# A Pin- and Power-Efficient 20.83Gb/s/wire 0.94pJ/bit Forwarded Clock CNRZ-5 Coded SerDes up to 12mm for MCM Packages in 28nm CMOS

A. Shokrollahi[1], D. Carnelli[1], J. Fox[2], K. Hofstra[1], B. Holden[1], A. Hormati[1], P. Hunt[2], M. Johnston[1], S. Pesenti[1], R. Simpson[2], D. Stauffer[1], A. Stewart[2], G. Surace[2], A. Tajalli[1], O. Talebi Amiri[1], A. Tschank[2], R. Ulrich[1], Ch. Walter[1], F. Licciardello[1], Y. Mogentale[1], A. Singh[1]

[1]Kandou Bus, Lausanne, Switzerland,
[2]Kandou Bus, Northampton, United Kingdom,

# Outline

- Introduction and motivation
- Signaling
- Macro architecture
  - Common block
  - Tx
  - Rx
- System Implementation
- Results
- Conclusion

# Motivation

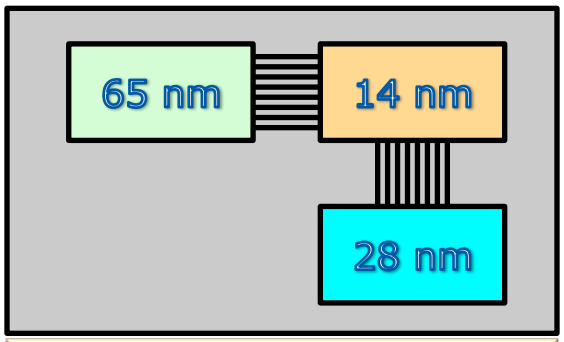

Heterogeneous die in package

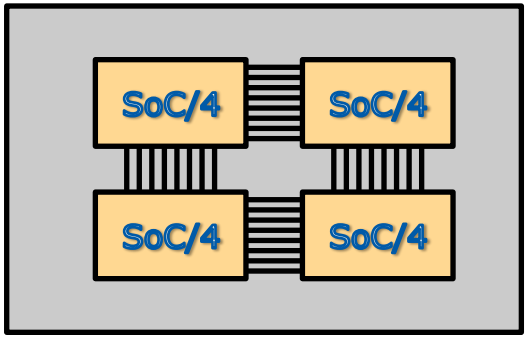

Lego-like SoC design

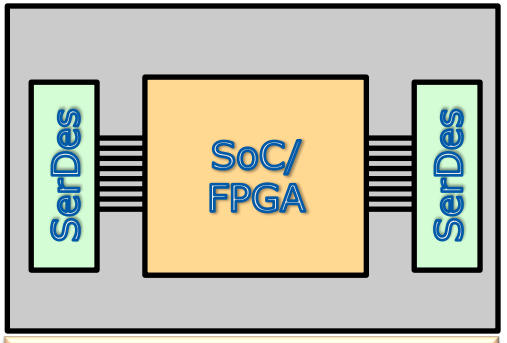

Interface to off-chip SerDes

- 2.5D integration is a means to connect multiple die in a low-cost package
  - Increase yield, lower manufacturing costs
- Efficient solution requires high-bandwidth, low-power SerDes

# Constraints of a Versatile 2.5D Solution

- Want very high throughput using few wires
  - Enable use of low cost of substrates
  - Use 20+ Gb/s per signal wire
  - Need very high bandwidth per mm die-edge
- Requires relatively large distance
  - Need up to 12 mm, scalability to longer distances desirable
- Demands very low power
  - Ideally same as or lower than on-chip links

# 2.5D Options

| Method | # wires | Distance | Power |
|---|---|---|---|
| Silicon interposers | Many | Very short | Good |
| Wafer level Processing | Many | Very short | Good |
| SerDes | Fewer? | Longer? | ? |

- Need SerDes with
  - Very high throughput per wire
  - Very low power at desired channel lengths

# SerDes: Prior Work

| Reference | [1] | [2] | [3] | This work |
|---|---|---|---|---|
| Year | 2009 | 2012 | 2014 | 2016 |
| pJ/bit | 1.9 | 0.54 | 1.4 | 0.94 |
| BW/pin (Gb/s) | 8.9 | 20 | 6 | 20.83 |
| Technology | 45nm | 28nm | 32nm | 28nm |
| Signaling | D | GRSE | D | CNRZ-5 |
| Channel loss | ≤ 20 dB | ≤ 1 dB | ≤ 3 dB | ≤ 3 dB |
| Substrate | SiC | MCM | Meg6 | MCM |
| Reach | ≤ 40 mm | ≤ 4.5 mm | ≤ 19 mm | ≤ 12 mm |

SiCa=Silicon carrier, SE = single ended, D = differential, GRSE=ground referenced single-ended, MCM = MCM organic substrate, Meg 6 = Megtron 6

# Outline

- Introduction and motivation
- Signaling
- Macro architecture
  - Common block
  - Tx
  - Rx
- System Implementation
- Results
- Conclusion

# Signaling, Implementation

- Choice of signaling scheme is a major ingredient for realizing performance targets, especially power
  - Signaling needs to have high pin-efficiency and built-in robustness to noise (common mode, SSO, ISI, EMI, etc.)

# Chord Signaling

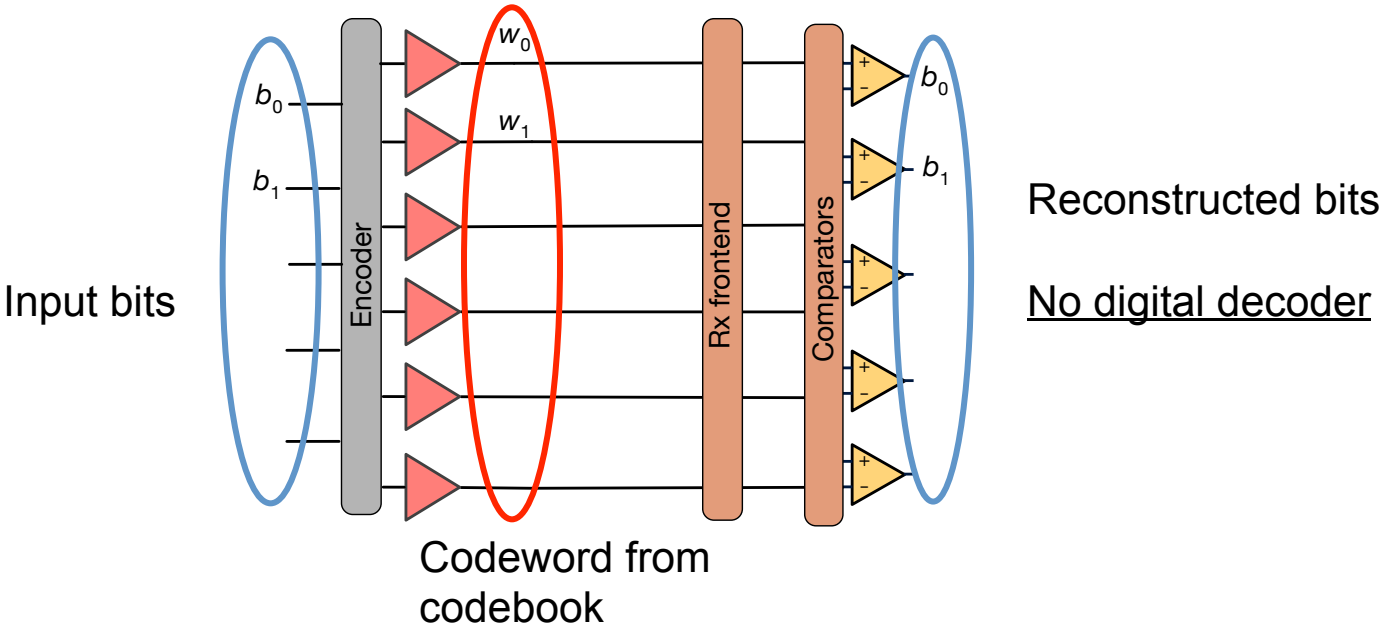


- Correlated signals on wires, matched comparator network
  - Signals on wires belong to a codebook
- Chordal code = codebook + comparators

# Chord Signaling

- Correct design of code
  - Increases throughput per pin
  - Reduces ISI
  - Eliminates SSO noise
  - Eliminates common mode noise
  - Reduces EMI noise

# CNRZ-5

- Transmits 5 bits on a collection of 6 wires in every UI.
- Codewords are judiciously chosen permutations of [+1,-1,+1/3,+1/3,-1/3,-1/3].
- Five comparators:
  - Two compare one wire value against another,
  - Two compare average of two wire values against a third,
  - One compares average of three wire values against the other three.

