+ sheet resistance and the RC charge sharing between strips may also move.

Q1

How large is the temperature coefficient of σ(Δt) after the correction chain — and can the chain remove it?

Q2

Which term carries it: per-strip jitter, per-strip residual, the common Landau term, or the reference time?

Q3

Is the requirement met over 21.3–38.5 ℃, and is the picture consistent with the laser TCT calibration?

DUT sensor and readout chain

HPK JFY25 BTOF baseline sensor, 500 μm pitch — read out by the FNAL 16-ch board into a 2 GHz scope

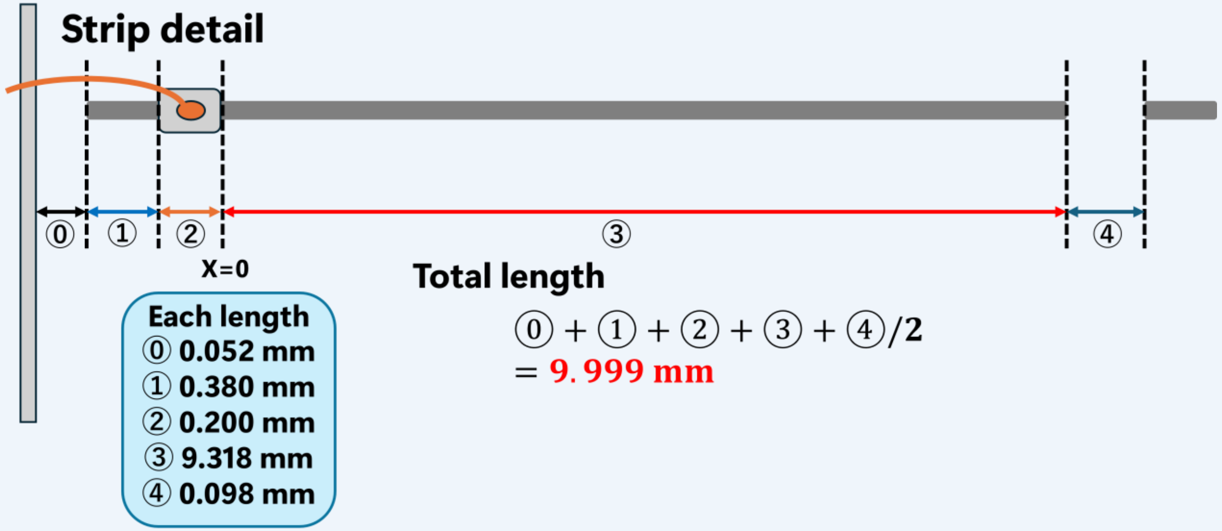

Strip geometry: 500 μm pitch, 50 μm metal width, 9.999 mm exposed length, 64 × 2 strips.

Sensor info

- HPK JFY25, 50 μm thick;
- n+ sheet 1600 Ω/□ (E-type);
- coupling 600 pF/mm²;
- SiO_2 50 nm + SiN 200 nm.

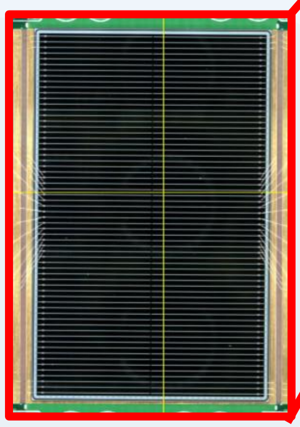

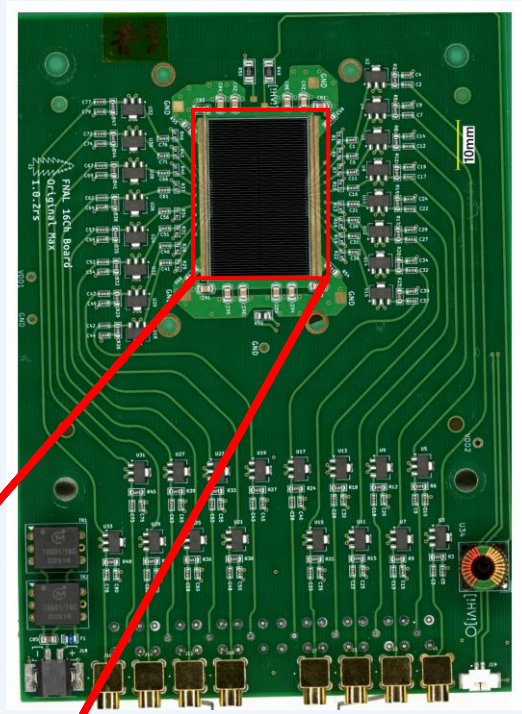

Amplifier

FNAL 16-ch board (unit #3), Mini-Circuits GALI-S66 per strip, MCX → coax.

Oscilloscope

LeCroy WaveRunner, 2 GHz / 10 GS/s / 12 bit, sequence mode.

C1–C7 = seven adjacent strips, C8 = MCP-PMT.

Vertical scale

Recovered from the .trc header: LSB 0.223 mV, quantisation noise 64 μV ≪ pedestal RMS 1.547 ± 0.028 mV.

Event selection

One cut set, applied identically to every run and every analysis chain

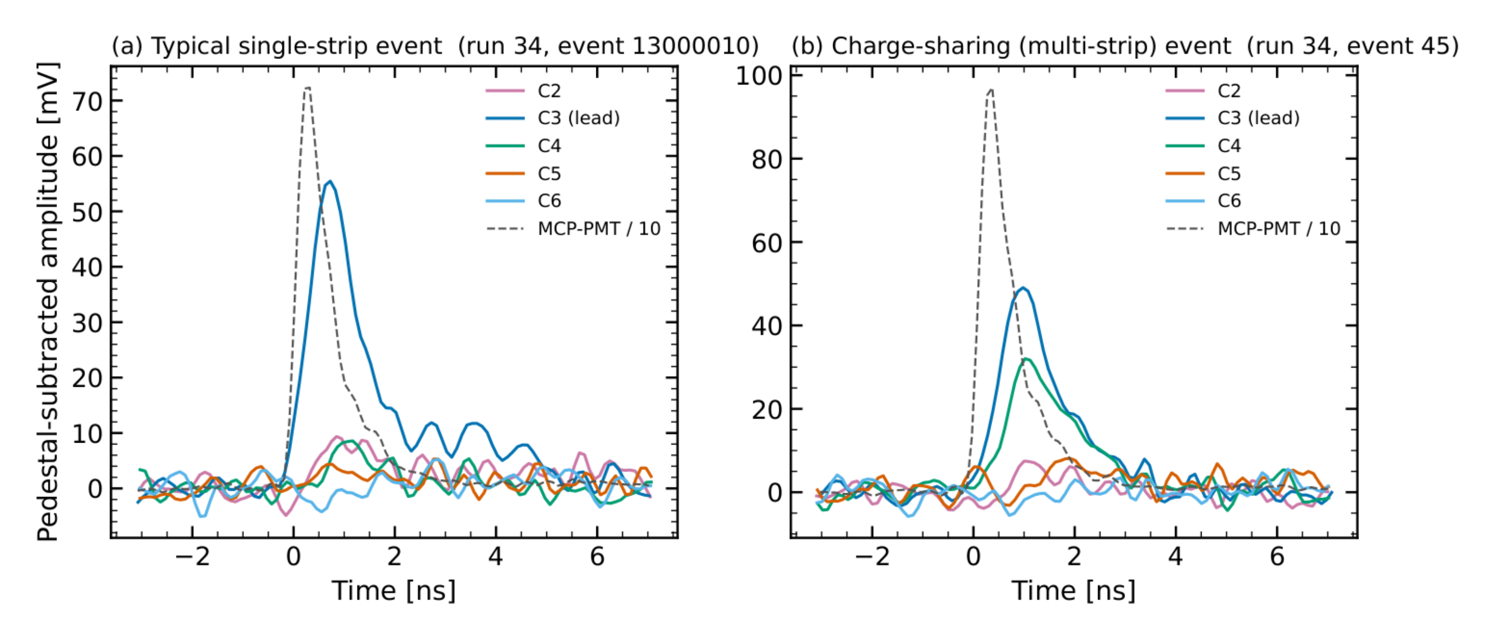

run 34 waveforms, pedestal subtracted, MCP scaled by 1/10: a typical single-strip event and a charge-sharing event.

