

Development and System Characterization of High-Speed Readout Electronics for Cluster Counting in Drift Chambers

INSTR2026 - October 12-16, 2026, Hefei

Hongbin Liu (hbliu@ihep.ac.cn) - on behalf of the drift-chamber group


Institute of High Energy Physics, Chinese Academy of Sciences
Q. Li, D. Xu, M. Ji, W. Pan, L. Wu, H. Zhang, Y. Zhao, M. Dong, Z. Sun



INSTR 2026


October 12-16, 2026 | Hefei, China

Why dN/dx - and the Scope of This Work




- A drift chamber is a **back-up option for the TPC** of the CEPC detector
- **Scope: R&D on the dN/dx readout method**, carried out on a 12-layer, 120-channel prototype

- **dE/dx saturates** at high momentum, where the resolution stops improving
- **Counting primary clusters (dN/dx)** cuts the ionisation-fluctuation term that dominates there




dE/dx : total charge




dN/dx : clusters per unit length



The Drift-Chamber Prototype



- **12 layers x 10 cells = 120 sense wires** (20 μm gold-plated tungsten); 416 field wires; **18 x 18 mm²** cells
- He/iC₄ H₁ o 90/10 at **1840 V**, gas gain $\sim 10^5$, drift velocity **25 $\mu\text{m}/\text{ns}$** , max drift 9 mm
- **60 cm** active length; the prototype exists to develop and verify the **dN/dx readout**





Requirement	Derived from	Design target	Measured
Channel count	12 layers × 10 cells	120 channels	120 built; 40-ch crate tested
Front-end architecture	avoid C_det broadening	low-Z + wideband	implemented
Transimpedance gain	signal amplitude vs noise	> 1000 Ω at the ADC	2232 Ω full chain
Full-chain bandwidth	≤ 10 % added rise, 2 ns edge	≥ 400 MHz	550 MHz
Linear dynamic range	MIP to large deposits	≥ 50 μA	~ 700 uA
Equivalent input noise	resolve clusters (loaded)	≤ 1.6 μA rms	0.83 μA rms
ADC	≥ 2 samples across a 2 ns rise	≥ 1 GS/s, ≥ 12 bit	1.3 GS/s, 14 bit
Electronics time jitter	≤ 3 % spatial degradation	< 1 ns	0.80 ns within-channel



- 120 channel waveform capture with a window ~ 1.6 us, with fix trigger to window latency

System Composition and Scalability




10-channel pre-amplifier board (left) and the 40-channel crate in the lab (right)



- **12 units x 10 channels = 120 channels**; the BEE is built from three standardised **40-channel** subsystems
- **FEE at the detector end** (12 pre-amplifier boards); 120 x SMA to the BEE
- The 120-channel system uses **replicated units**; full-detector integration remains future work

The Analog Front End: Current to ADC





1233 Ω
FEE board alone (24.66 V/V)

2500 Ω
after $\times 2$, at the ADC driver output

2232 Ω calculated
including the ADC input network

2248.7 Ω measured
full chain, single channel

One 10-Channel Readout Unit



- **One ADC mezzanine** (5 x AD9695) + **one ZU15EG** board = **10 channels**
- ADC: AD9695, 1.3 Gsps, 14-bit, 1.6 Vpp DR, JESD output
- LMK04828 clock generator;
- **Analog front end** on the FEE board in front of the mezzanine: two LMH6629 gain stages
- **8 x 25 cm**, 12 V DC, 30 W

10

channels per unit

1.3 GS/s

ADC, 14 bit


10 GbE

UDP waveform data

1 GbE

TCP slow control

Control and Data Acquisition



- The **ZU15EG processing system (PS)** runs the **control firmware** on every board: configuration and monitoring over **1 GbE TCP**
- All waveform data leave over **one 10 GbE UDP optical link** per board
- Each unit instantiates **two IPs - one control, one data**; the data IP's destination address is configurable, so the 12 links can be received by different DAQ NICs and acquisition is not limited by one port

12 x 1GbE

TCP control

12 x 10 GbE

UDP data links




3 x MPO/LC

fibre per crate

2 IPs / unit

one control, one data

Crate Characterization: Noise and Response



0.470 mV
mean noise RMS, 36 ch

5.57 μV
channel-to-channel sigma (1.2 %)


~ 550 MHz
full-chain -3 dB (target 400)

0.82 %
gain dispersion, 120 ch

- Crate, **no detector load**: 36 baseline channels x 129 events x 2000 samples; the 4 channels carrying the **10 MHz test signal are excluded**



Position from the drift time inside the acquisition window

The leading edge marks the arrival of the first drifted electron. Its position measured from the window start gives t_{drift} , and $x = v_{\text{drift}} \times t_{\text{drift}}$.



This is only well defined if the window start sits at the same physical time after the trigger in every event: any jitter δ there becomes a position error $v_{\text{drift}} \cdot \delta$. Hence the fixed-latency requirement.