# CNRZ-5

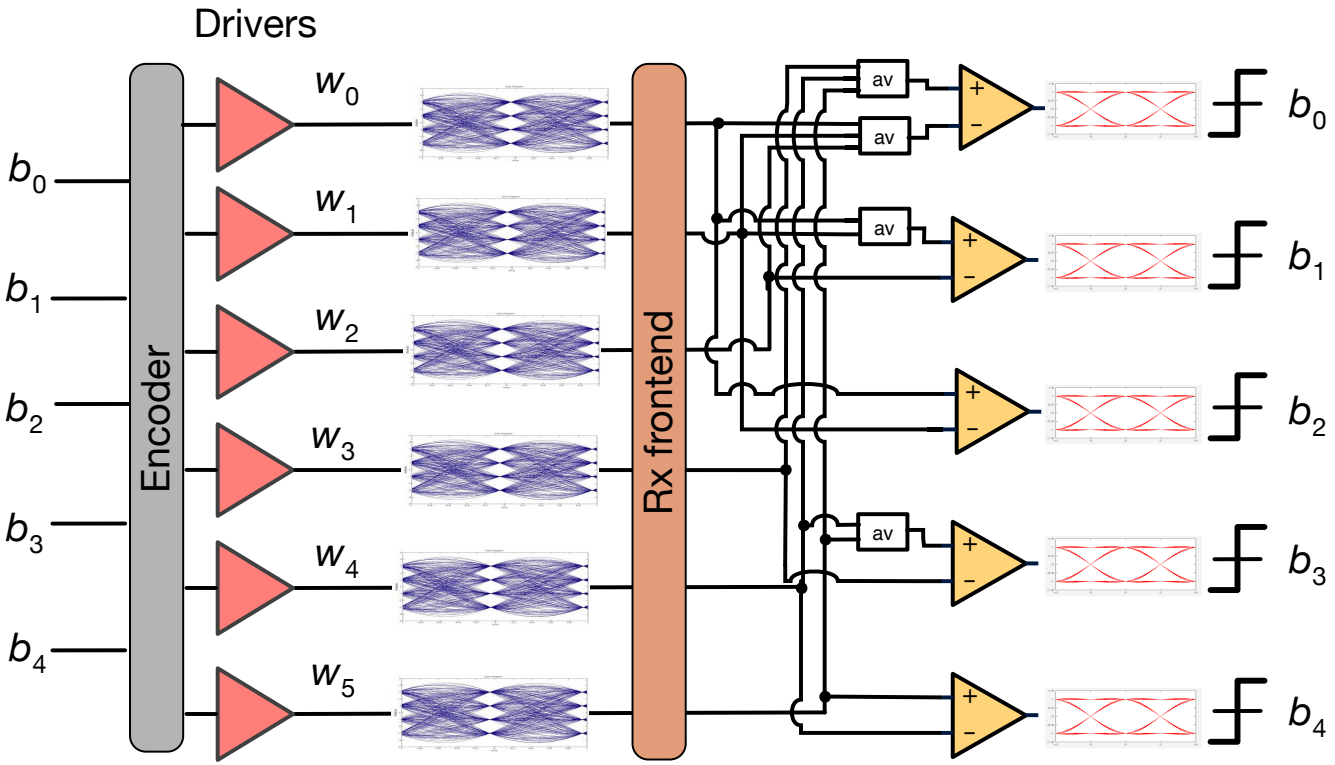

# CNRZ-5

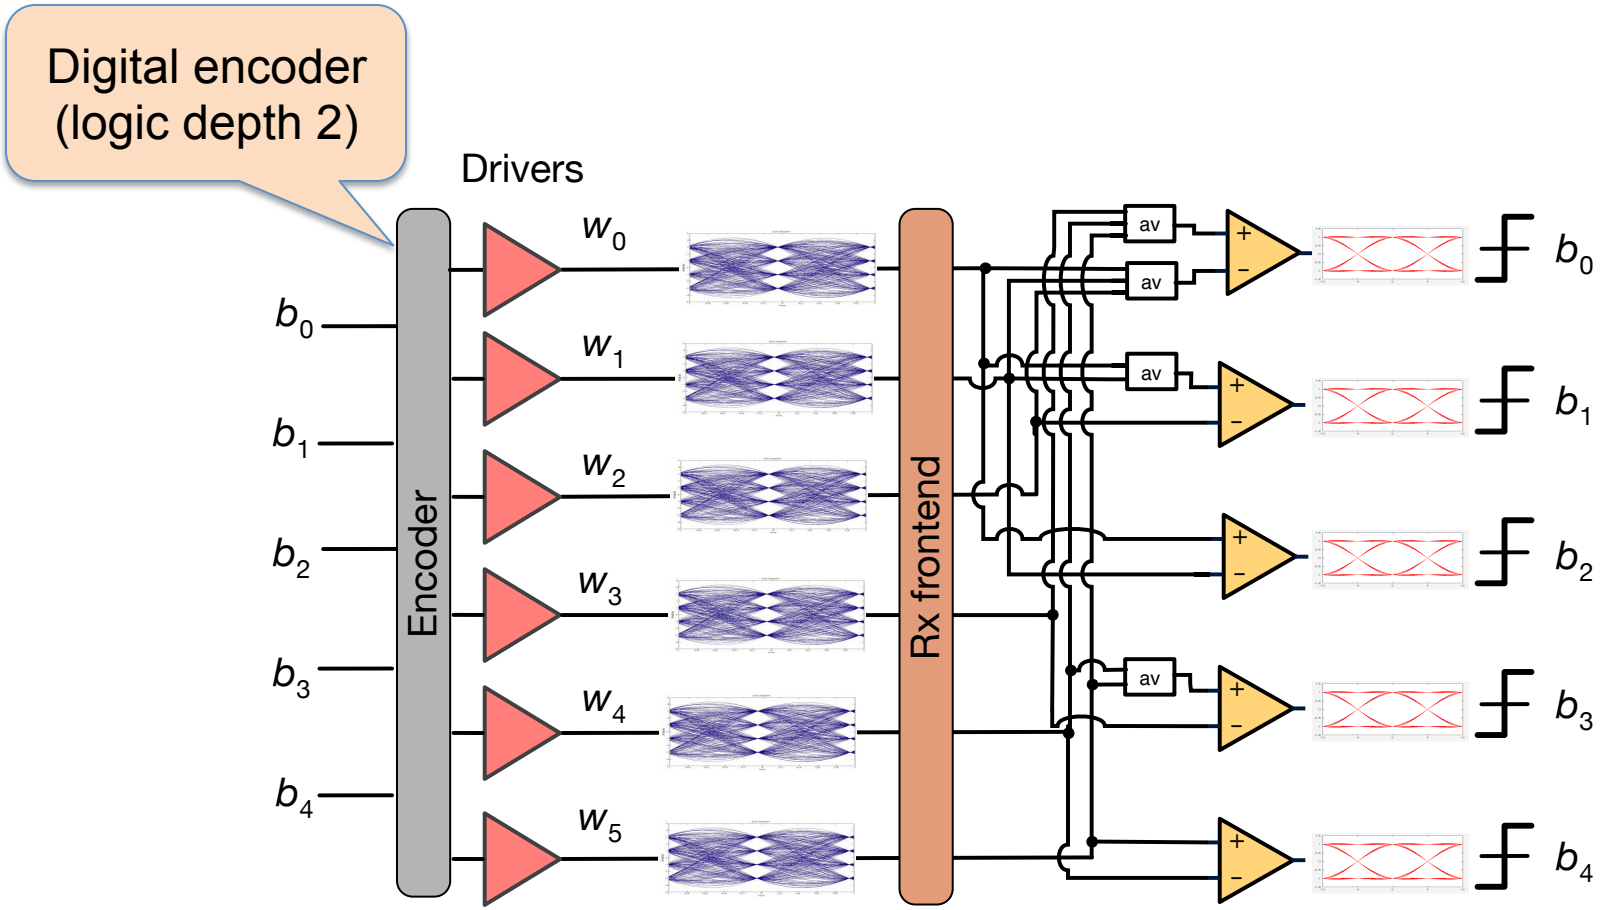

# CNRZ-5

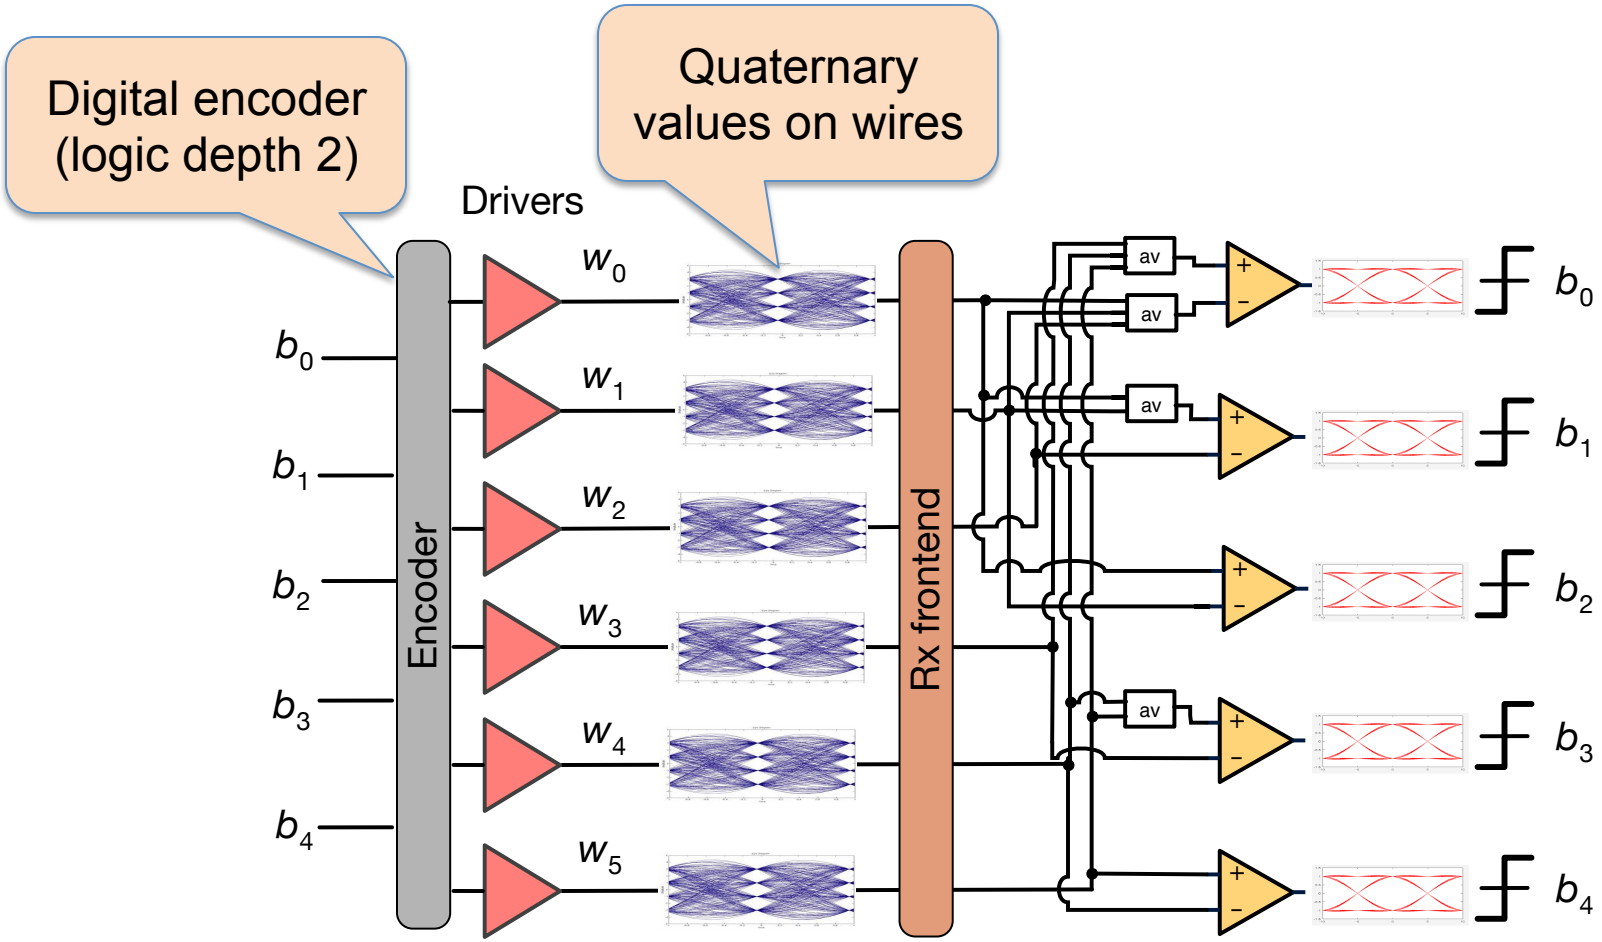

# CNRZ-5

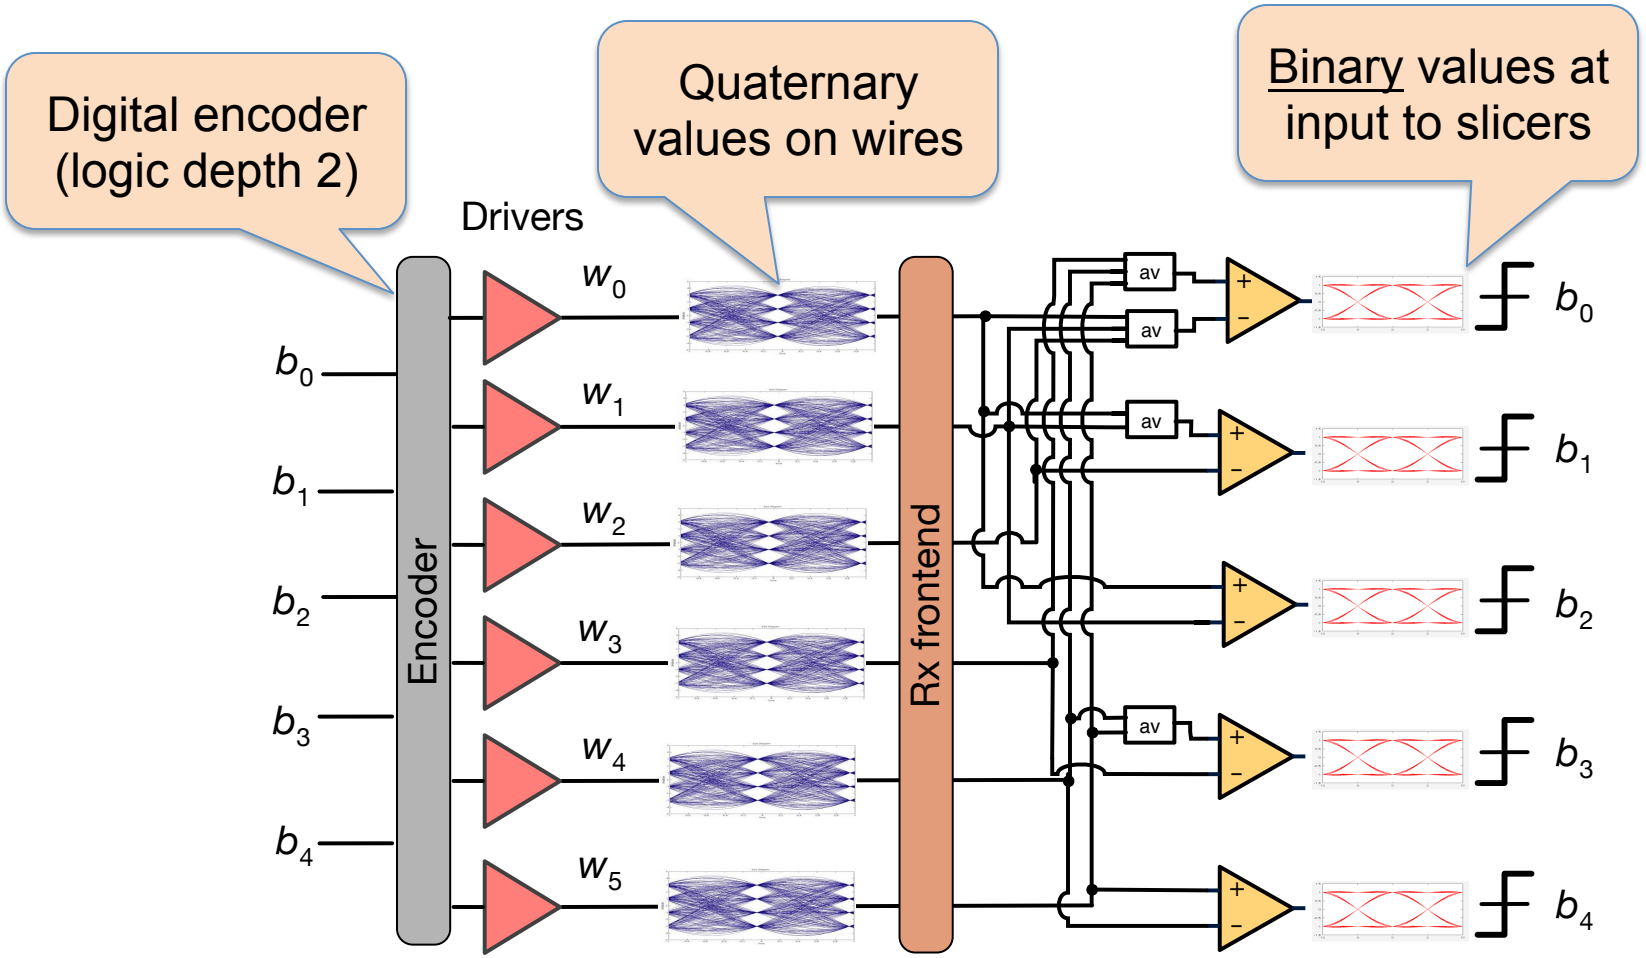

# CNRZ-5 Codebook

| | |
|---|---|
| [1/3, -1/3, -1, -1/3, 1/3, 1] | [1/3, -1/3, -1, -1/3, 1, 1/3] |
| [1, 1/3, -1/3, -1, -1/3, 1/3] | [1, 1/3, -1/3, -1, 1/3, -1/3] |
| [-1/3, -1, 1/3, -1/3, 1/3, 1] | [-1/3, -1, 1/3, -1/3, 1, 1/3] |
| [1/3, -1/3, 1, -1, -1/3, 1/3] | [1/3, -1/3, 1, -1, 1/3, -1/3] |
| [-1/3, 1/3, -1, -1/3, 1/3, 1] | [-1/3, 1/3, -1, -1/3, 1, 1/3] |
| [1/3, 1, -1/3, -1, -1/3, 1/3] | [1/3, 1, -1/3, -1, 1/3, -1/3] |
| [-1, -1/3, 1/3, -1/3, 1/3, 1] | [-1, -1/3, 1/3, -1/3, 1, 1/3] |
| [-1/3, 1/3, 1, -1, -1/3, 1/3] | [-1/3, 1/3, 1, -1, 1/3, -1/3] |
| [1/3, -1/3, -1, 1, -1/3, 1/3] | [1/3, -1/3, -1, 1, 1/3, -1/3] |
| [1, 1/3, -1/3, 1/3, -1, -1/3] | [1, 1/3, -1/3, 1/3, -1/3, -1] |
| [-1/3, -1, 1/3, 1, -1/3, 1/3] | [-1/3, -1, 1/3, 1, 1/3, -1/3] |
| [1/3, -1/3, 1, 1/3, -1, -1/3] | [1/3, -1/3, 1, 1/3, -1/3, -1] |
| [-1/3, 1/3, -1, 1, -1/3, 1/3] | [-1/3, 1/3, -1, 1, 1/3, -1/3] |
| [1/3, 1, -1/3, 1/3, -1, -1/3] | [1/3, 1, -1/3, 1/3, -1/3, -1] |
| [-1, -1/3, 1/3, 1, -1/3, 1/3] | [-1, -1/3, 1/3, 1, 1/3, -1/3] |