Cuts

MCP amplitude ≥ 500 mV
Leading strip in C2–C6, 20 ≤ A ≤ 150 mV
|Δt_raw − peak| < 1 ns
Reject if C1 or C7 exceeds the leading strip

4.285 ± 0.050 %

selection efficiency (run 34: 163 999 → 7 028)

82.2 %

of the rejects are A_lead < 20 mV
— beam outside the 7-strip window

Beam spot and occupancy

Absolute hit distribution: RMS ≈ 571 μm, spread over five strips (C4 carries 25 % of the leading hits). Saturation A_lead > 150 mV is only 0.17 % at run 34.

Expected questions (1): results and physics

Each question with a short answer and the slide that shows it

Question	Short answer	Slides
Why does the resolution still rise with temperature after the corrections?	Only the gain moves: amplitude falls by $-1.59\ \%/°\mathrm{C}$ while noise and rise time stay flat, so the jitter grows along the 1/A law.	15 · 18 · 19 · 39
Does the jitter slope $(+0.31\ \mathrm{ps}/°\mathrm{C})$ explain the resolution slope $(+0.58\ \mathrm{ps}/°\mathrm{C})$?	Not by itself. In quadrature at ≈50 ps, the jitter contributes only $\sim 0.1\ \mathrm{ps}/°\mathrm{C}$, so other terms must also move. The decomposition is ongoing with the laser.	19 · 39 · 40
What limits the ~50 ps — what is in the remaining budget?	The 3-detector split gives a large common “bucket” (MCP ⊕ beam bucket ⊕ common Landau); Landau + common residual is 33–36 ps.	13 · 38 · 42
Is the requirement met at the real (cooled) ePIC operating point? (< 20°C?)	All measured points are above $21\ °\mathrm{C}$, so the cooled point is an extrapolation. Extending the scan below $20\ °\mathrm{C}$ is a stated next step.	15 · 22 · 34
Can the gain drop be compensated in operation?	To first order yes: raising the bias by $1.17\ \pm\ 0.19\ \mathrm{V}/°\mathrm{C}$ holds the gain constant. The HV series shows the temperature ↔ gain equivalence.	22 · 35

Expected questions (2): method

Each question with a short answer and the slide that shows it

Question	Short answer	Slides
Are the in-situ resolutions optimistic, since the LUTs come from the same data?	Yes, slightly — that is why every number is quoted twice. The half-split test gives ~60 ps for the leading strip; the laser even/odd LUT split shows no over-training.	11 · 12 · 33 · 41
Could one global LUT be used for the whole temperature range?	No. Applying the run 34 time-walk LUT to run 43 costs 52 ps RMS, so the LUTs are derived per temperature.	31 · 33
Why CFD at 50 % and not another fraction or a fixed threshold?	A CFD scan finds the minimum at 65 %, only 0.6 ps better than 50 %. The best fixed threshold (LE 10 mV) is a factor 1.38 worse.	10 · 31
How reliable is the temperature axis?	Two hand-read points per run; the mean in-run drift of $1.4\ ^\circ\mathrm{C}$ gives $\pm 0.7\ ^\circ\mathrm{C}$ on the slopes. No in-plane map; continuous logging is planned.	6 · 33 · 34
How well is the hit position known without a tracker?	η_3 with the laser effective pitch of 686 μm; an independent η-CDF calibration agrees to 4 μm rms, and the choice changes the chain σ by ≤1 ps.	5 · 9 · 33

Same quantities, laser TCT vs this beam test

Jitter, single-strip and sensor terms agree; the beam adds the Landau and common terms on top

Quantity	Laser TCT (matome, −185 V, ~21 ° C)	Beam run 34 (−185 V, 21.3 ° C)	Comment
Leading amplitude MPV	44.7 ± 0.1 mV (uniform scan)	51.5 mV	laser has no Landau fluctuation
Noise RMS / rise time	1.67 mV · 610–650 ps	≈1.5 mV · 636 ± 4 ps (MPV)	same amplifier chain and regime
Leading-strip jitter MPV	17.3 ± 0.2 ps	16.1 ± 0.4 ps	consistent via the noise ratio
Jitter vs amplitude	source of the law: 781/A + 5.42 ps (MPV)	follows the same law (~20 % offset)	offset from the slope definition
Single strip, no Landau / floor	$\sigma(\Delta t)/\sqrt{2}$ = 36.6 ± 0.2 ps	σ_lead = 35.9 ± 4.0 ps	same level
Two-strip sensor σ	23.8 ± 0.4 ps	σ_sensor = 30.8 ± 2.5 ps	difference = beam per-strip residual
System σ vs reference	— (no reference detector term)	50.2 ps (60.0 half-split)	Landau + common 33–36 ps add here
Propagation speed v_eff	0.933 ± 0.003 μm/ps (pair η)	0.93 (pair η) / 0.84 (η_3 axis)	η_3-axis nonlinearity — open item
Effective pitch	686 ± 37 μm (measured)	686 μm adopted (transferred)	η-CDF check agrees to 4 μm rms
CFD fraction	50 % near-optimal	50 % near-optimal (65 % only −0.6 ps)	

Run list used in this analysis

Main 185 V series in bold; excluded runs and reasons on the next slide

run	HV [V]	T [° C]	n_clean	series
16	185	37.10	4 819	excluded
18	165	36.85	2 799	sub 165 V
21	185	34.40	5 478	main 185 V
25	165	31.40	3 688	sub 165 V
27	185	25.05	4 907	main 185 V
29	165	26.05	3 623	sub 165 V
31	185	28.05	4 852	main 185 V
32	185	23.15	5 055	main 185 V
33	185	22.15	5 375	main 185 V
34	**185**	**21.30**	**5 280**	**main 185 V · HV series · reference**
35	165	20.80	4 060	sub 165 V · HV series
36	145	20.65	2 314	HV series
37	190	20.95	4 267	HV series
38	185	28.25	2 281	main 185 V
40	185	34.50	2 139	main 185 V
42	185	37.50	2 104	main 185 V
43	185	38.50	2 051	main 185 V
44	185	37.40	4 076	main 185 V

Excluded runs and the cut flow

run 34 as the example of the selection

Excluded runs

16	settings not recorded (threshold column empty; first run after the collimator was installed)
17, 26, 30, 39, 41	low statistics (n_clean < 2000)
22	end temperature missing in run_info.csv
23	truncated data (2 files/ch)
24	oscilloscope variable gain was ON
28	marked skip in run_info.csv
45	195 V at 35.5 ° C — no temperature-matched partner for the HV series

