IDEC Chip Design Contest

A 10-Gb/s Memory Controller PHY for DDR5 RDIMM Interfaces with On-Die Loopback BIST and I3C-Configurable Mixed-Signal Control Words

Saransh Rajjarwal, YoungJun Byun, and Gyungsu Byun
Department of Electrical and Computer Engineering, Inha University

Introduction

1. Problem. AI-class servers push DDR5 RDIMM to 8800 MT/s, but the host MC PHY must drive a Gen-6 RCD that fans out to 40 DRAM packages on a 2-rank module — within a 2.5 pJ/b socket envelope and a millisecond-class warm-reset budget that today's I²C-firmware-driven training (12–40 ms) cannot meet.

2. Prior-art gap. SST 2-tap pre-emphasis cannot suppress the RCD fan-out's reflective postcursor; pulse-based FFEs need a DRAM-side damping resistor unavailable on JEDEC RDIMM parts; receiver-DFE-only PHYs leave C/A unequalized and cost >3 pJ/b; open-loop ZQ drifts ±8 % over PVT.

3. This work. A 28-nm MC PHY that closes all four gaps in one die: (i) C/A slice-based CTLE + adaptive 3-tap DFE, (ii) DQ 2-tap SST FFE + 1-tap embedded-DFE slicer, (iii) **on-die BIST that exploits the RCD's loopback / pass-through self-test mode** to complete 7-stage cold-start training in **14 µs (850× speed-up)**, and (iv) **I3C MC_CW sideband at 12.5 MHz** (I²C-FM+ backward-compatible) mirroring the JEDEC RCD register map so one firmware driver covers MC + RCD.

Body

A. System architecture (Fig. 1). The PHY splits into two deliberately asymmetric datapaths matched to their channels. The **C/A path** drives the Gen-6 RCD that fans out to 40 DRAM packages — a reflective, dispersive channel — so it carries the heavier 3-slice CTLE + adaptive 3-tap DFE and is the path closed by RCD loopback during training. The **DQ path** is point-to-point with forwarded DQS (no receiver CDR), a cleaner channel, so it needs only a 2-tap SST FFE + embedded-DFE slicer. A central **LTFSM** is the integration anchor: it embeds the BIST engine, sequences the 7-stage cold-start and thermal-retune, arbitrates MC_CW access between the I3C sideband (firmware) and BIST (internal), and keeps the link online during retune. Timing is set by a 5-GHz ADPLL feeding per-lane 0.78-ps DCDLs. 28-nm CMOS, 80 DQ + 56 C/A lanes, 20.9 mm².

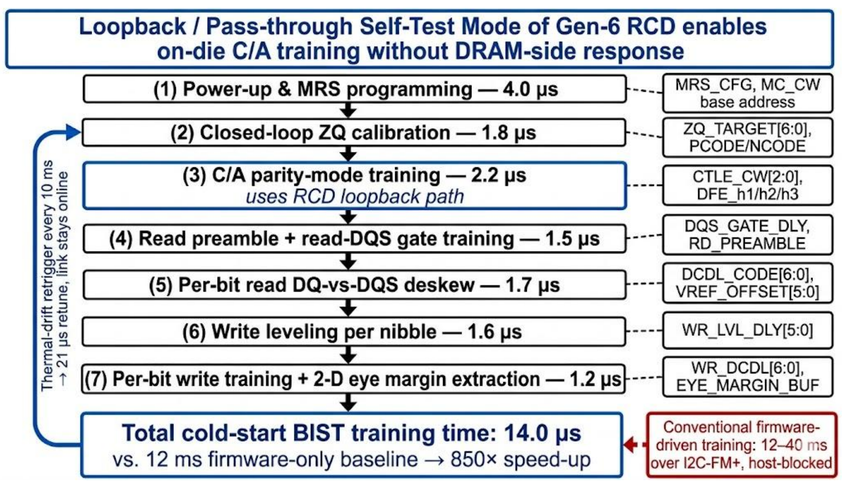

(Fig. 2).

B. Loopback-driven on-die BIST

•**The Gen-6 RCD's pass-through / loopback self-test mode lets the BIST close the C/A training loop at the DIMM connector**, so the MC can characterize and tune the C/A channel without depending on DRAM-side response paths — eliminating the per-platform firmware bring-up that conventional DDR5 RCDs require. 4.6-kgate FSM + 2-kbit pattern memory; reuses each lane's DCDL → only **5.3 % lane-area overhead**. Cold-start training: **14.0 µs** vs. 12 ms firmware baseline → **850×**. Thermal retune: **21.0 µs** with link kept online; verified zero errors over 10^{11} symbols during the retune window.