| [-1/3, 1/3, 1, 1/3, -1, -1/3] | [-1/3, 1/3, 1, 1/3, -1/3, -1] |

# CNRZ-5 Encoder

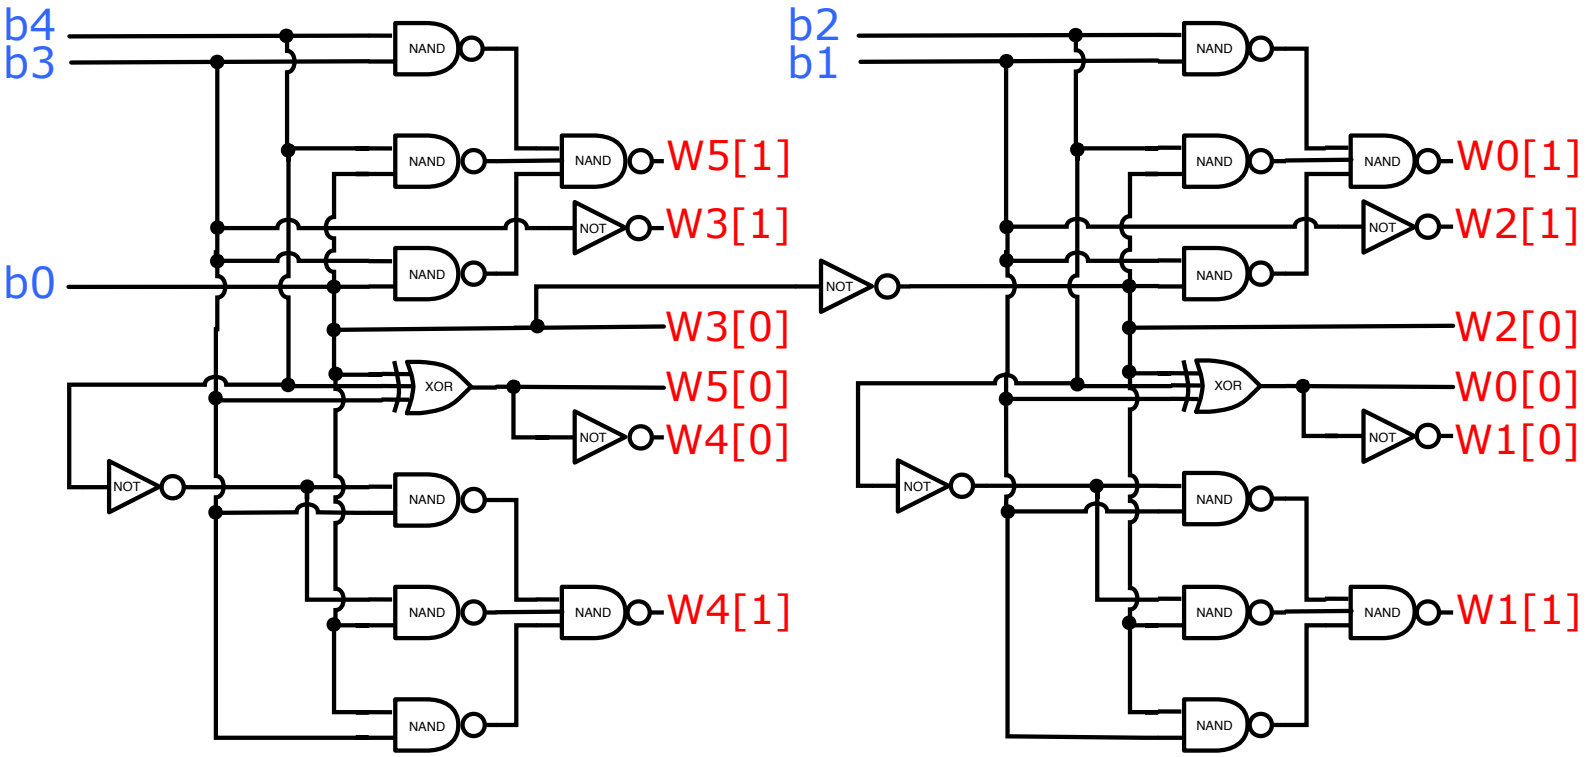


- Computes two bits Wi[0], Wi[1], i=0..5 from input bits b0,..,b4, which are used by Tx output driver to create codeword values w0,..,w5 on wires

# Binary Values at Input to Slicers

- Core concept is that of *ISI-ratio* [4], [5].
- This property massively reduces ISI noise

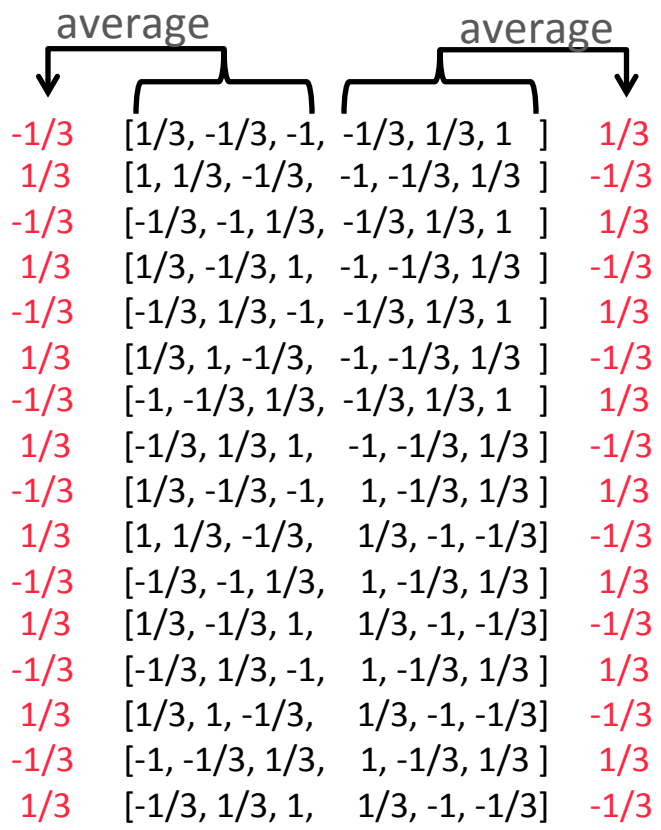


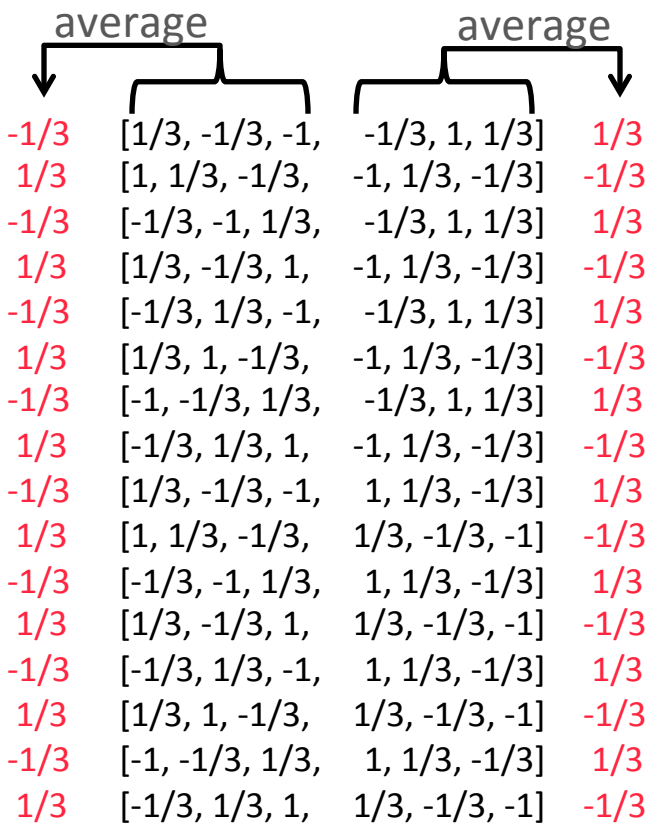

# Signal Integrity through Coding

| CM Noise | EMI | Crosstalk |
|---|---|---|
| • Comparators designed to tolerate common mode noise<br>• Balanced values across wires | Codewords designed to minimize EM-field | • Code designed to tolerate some crosstalk<br>• Complemented by sensible design |
| **SSO Noise** | **ISI** | **Reference-less** |
| Code designed so that all codewords draw same current from source | Values at input to slicers are binary, even though wire values are quaternary | • Comparators are self-referencing and binary<br>• No need for creating levels |

# Outline

- Introduction and motivation
- Signaling
- Macro architecture
  - Common block
  - Tx
  - Rx
- System Implementation
- Results
- Conclusion