Cut flow — run 34

Triggers (MCP entries)	**163 999**
− MCP < 500 mV	20 704
− leading strip < 20 mV	134 754
− leading strip > 150 mV	277
− Δt window ±1 ns	24
− C1/C7 veto	1 212
Selected	**7 028 (4.285 ± 0.050 %)**
Clean events for the chain	**5 280**

Data quality vs. temperature

Selection efficiency, noise and rise time are stable; the saturation fraction falls with the gain

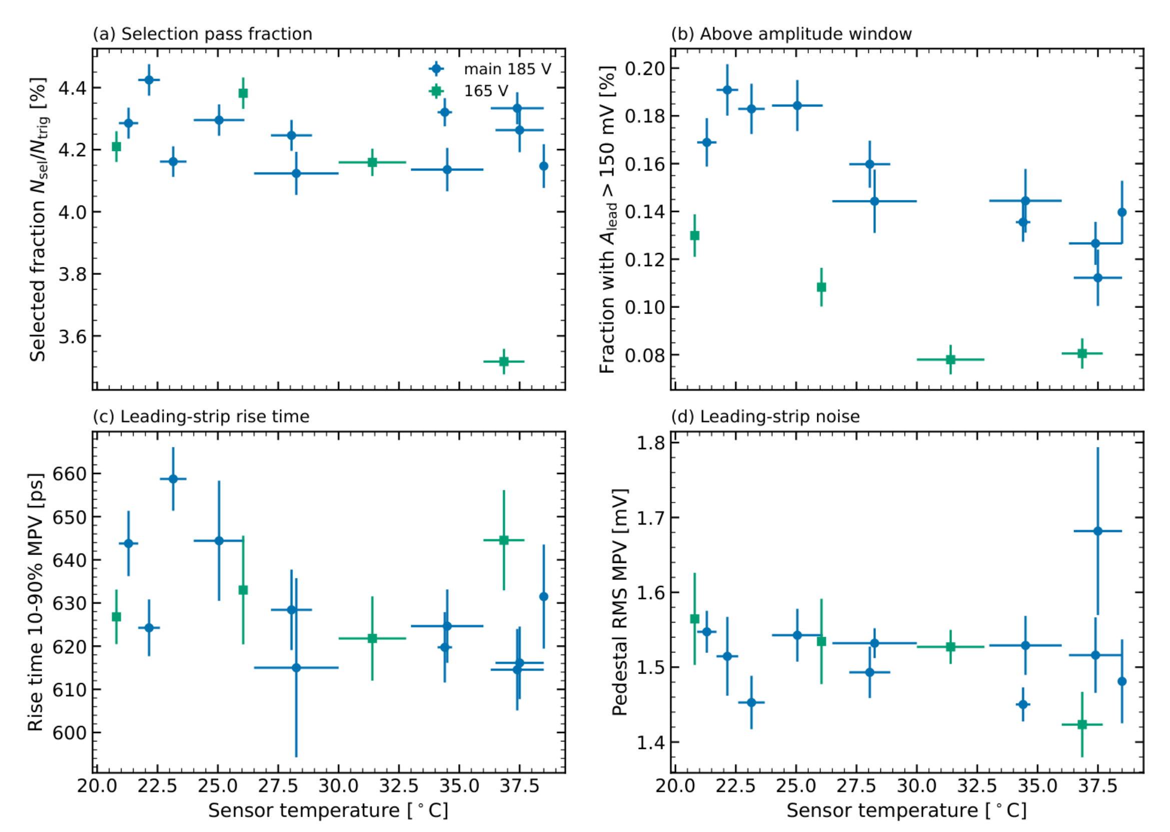

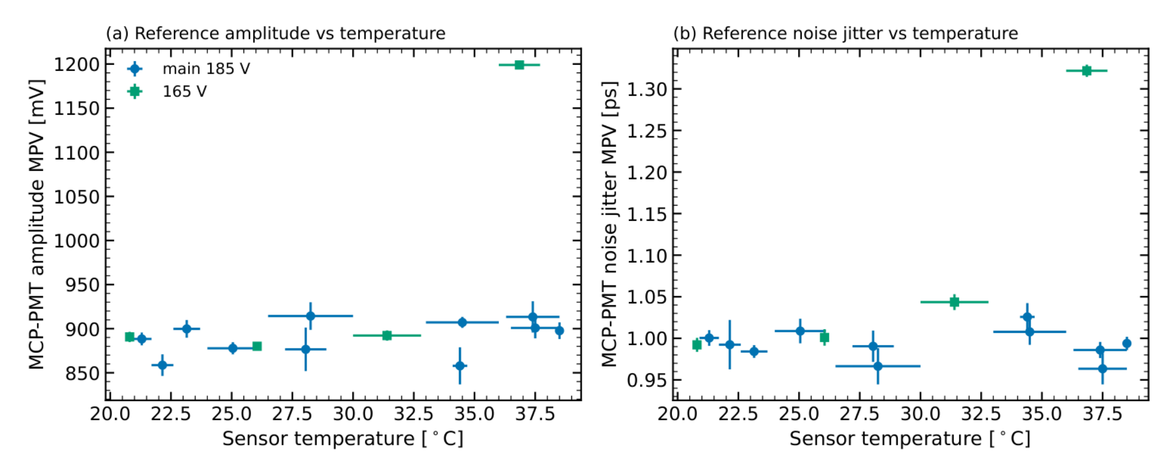

Reference MCP-PMT

Amplitude MPV +0.147 ± 0.073 %/° C, jitter MPV −0.000 ± 0.001 ps/° C — stable over the same range.

Sensor quantities

Rise time MPV 636.1 ± 4.2 ps at 25 ° C, slope −1.40 ± 0.54 ps/° C. Pedestal noise −0.0028 ± 0.0024 mV/° C, i.e. constant — the electronic noise is set by the amplifier, not by the sensor leakage.