Fixed Latency for the Position Measurement



87.90 ns

trigger to ADC, 197 events

1.58 ns

range of channel mean latencies

0.39 ns


range of board mean latencies

0.09 ns

cycle-mean range, 3 cycles

- Position comes from **drift time** (25 $\mu\text{m}/\text{ns}$, i.e. **1 ns = 25 μm**), so timing is important for dX reconstruction
- The **1.58 ns channel-to-channel spread ($\approx 40 \mu\text{m}$) must be calibrated out;**
- Those offsets stayed put over **the three power cycles tested (0.09 ns)** - that stability is exactly what makes a fixed per-channel calibration valid

Data Throughput and Link Headroom



- The observable is the **event rate reconstructed from the received packets**, not the trigger rate
- **No event with a missing packet, up to a 70 kHz input trigger rate**; missing packets at higher rates
- At **4008 Byte per event** (10 channels per event) the left axis converts to the right axis: 10 kHz needs **3.2 Gbit/s**, and the measured **461 MB/s** ceiling is reached at **11.5 kHz**

70 kHz

formats with no missing packets




11.5 kHz

max reconstructed event rate

31 kHz

10 Gbe nominal payload limit

Cosmic-Ray Operation



- **Two-layer scintillator coincidence** gives the global time reference
- Chamber at 1840 V, gas gain $\sim 10^5$; **8 layers instrumented, 40 readout channels connected**
- **~ 300 K cosmic events** analysed
- Mean noise RMS **1.4 mV**, mean SNR **150.46**

~ 300 K

cosmic events analysed

1.4 mV

mean noise RMS


150.46

mean per-pulse SNR, cosmic

40


readout channels connected

Beam Test at the CERN PS T9 Line




1.61 mV

noise RMS




30.12 mV

signal amplitude




1.89 pC

collected charge





19.04

mean per-pulse SNR,
beam



1.12 ns


mean rise time, ± 0.63 ns




Online Cluster Counting: Why and How

- At **5 % occupancy** one full-detector trigger is **~7 MB**, i.e. **~7 GB/s** at 1 kHz if every waveform is read out - so peak finding and clustering have to move **into the front-end ASIC**
- **Per pulse**: 2nd-difference candidate positions, then a **CNN** suppresses the noise peaks. **Per event**: the cluster count is regressed from the **peak count, times and intervals**





per-pulse processing



per-event aggregation



Online Peak Finding on 10 GeV Data



- **Preliminary:** the peak-finding model applied to a **10 GeV event sample** (Data0040)
- **12, 20, 28 and 36 peaks** in the four example events, from four different boards and channels
- The per-pulse stage returns the **peak positions**; the per-event stage turns them into the **cluster count**

12 - 36

peaks per example
event


4

events shown
(boards 9-12)


Red markers show where the peaks were found. The per-pulse stage outputs the position of each peak and the count n is carried in the legend; the per-event stage converts them into the cluster count. Offline model result, not the RTL implementation.

Online Cluster Counting: RTL Implementation


(a) Per-event processing time



(b) Total utilisation on the ZU15EG



(c) On-chip power, total 12.37 W



0.213 ms

RTL simulation, per event

73.2 x

faster than Python/NumPy

100

events, final count agrees

29.8 % / 63.2 %

total LUT / DSP, two channels

- Over 100 events the **final cluster count** from the RTL implementation agrees with the offline Python/NumPy reference **event by event**
- The two cluster-counting instances account for **4.919 W of the estimated 12.37 W** on-chip total, and the CNN inference occupies **63.2 %** of the device DSPs - that DSP cost is what the accelerator work targets
- These are **RTL simulation and implementation results** on the ZU15EG, **not yet detector measurements**

» Established

- A modular chain: **10-channel unit -> 40-channel crate -> 120 channels** of readout capacity, on one clock and trigger
- Noise RMS **0.470 mV** with a **5.57 μ V** channel-to-channel sigma; **550 MHz** full-chain bandwidth; **0.82 %** gain dispersion
- Testbeam run this summer verified the system's stability

» Challenges

- The **dN/dx analysis is ongoing**: peak-finding efficiency, the wire-to-electronics interface, noise and shielding are still under study
- The 10 GbE link is estimated at **~30 kHz** against **11.5 kHz** measured; the ceiling is understood to sit in the **IP core**
- The wideband front end is sensitive to **ambient HF interference** and shielding was incomplete early in the run

» Next steps

- Improve the chamber-electronics interface and shielding, and **commission the full 120-channel readout**
- Optimise the **network IP core** and take online cluster counting **towards detector data**
- **Hardware-accelerated algorithms** and **scaled-down models** to cut power and resources, towards an **ASIC**

Thank you for your attention

Hongbin Liu - hbliu@ihep.ac.cn

Institute of High Energy Physics, Chinese Academy of Sciences



INSTR 2026

October 12-16, 2026 | Hefei, China