# Macro Architecture

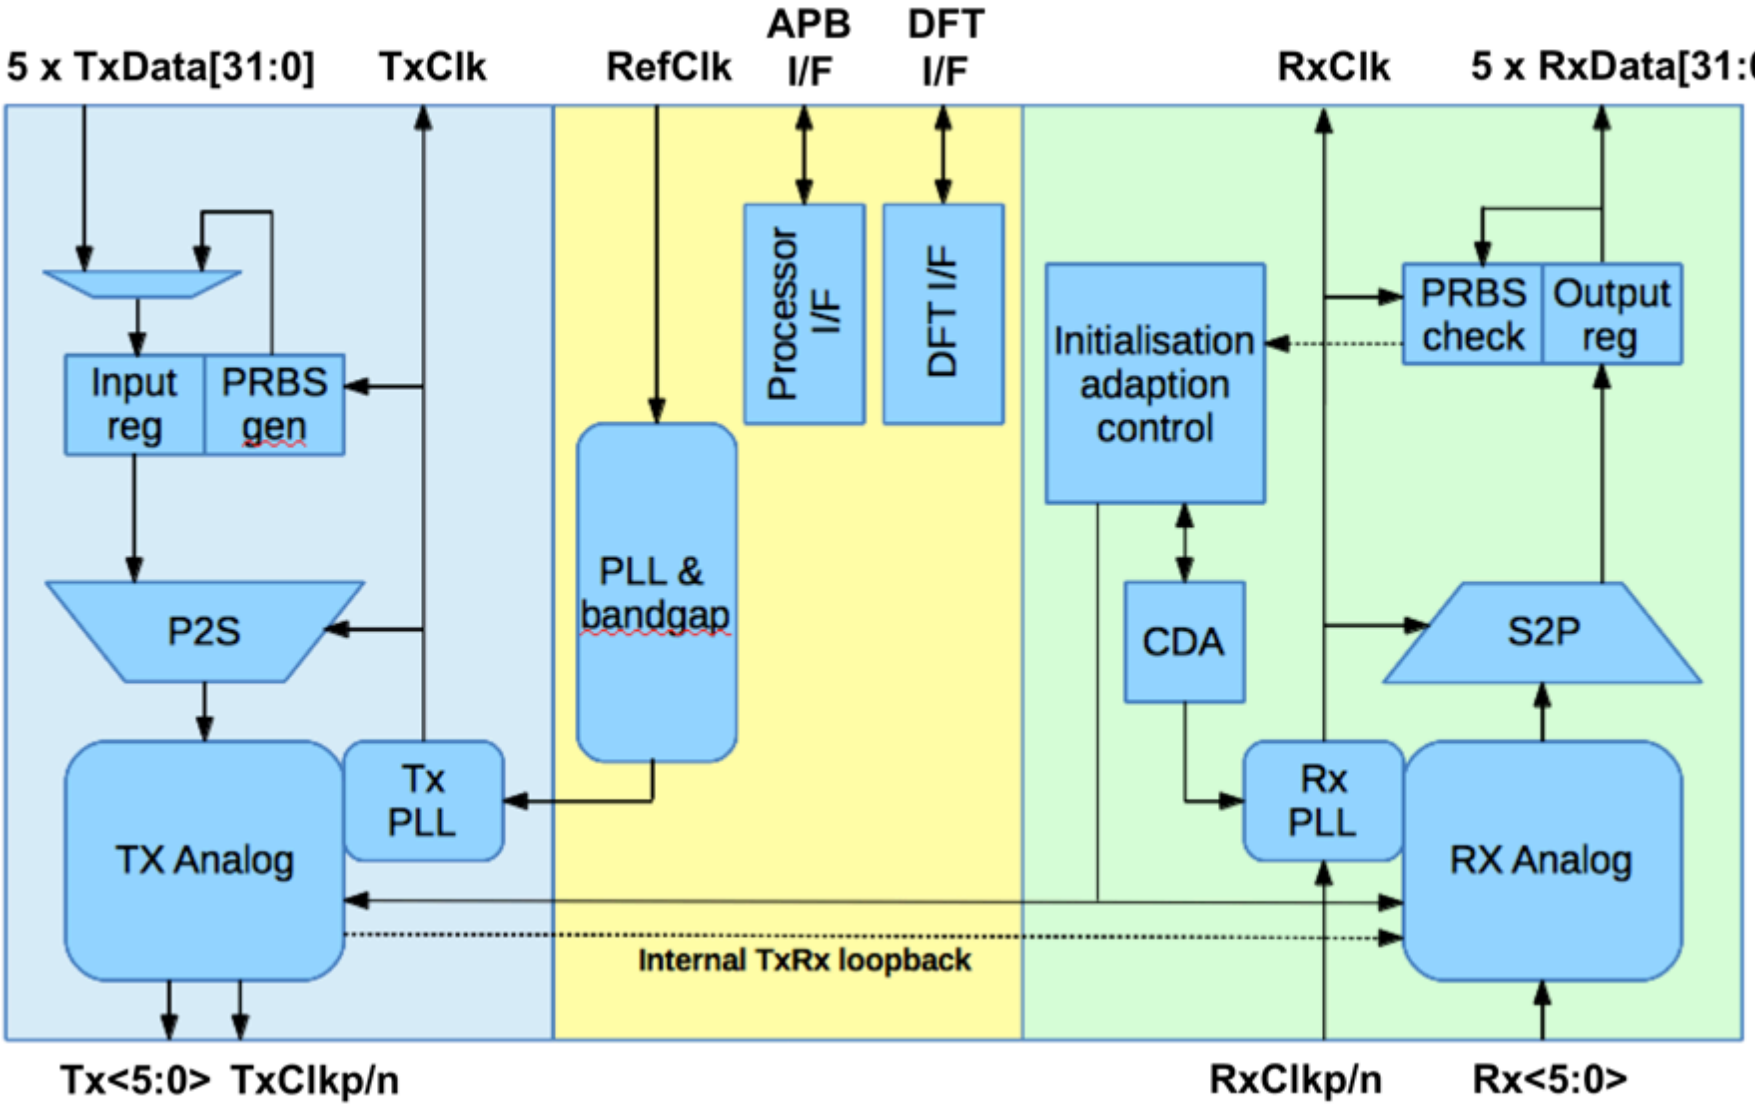

# Macro Architecture

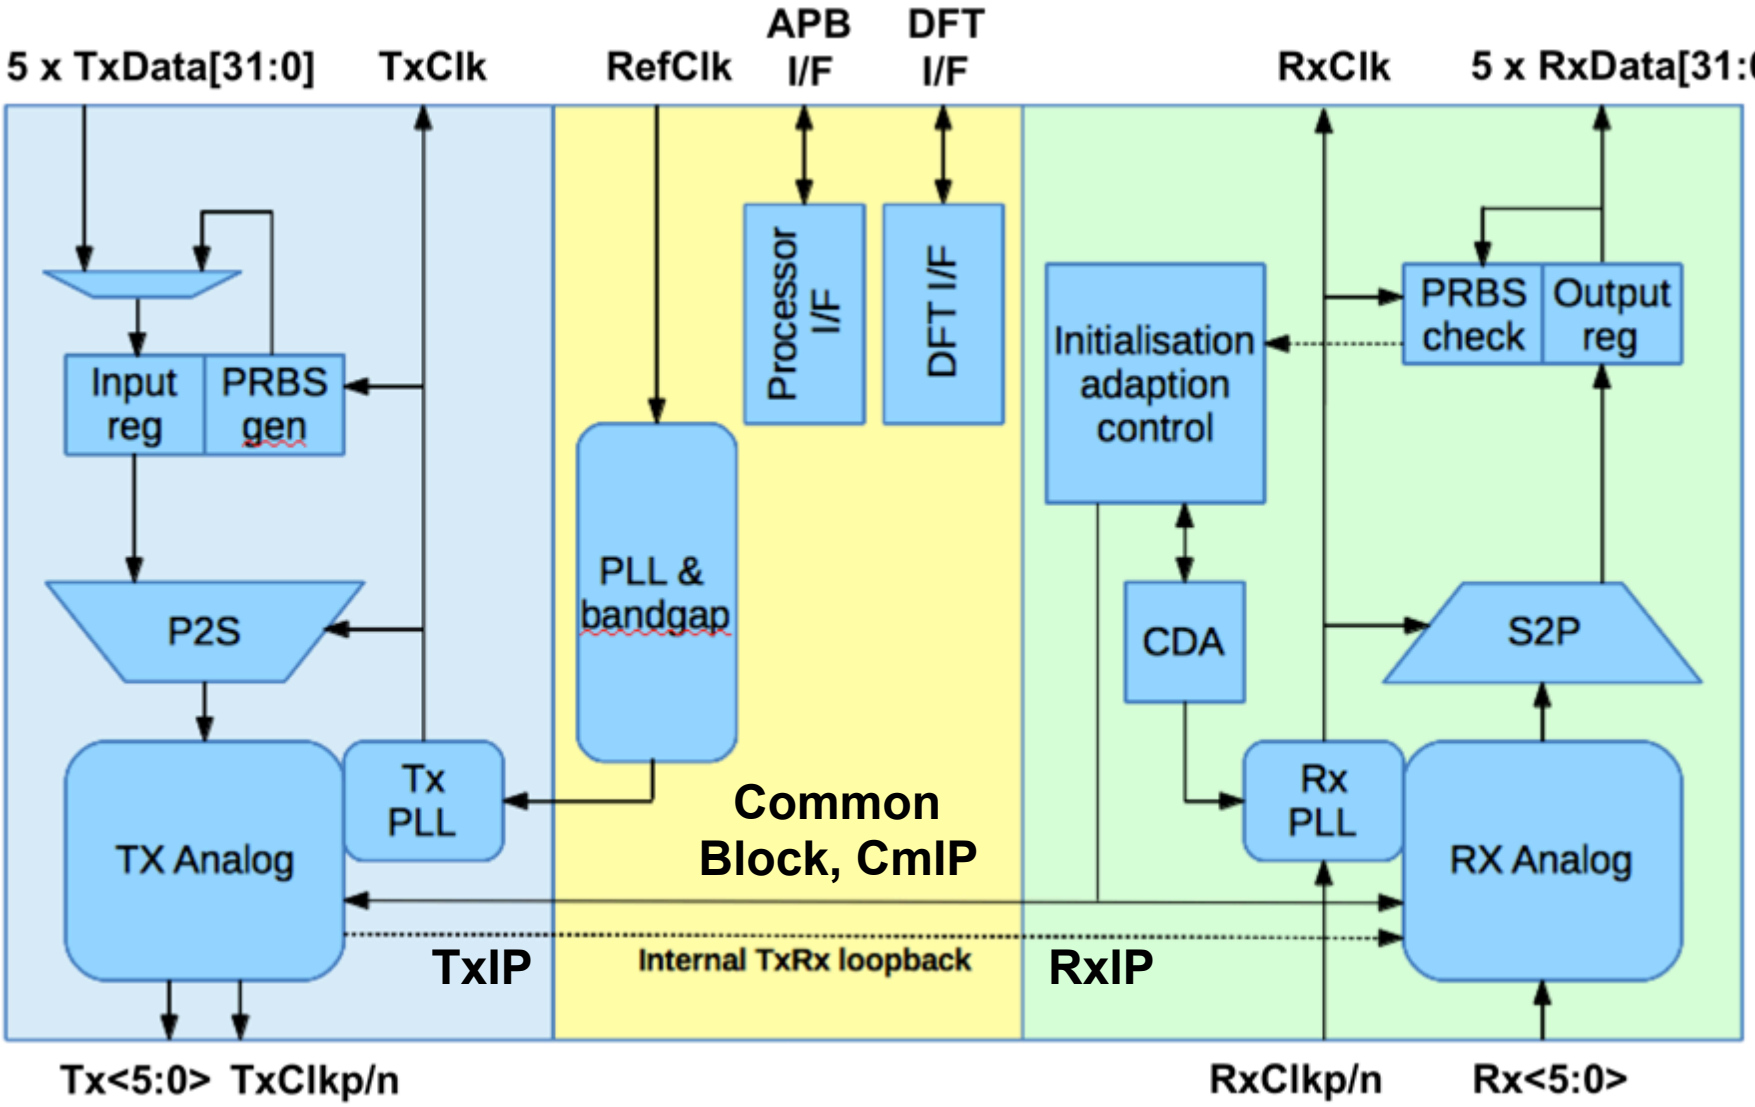

# Macro Architecture

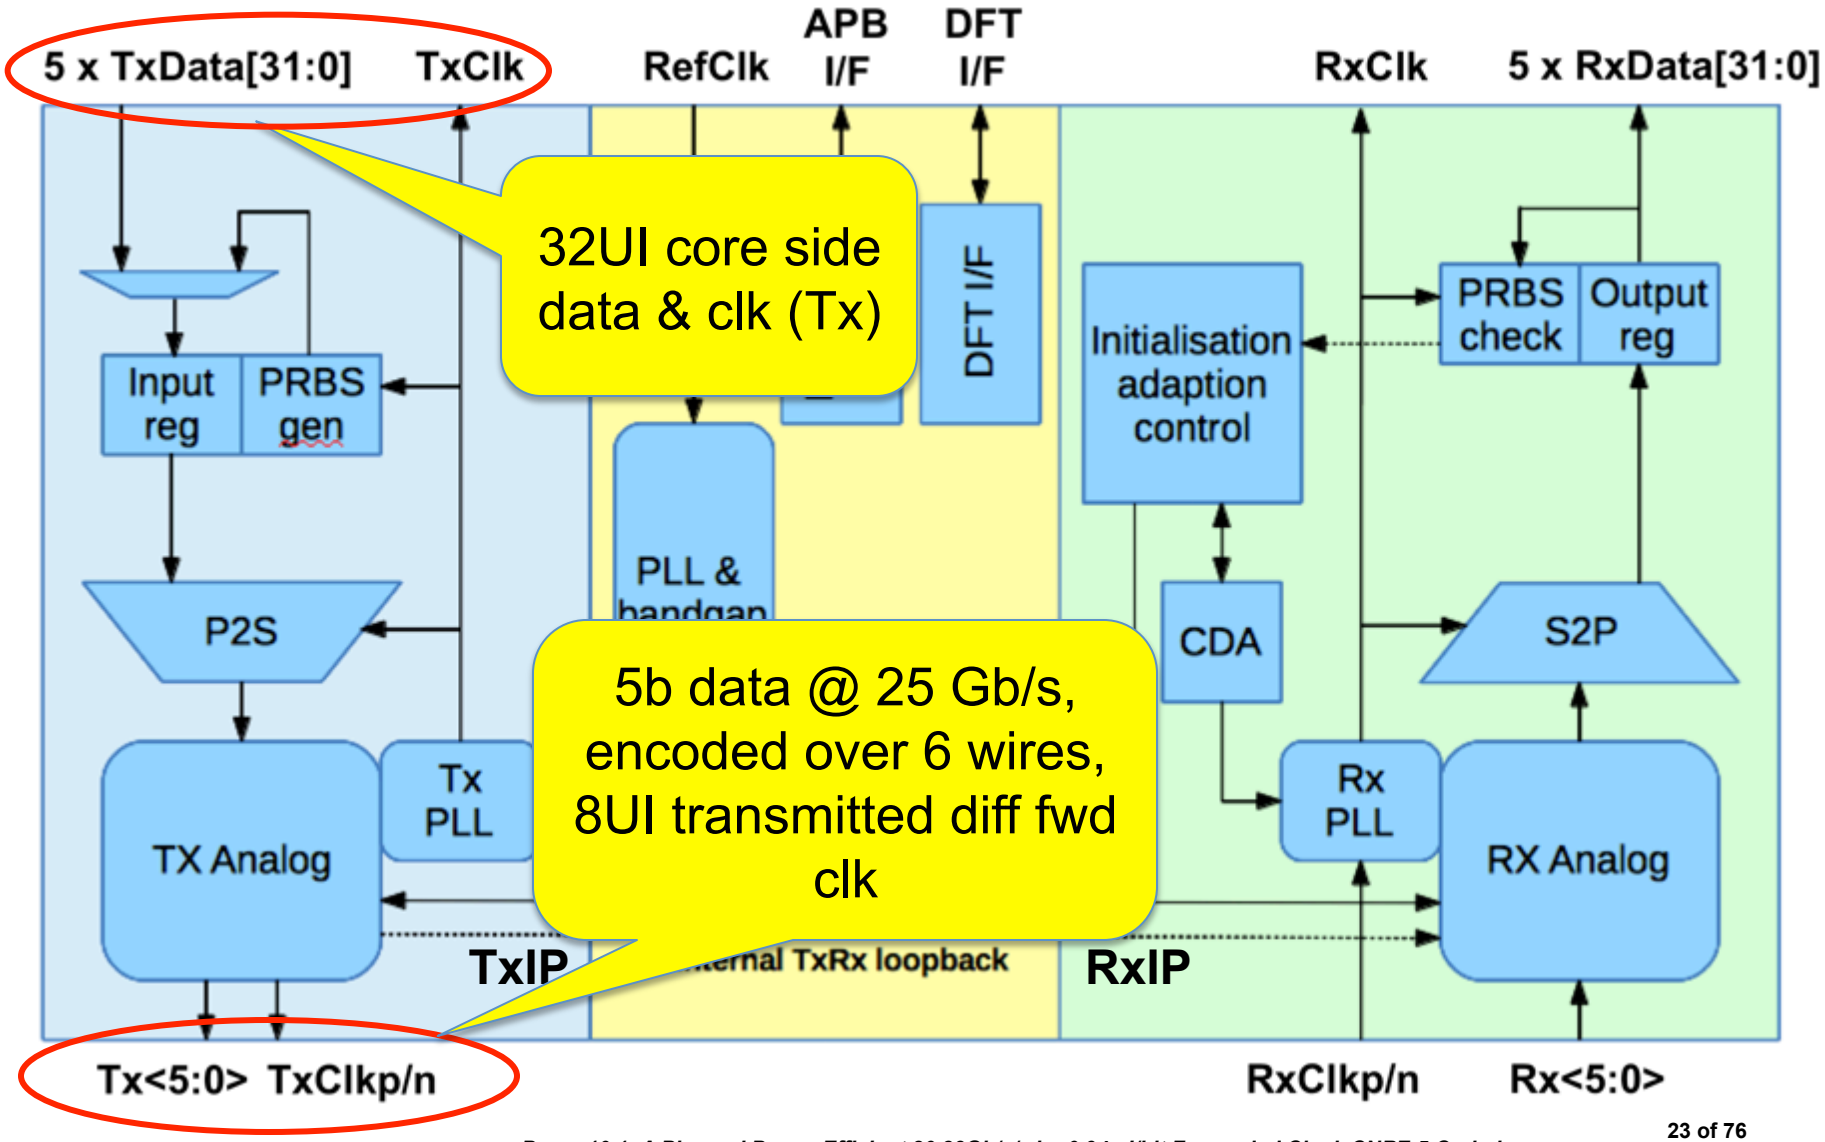

# Macro Architecture

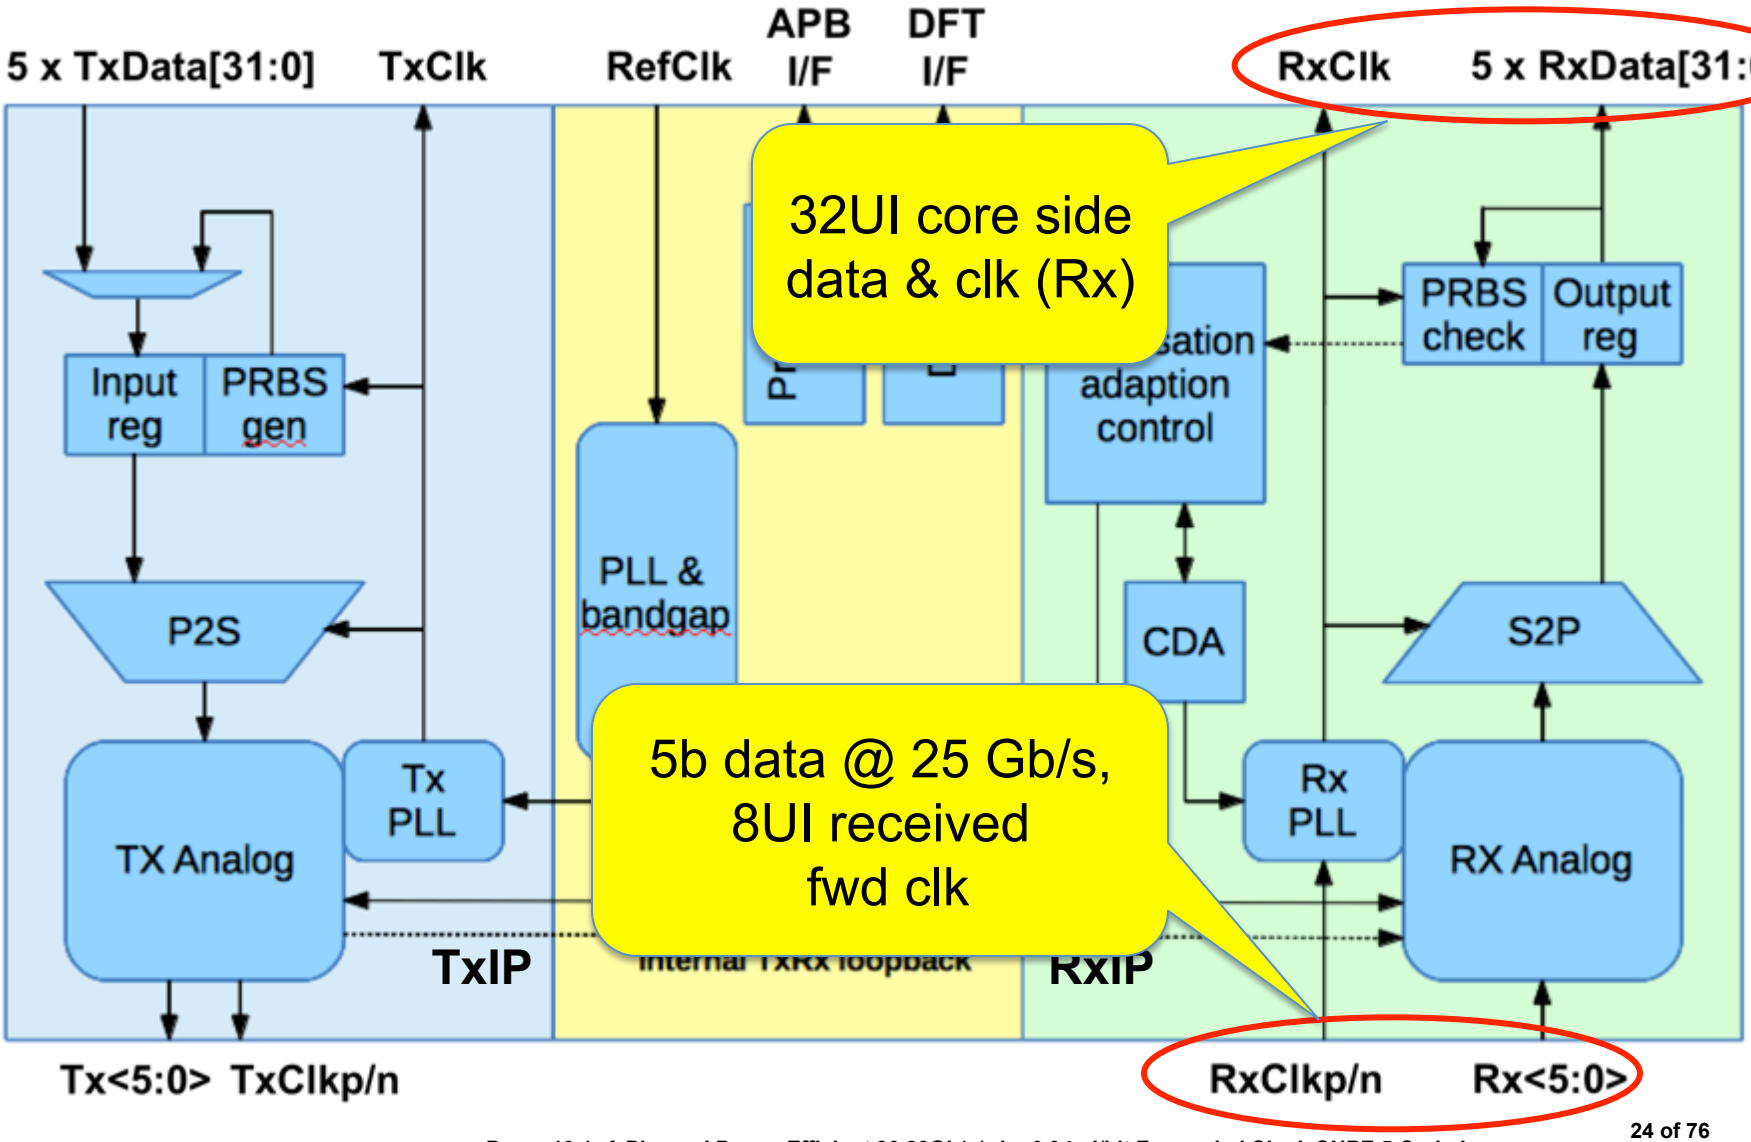

# Common Block

- CmIP consists of:
  - Main PLL
    - Ring Oscillator, produces 2 phase 8UI clock (3.125GHz)
    - Programmable to cover half rate to full rate links (25Gb/s to 12Gb/s)
  - Bandgap
  - Temperature sensor