PRELIMINARY

LUT transferability and CFD choice

The LUT shape moves with temperature, and CFD50 is near-optimal

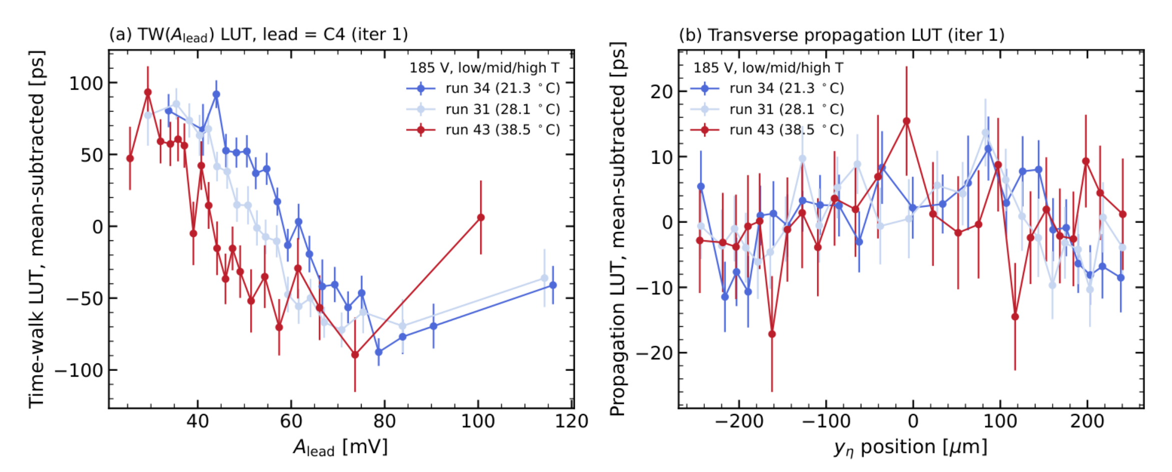

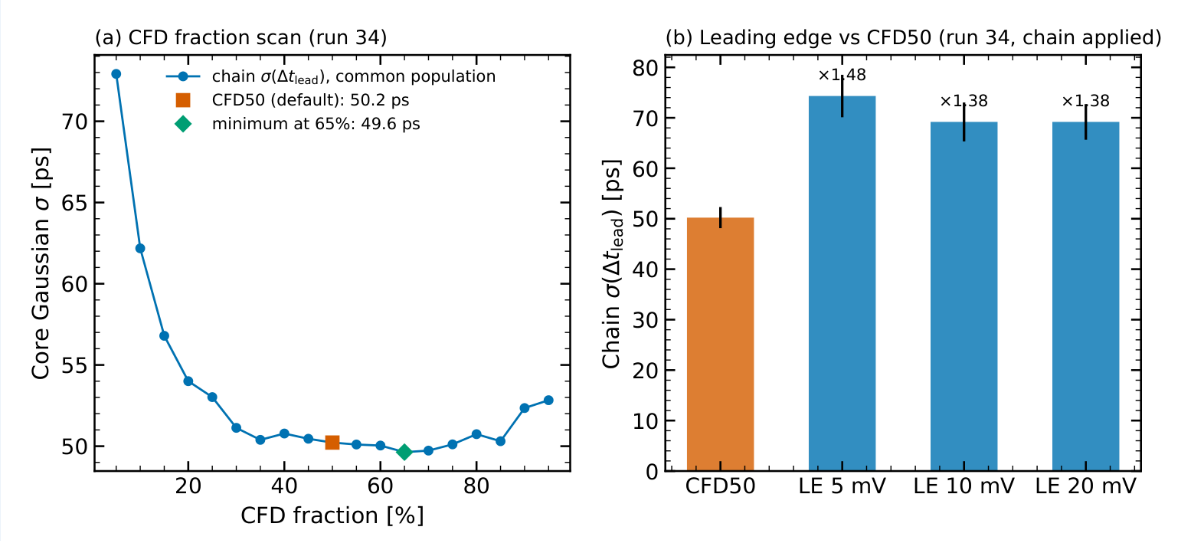

Correction LUTs vs. temperature

Applying the run 34 time-walk LUT to run 43 costs 52.1 ps RMS (C4): a single global LUT is not transferable across the range. The propagation LUT is small and its temperature dependence is not significant.

Arrival-time definition

CFD fraction scan: minimum at 65 % (49.6 ps), only 0.6 ps below the default CFD50 (50.2 ps). The best fixed-threshold leading edge, LE 10 mV, gives 69.2 ± 3.9 ps — a factor 1.38 worse. CFD50 is used throughout.