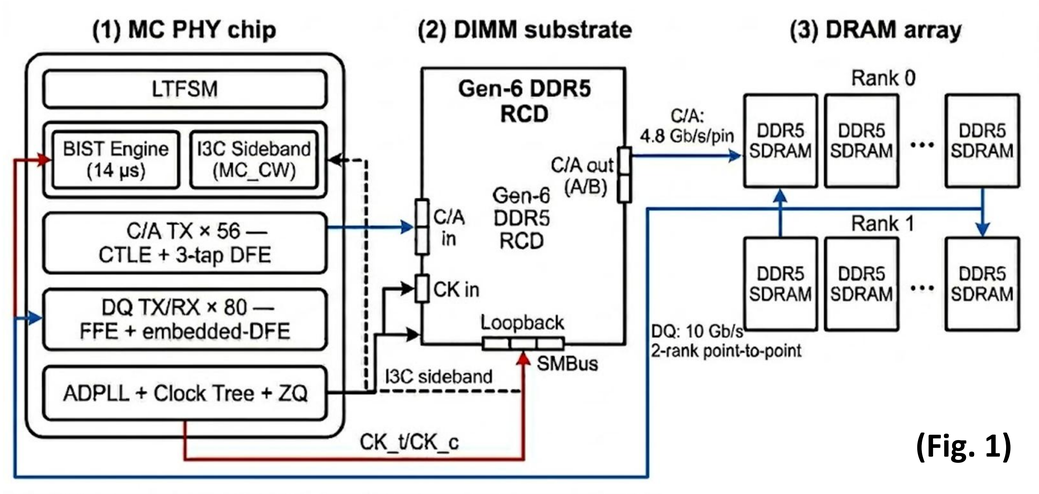

(Fig. 1)

C. I3C MC_CW sideband

256-byte mixed-signal control-word file: driver impedance (0x00–0x1F), EQ coefficients FFE/DFE/CTLE (0x20–0x5F), V_REF + DCDL init (0x60–0x9F), BIST control + 2-D eye buffer (0xA0–0xCF), ZQ targets + thermal thresholds + vendor scratchpad (0xD0–0xFF).

•**Address map mirrors JEDEC RCD register-set** → one platform-firmware driver, one DMA descriptor, one IBI handler covers both MC and RCD.

•I3C SDR @ 12.5 MHz; falls back to I²C-FM+ @ 1 MHz for legacy boards.

•Three transaction types: *private write* with end-of-transaction commit strobe (kills partial-update glitches that single-byte I²C writes cause on multi-byte coefficient fields); *private read* (BIST status, eye buffer, thermal sensor); *in-band interrupt* (IBI) raised on thermal / margin breach.

•Hardware OTP write-lock on driver-impedance and ZQ-target bytes prevents firmware faults from driving the output stage out of the linear region.

Die Photo

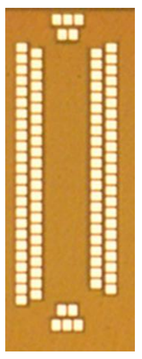

Eye diagram

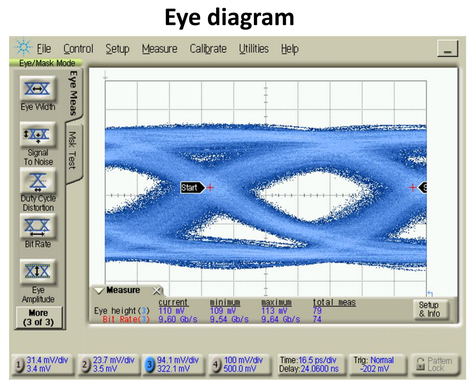

TABLE I — Measured Performance Summary

Parameter	This Work
Process	28-nm CMOS
Application	DDR5 RDIMM (Gen-6 RCD)
Data rate per pin	**10 Gb/s (DQ) / 4.8 Gb/s (C/A)**
Channel	2-rank RDIMM via Gen-6 RCD
Equalization	C/A: CTLE + 3-tap DFE | DQ: 2-tap FFE + 1-tap embedded DFE
DQ timing margin (R / W)	**0.61 / 0.66 UI**
DQ voltage margin (R / W)	**235 / 248 mV**
Energy efficiency	**1.62 pJ/b**
ZQ drift over PVT	±1.6 % (5× better than open-loop baseline)
BIST cold-start training	**14 µs** (vs. **12 ms** firmware baseline → **850× speed-up**)
Sideband	I3C 12.5 MHz + I2C-FM+ (256-byte MC_CW)

D. Equalization (compact) C/A: 3-slice CTLE programmable at +6 / +9 / +12 dB peaking at 4.8 GHz, in series with a 3-tap DFE sharing one SS-LMS engine across taps (−38 % loop area vs. per-tap LMS). DQ: 2-tap SST FFE on TX co-adapted with a stack-reduced strong-arm slicer that merges the 1-tap DFE summer into the input pair (−49 % slicer power vs. summer-then-slicer cascade). All coefficients pre-loaded from MC_CW, then refined by BIST in 1.2 µs/byte-lane.

Conclusion

1. A 28-nm 10-Gb/s/pin MC PHY is, to our knowledge, the first to combine host-side C/A + DQ adaptive EQ, **RCD-loopback-enabled on-die BIST**, closed-loop ZQ, and a **JEDEC-RCD-compatible I3C MC_CW register map** in one die — collapsing cold-start training to 14 µs (850×) and runtime thermal retune to 21 µs (67×) without taking the link offline.

2. Measured **0.61 / 0.66 UI** read/write margin and **235 / 248 mV** voltage margin on 2-rank DDR5 RDIMM at 10 Gb/s, all at **1.62 pJ/b** — 27 % below the 2.5 pJ/b AI-class socket envelope.