# Transmitter

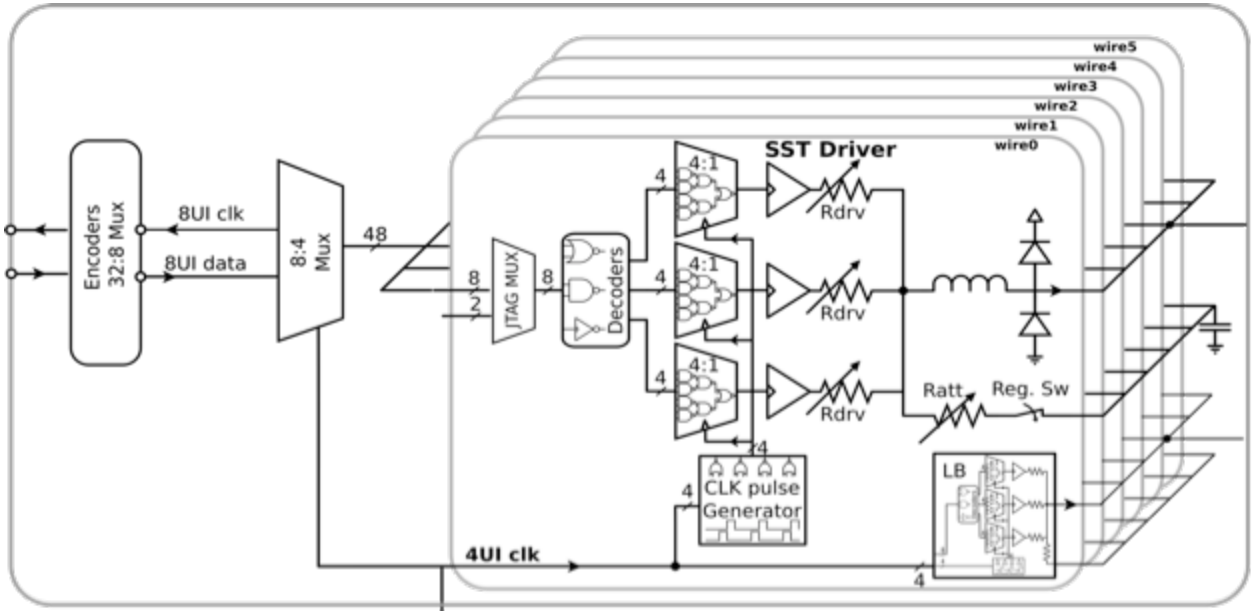

# Transmitter

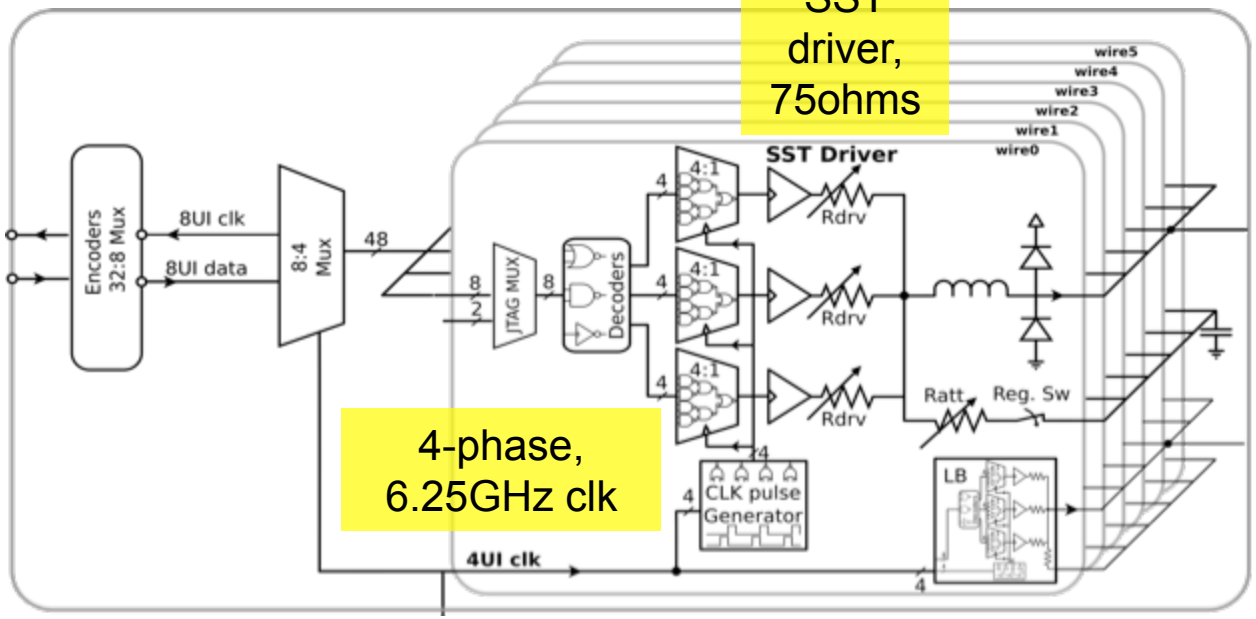

# Transmitter

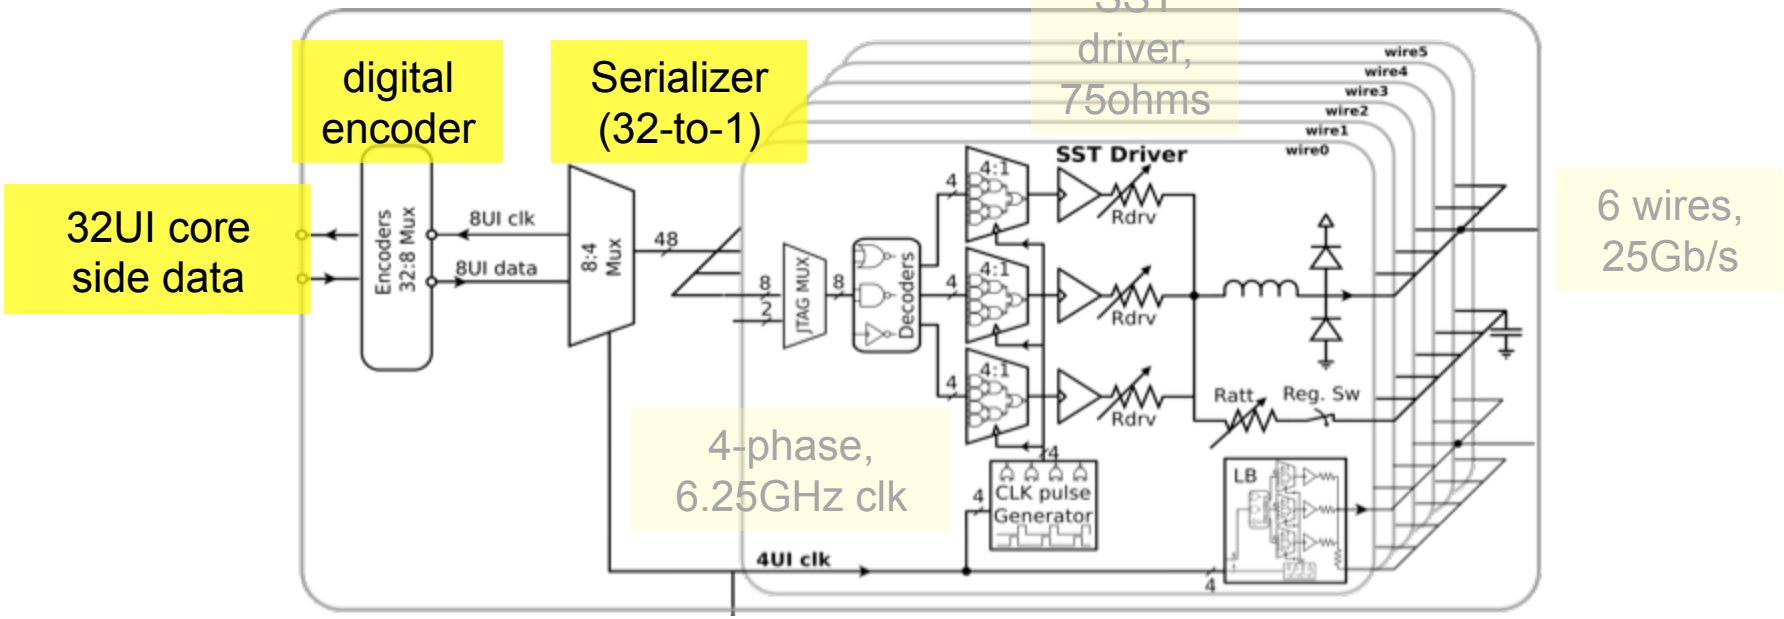