PRELIMINARY

Robustness: full panels

Attribution, statistics correlation, charge sharing and in-run slices

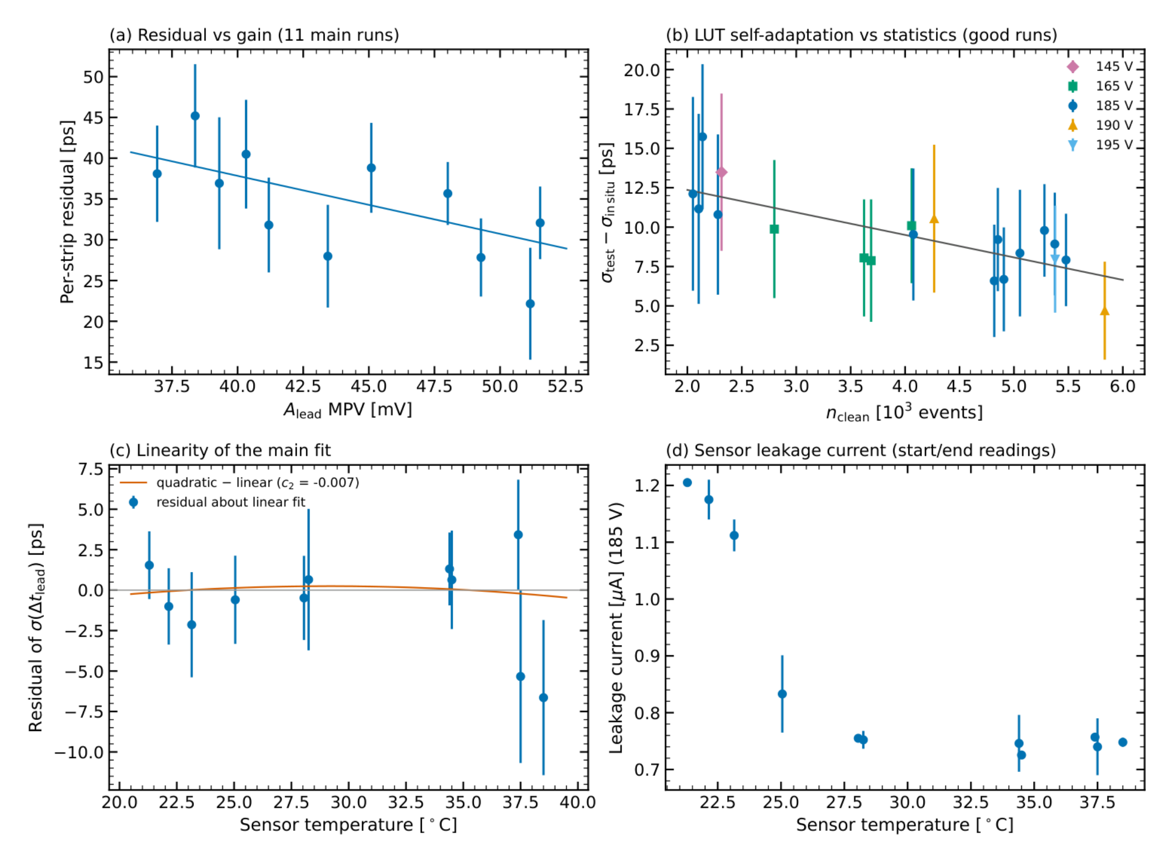

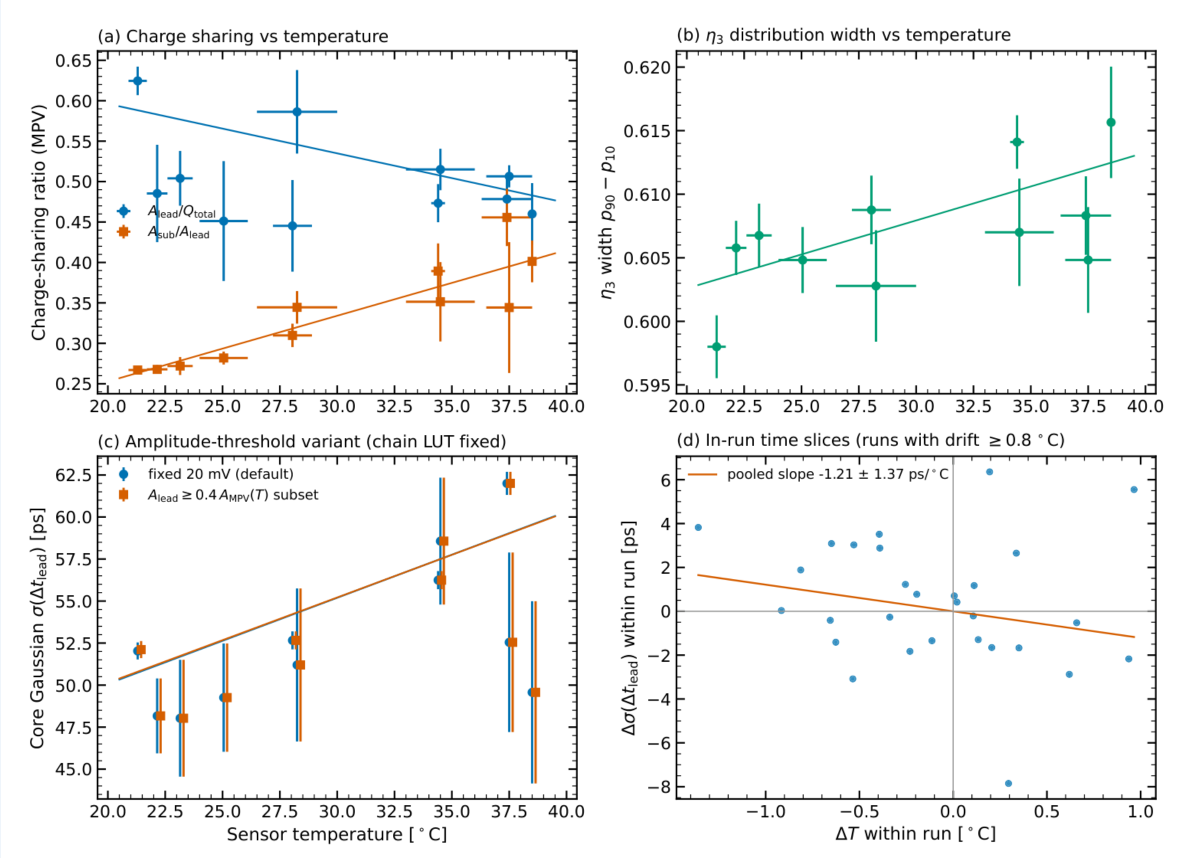

PRELIMINARY

Systematics — full table

Items that were not evaluated are stated as such

Source	Affected quantity	Size	Evaluated from
LUT self-adaptation	σ(Δt_lead)	**9.8 ps**	in-situ vs. half-split test
σ estimator choice	Landau term	**11.0 ps**	core Gaussian vs. truncated RMS ±2–3σ
σ estimator choice	bucket	**10.7 ps**	same
CFD fraction	σ(Δt_lead)	**0.6 ps**	CFD50 vs. best fraction (65 %)
3-detector allocation	σ_lead / bucket	**5.2 / 6.0 ps**	chain vs. symmetric LUT method
Position variable	chain σ	**≤ 1.0 ps**	η-CDF vs. η_3 linear scale
Run-to-run reproducibility	σ vs. T fit	**2.9 ps (RMS)**	fit residuals, 11 main runs
Temperature axis	slopes	**± 0.7 ° C**	in-run drift; in-plane gradient not evaluated
TW LUT shape vs. T	chain transferability	**52 ps (RMS)**	LUT difference run 34 vs. run 43 (C4)
Amplitude window	slope vs. T	**−0.00 ps/° C**	threshold ∝ A_MPV(T) variant
Laser 1/A transfer	jitter prediction	**~20 %**	definition-matched ratio 0.76 / 0.84
σ_MCP	bucket decomposition	**range 5.5–10 ps**	Sr-90 floor vs. literature (a range, not an
σ_long	bucket decomposition	**range 3.4–10 ps**	tracker-based laser analysis
Thermometer calibration	absolute T	**not evaluated**	single hand-held sensor
Rate / beam intensity	all	**not evaluated**	no rate scan available

Limitations of this analysis

Stated explicitly so the numbers are not over-interpreted

- The temperature scan is essentially a single bias point: 185 V has 11 runs, 165 V only 4, and 145 V is unusable at high temperature.
- Without a tracker, the absolute hit efficiency and the absolute position resolution cannot be measured; the selection efficiency and the η_3 checks stand in for them.
- Temperature comes from two hand-read points per run — no continuous log and no in-plane distribution.
- run 16, the hottest 185 V point, is excluded because its settings were not recorded; the official summary slope includes it.
- No rate or beam-intensity dependence was tested.
- The sensor die ID, the MCP-PMT model and its bias during the beam time have no primary record and are being traced. If the die cannot be identified, comparing this temperature dependence with another die is not possible.
- All measured temperatures are above room temperature, so applying the result to a cooled operating point is an extrapolation.

Temperature ↔ gain equivalence

Lowering the bias reproduces the same degradation as heating — a ~1.2 V/℃ bias ramp would hold the gain constant

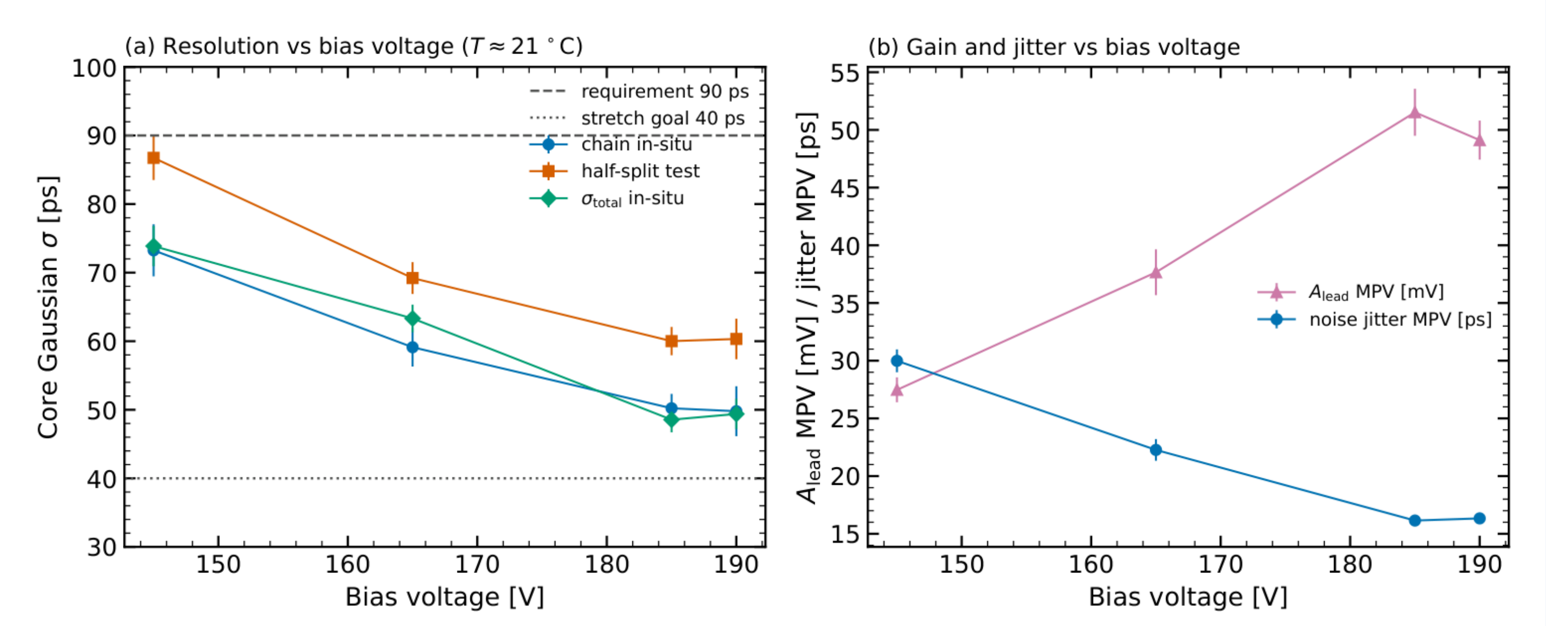

HV series at T ≈ 21 ° C (runs 36/35/34/37). run 45 at 195 V is excluded — no temperature-matched partner.

145 V → 185 V

gain ×1.88, σ(Δt_lead) 73.3 ± 3.8 ps → 50.2 ± 2.1 ps

165 V sub series

+0.80 ± 0.26 ps/℃, steeper than 185 V — the lower the gain, the harder the 1/A law bites.

Constant-gain bias tracking

d ln A/dV = 1.37 ± 0.18 %/V vs. −1.59 ± 0.15 %/℃

→ 1.17 ± 0.19 V/℃ (first-order estimate)

Which term carries the temperature dependence?

Per-strip terms and the jitter rise with temperature; the common bucket is flat within errors

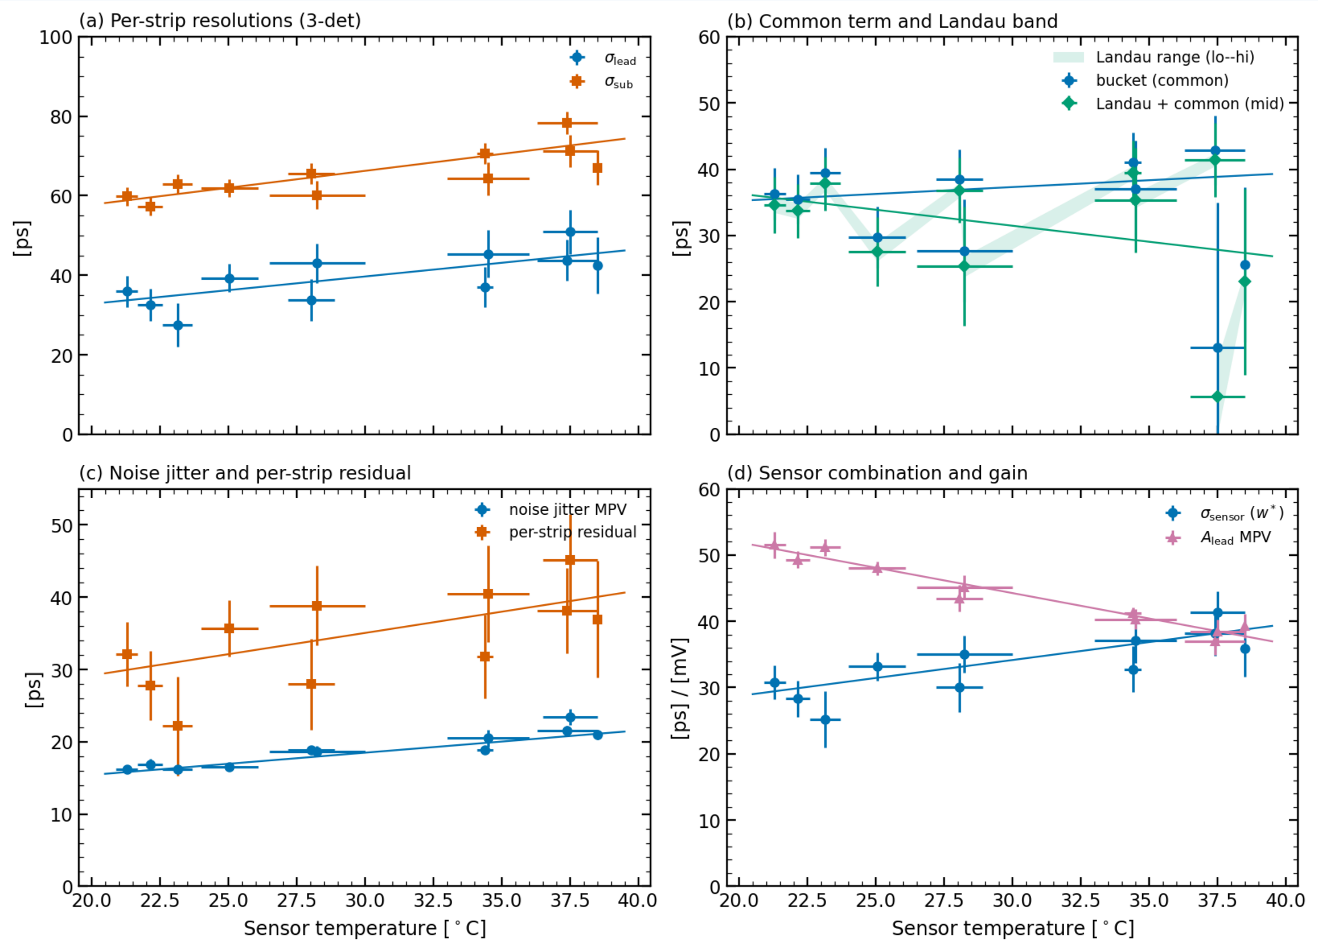

PRELIMINARY

Slope closure: jitter alone is not enough

Jitter-only growth gives +0.14 ps/° C; adding the per-strip residual reconstructs +0.66, matching the measured +0.69 ± 0.24