# Transmitter

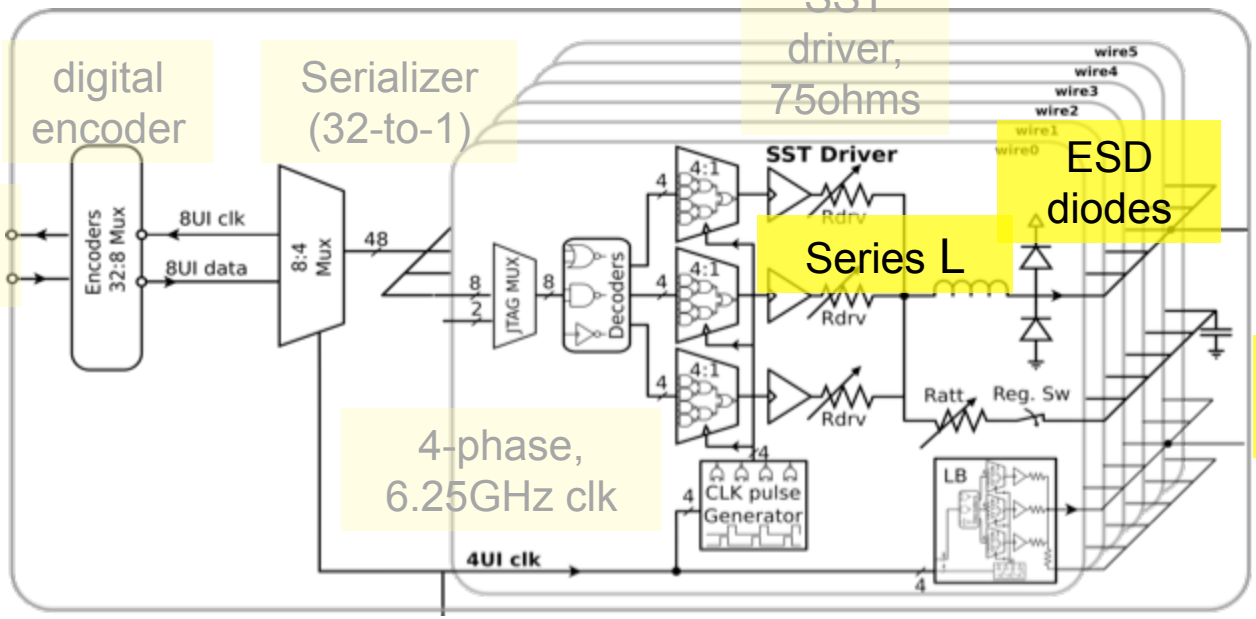

# Transmitter

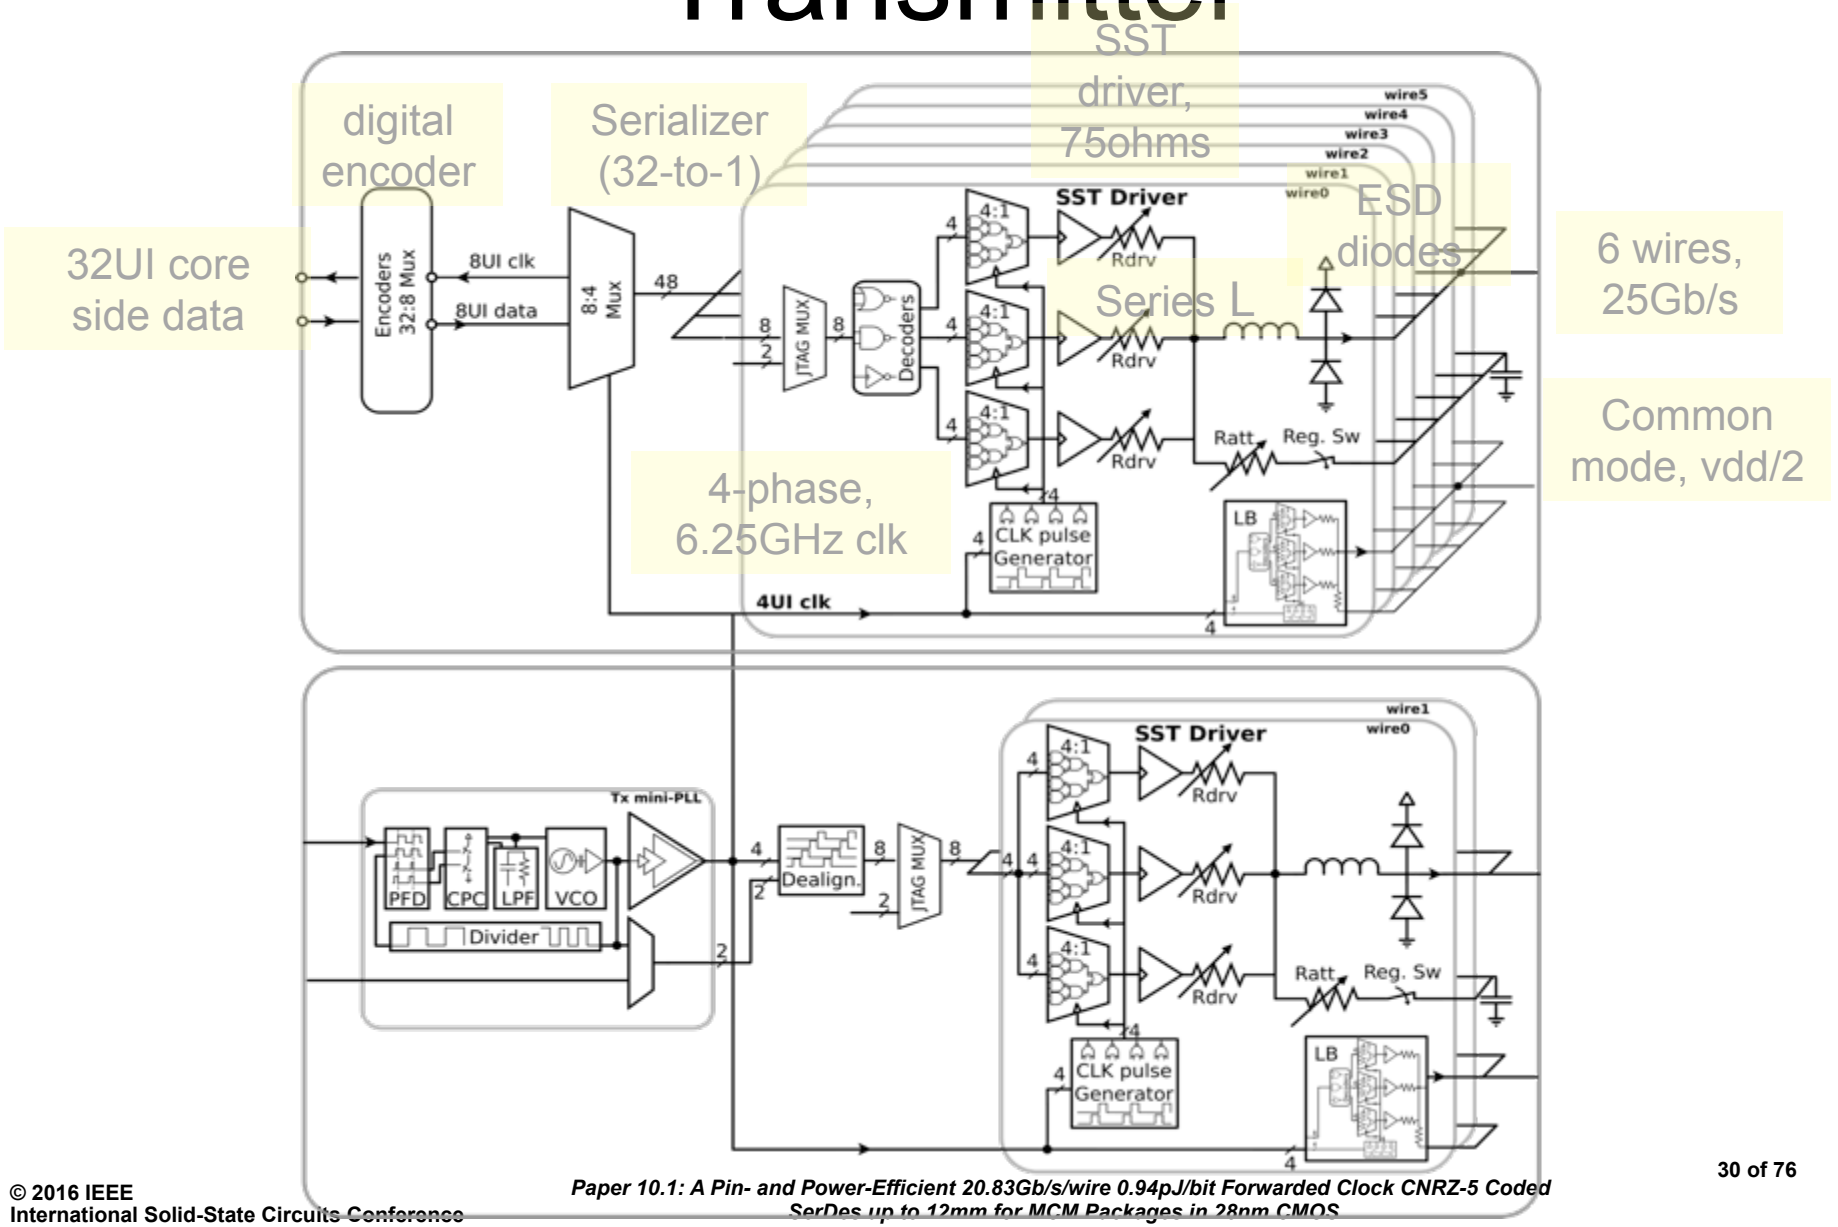

# Transmitter

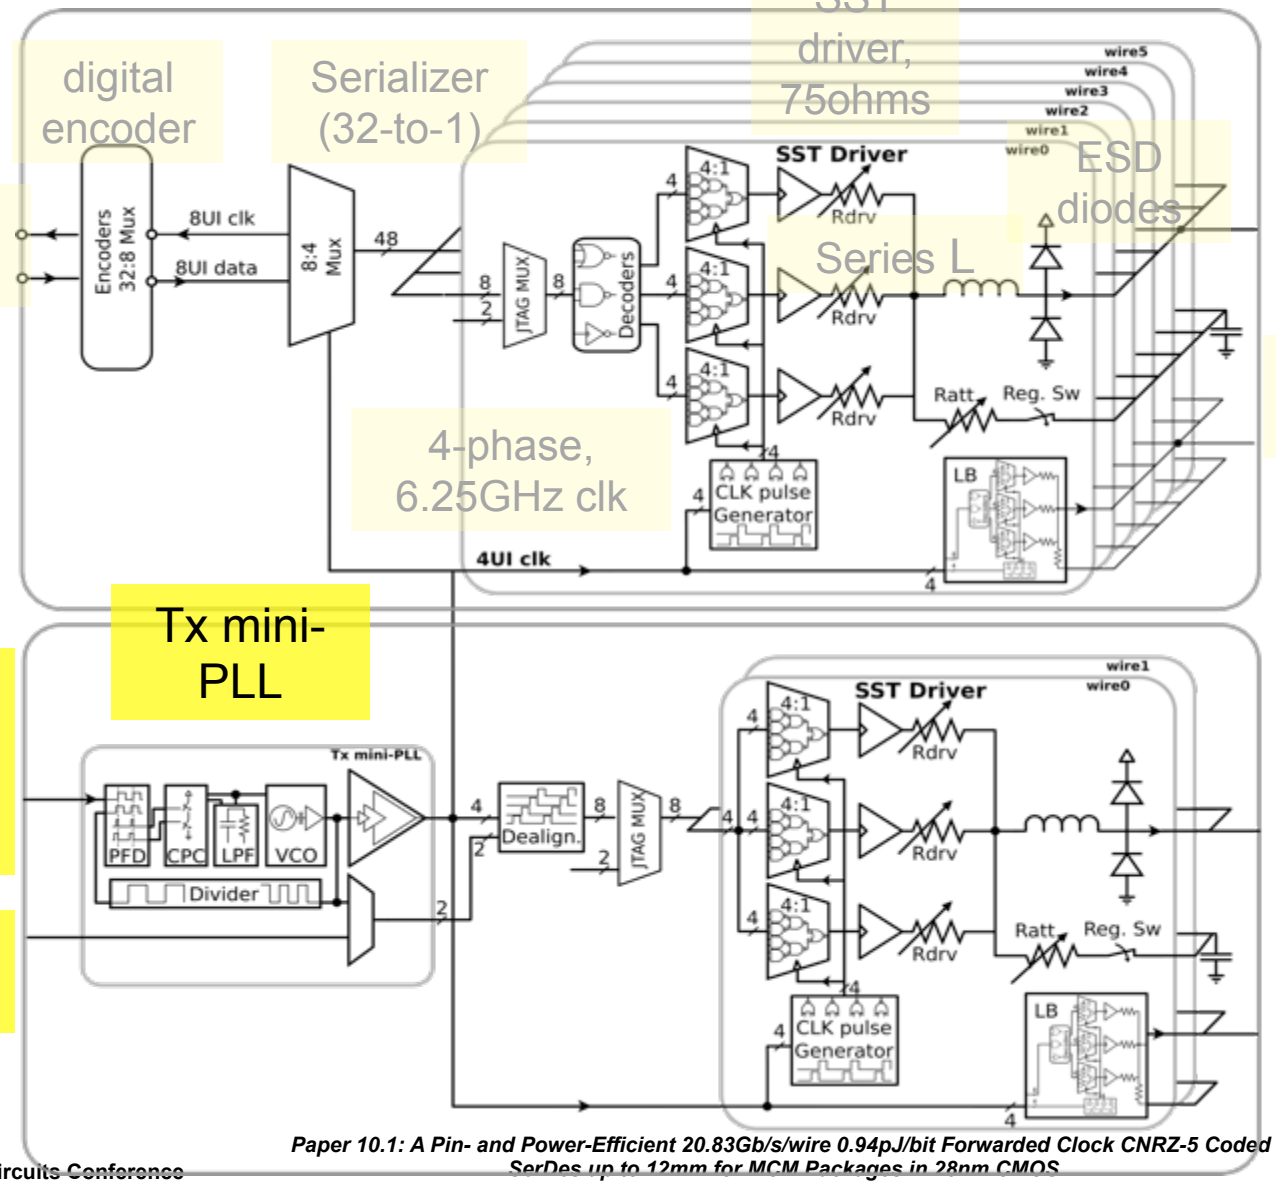

# Transmitter

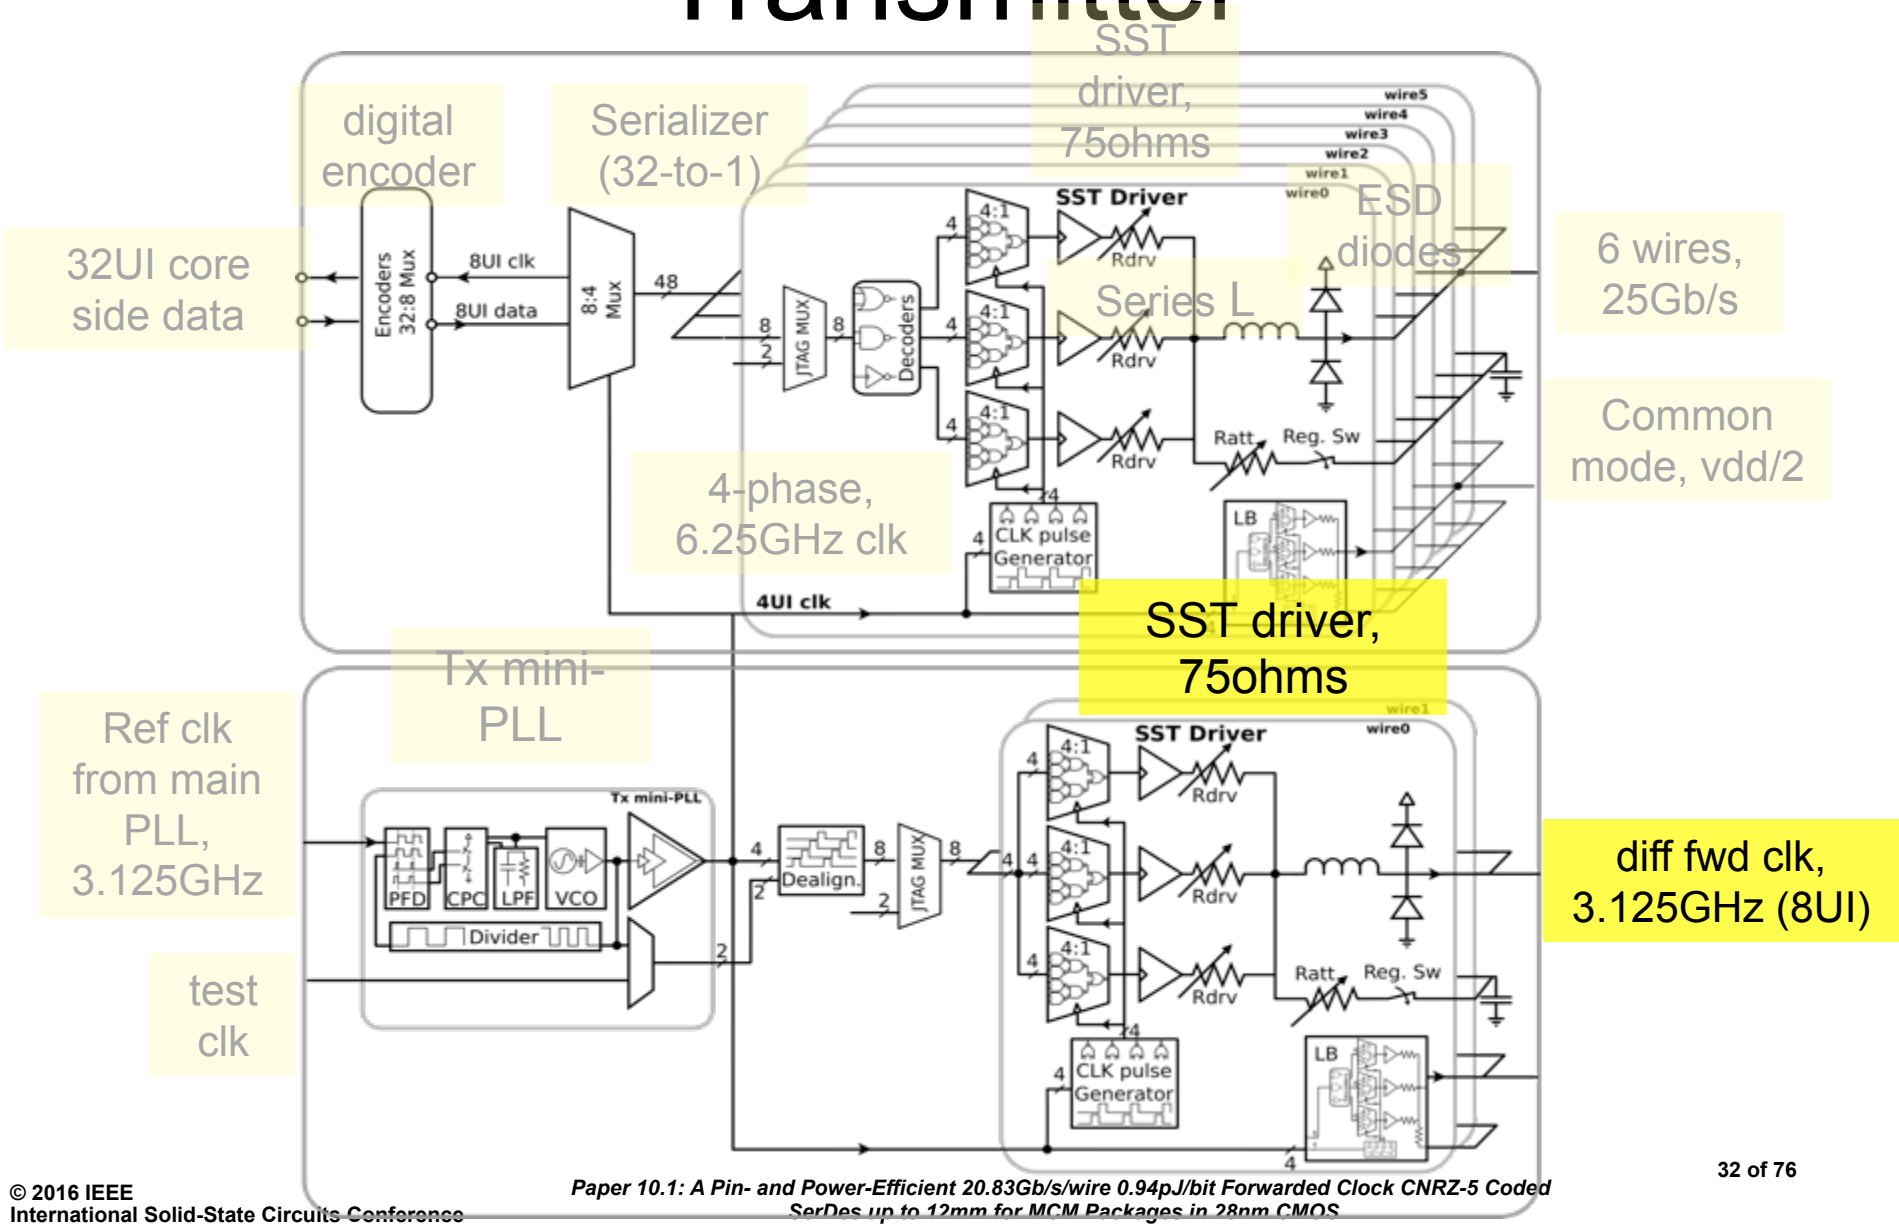