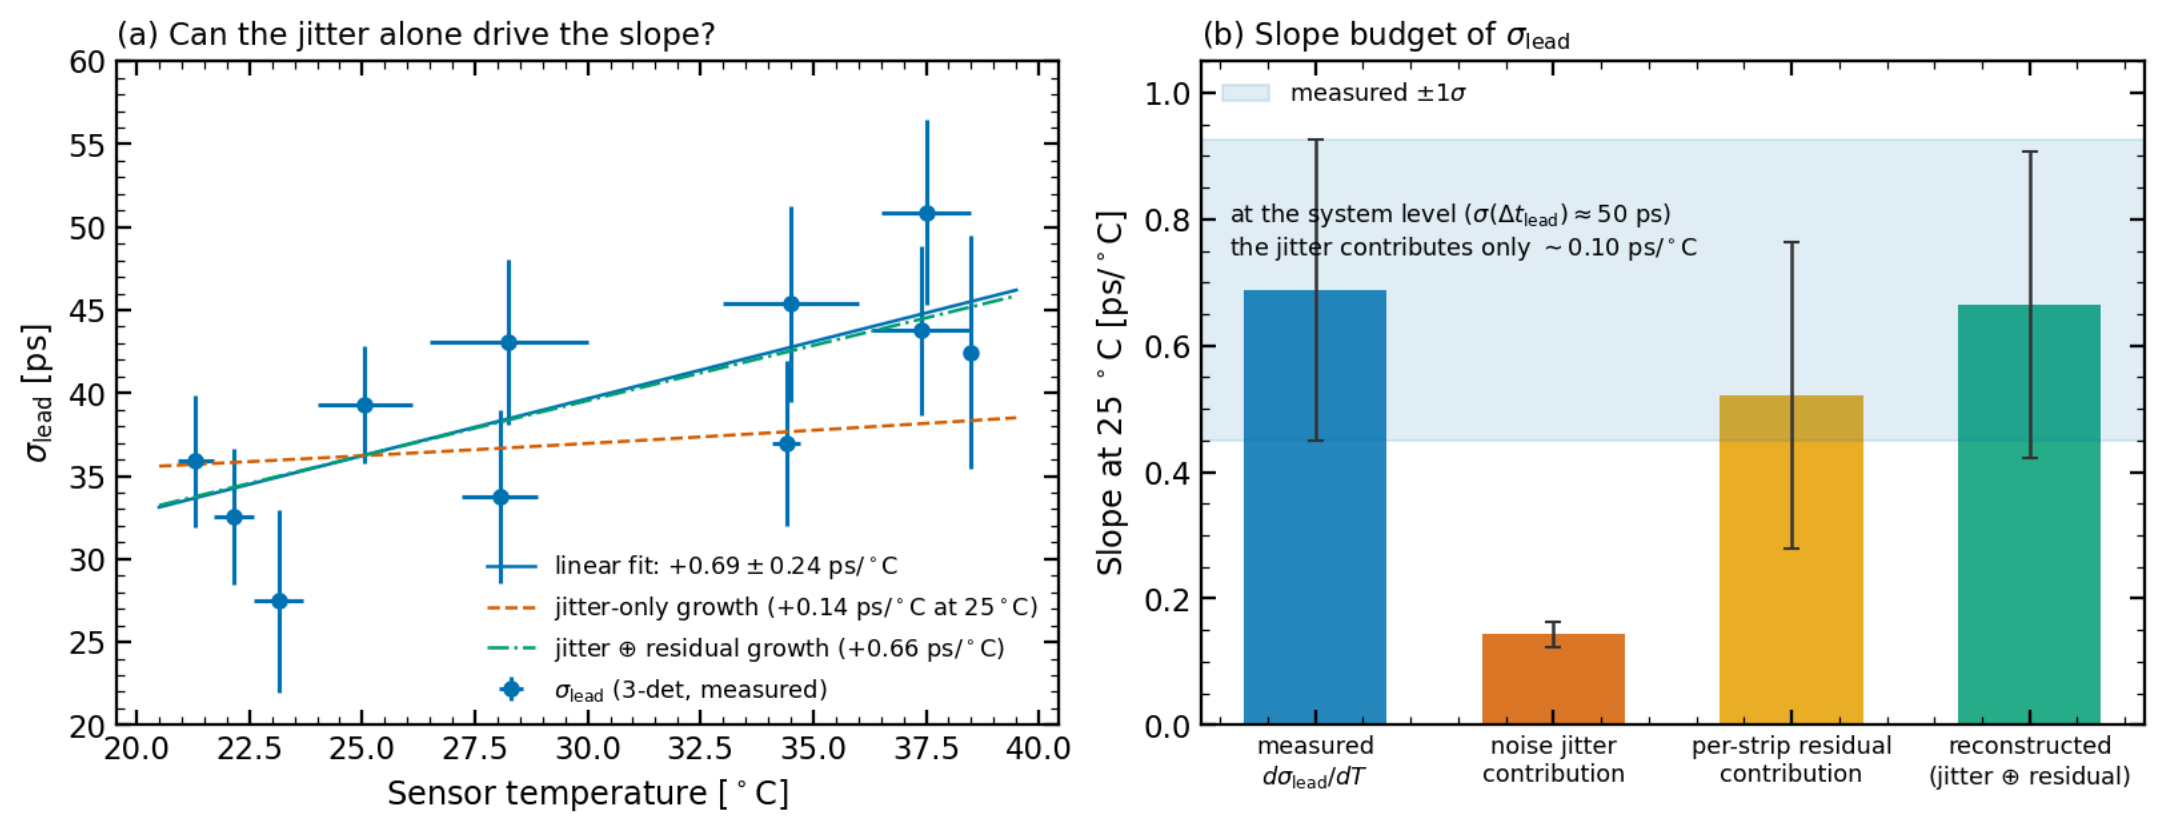

PRELIMINARY

The slope survives the half-split test

The half-split test shifts the level by ~9 ps but not the slope (+0.55 in-situ vs +0.57 test) — the trend is not a LUT artifact

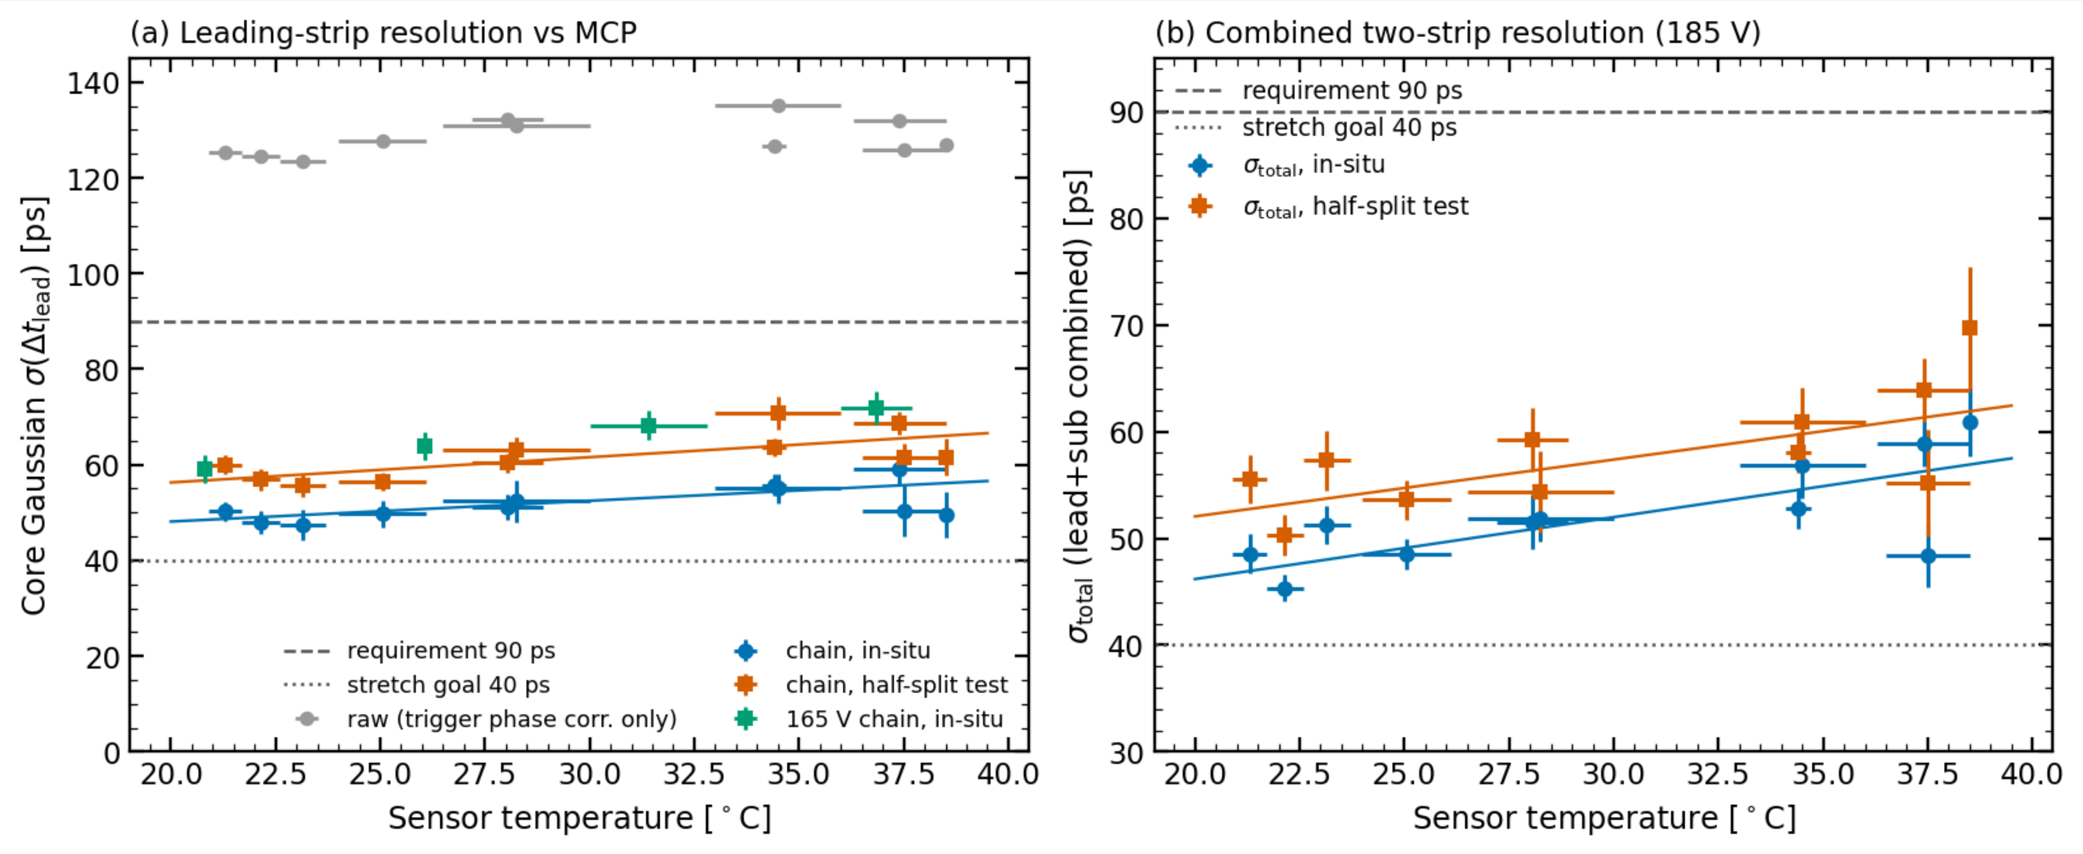

PRELIMINARY

Splitting the Landau term with the laser and Sr-90

The beam-only excess is amplitude-independent: 28.9 ± 4.4 ps per strip — a Landau shape fluctuation, not noise

Splitting the Landau term: common vs uncommon (run 34 + laser + Sr-90)

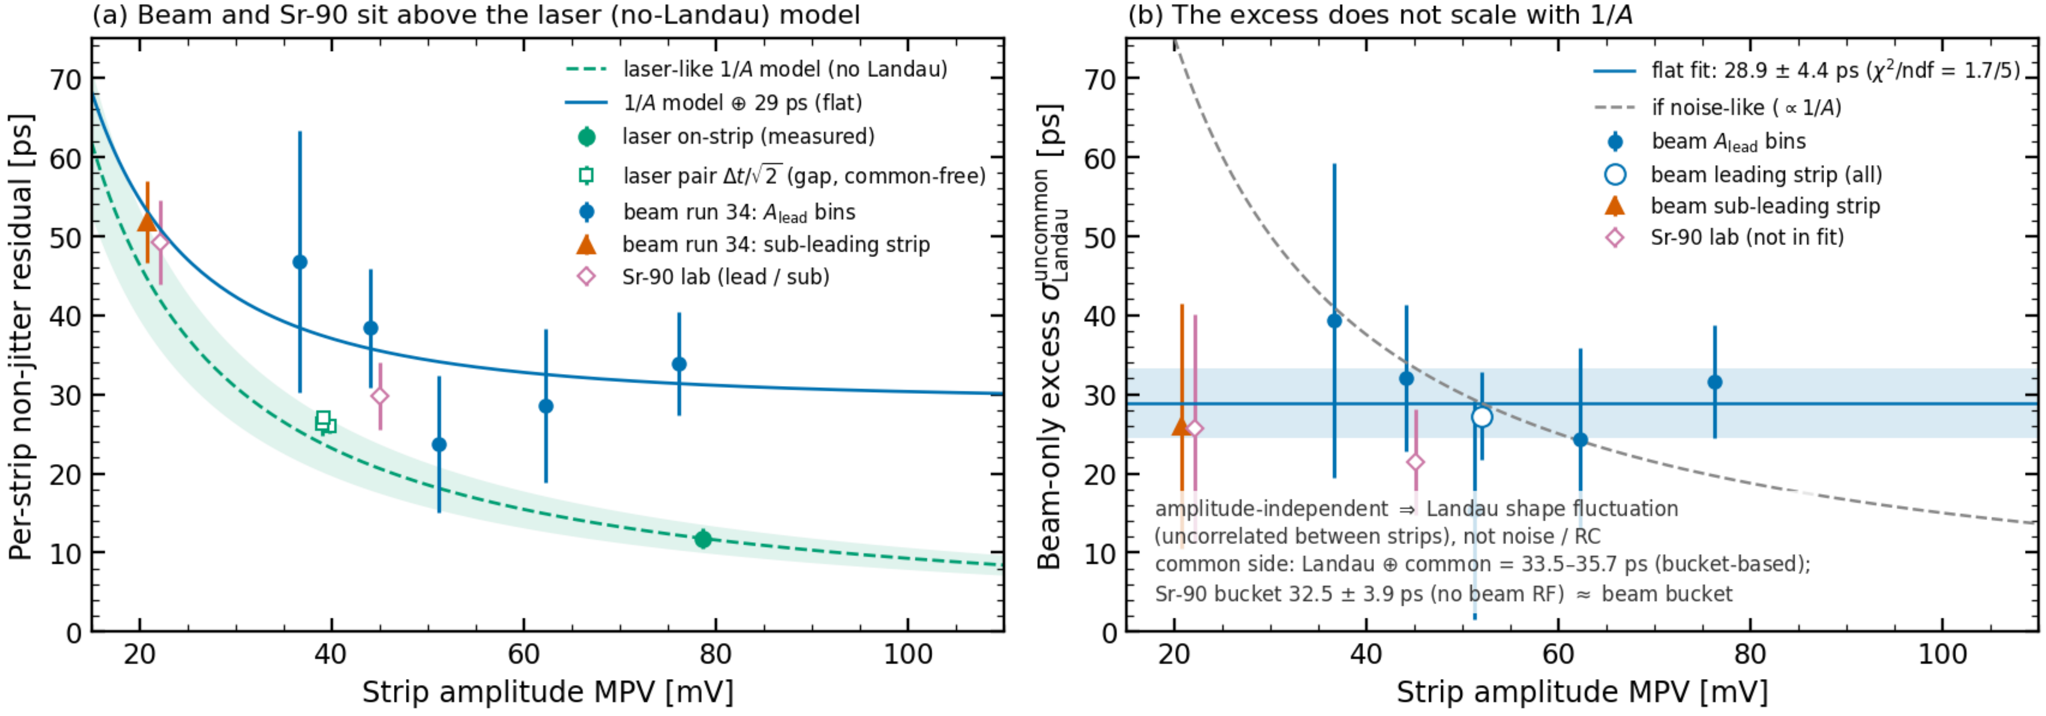

PRELIMINARY