# Transmitter

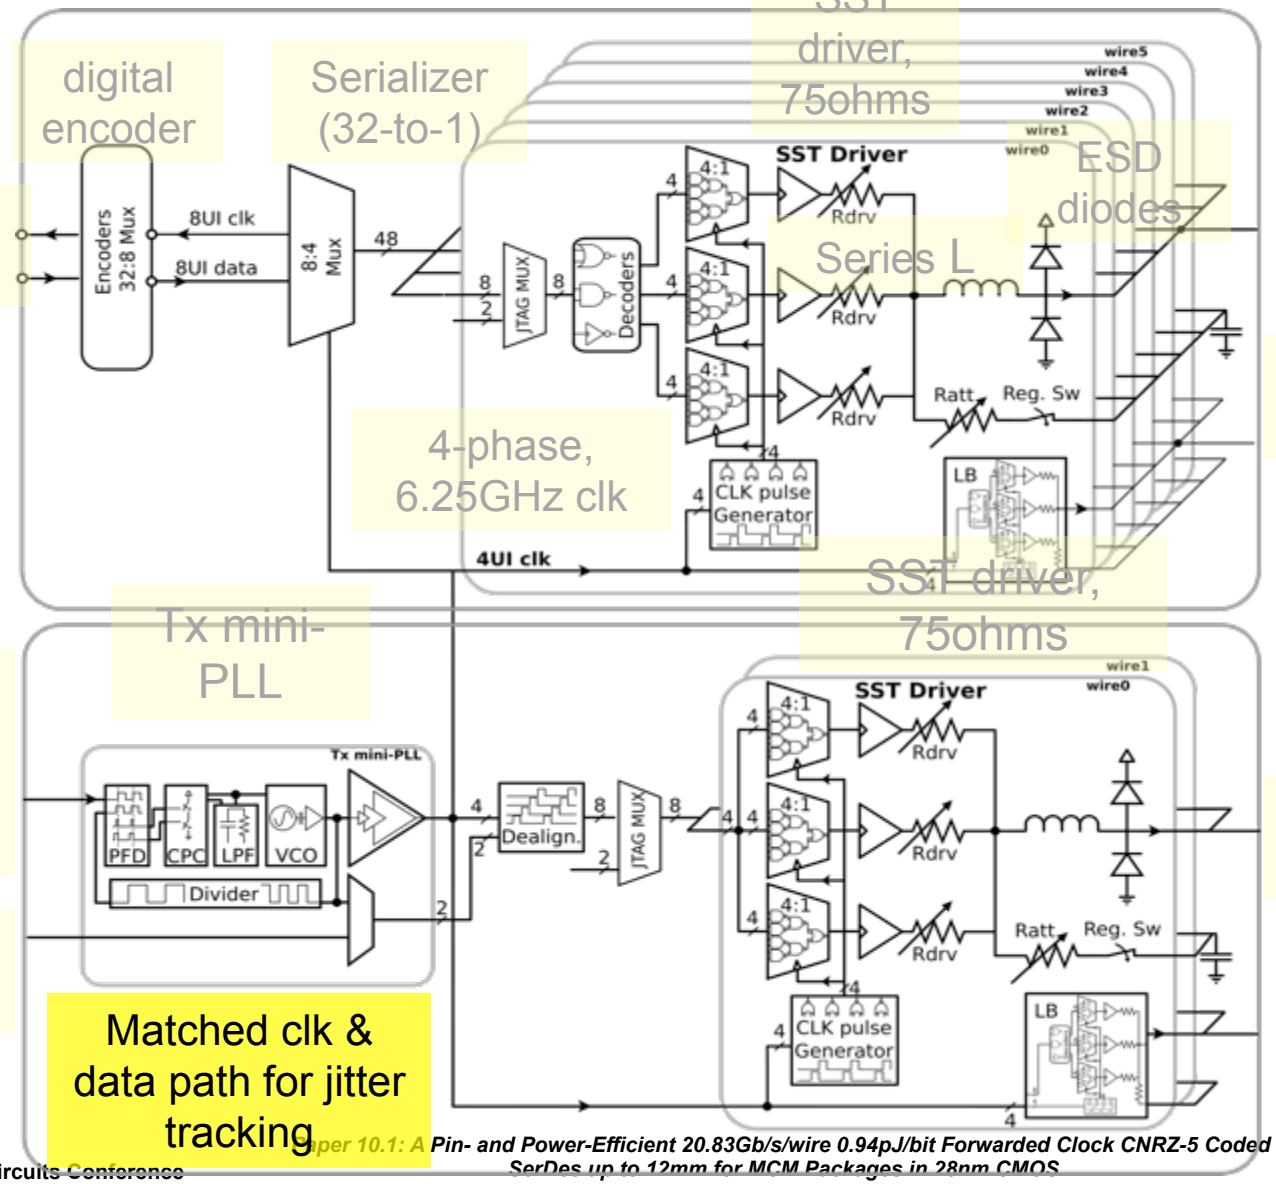

# Transmitter Blocks

- SST output driver

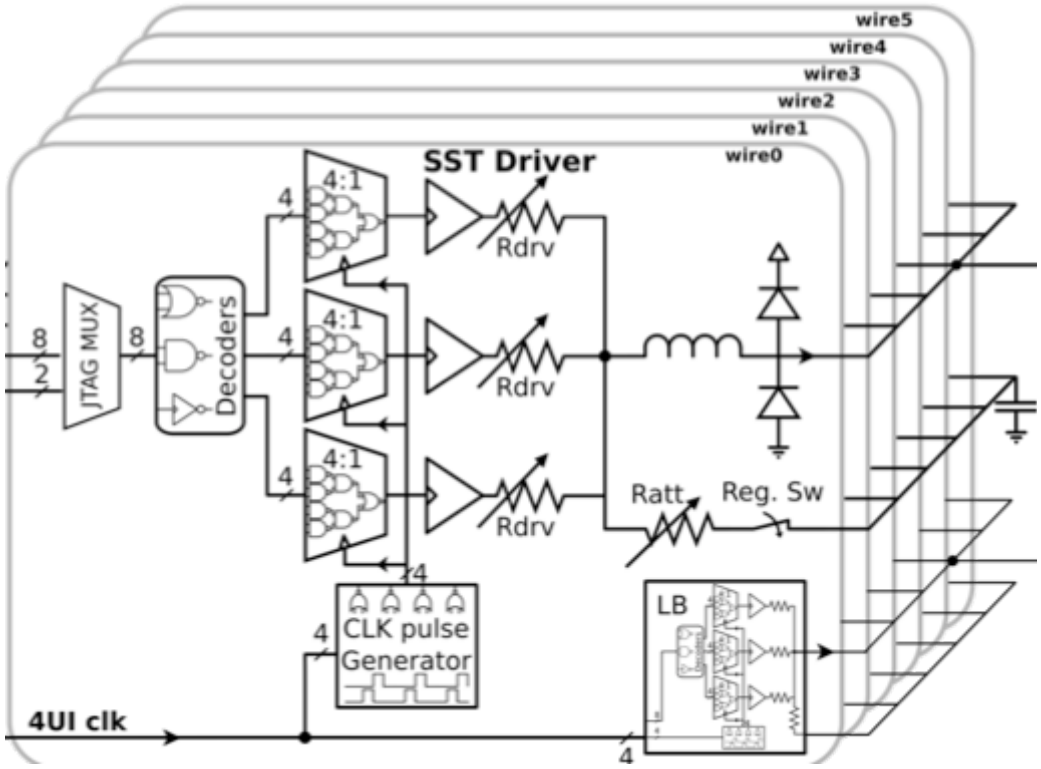

# Transmitter Blocks

- SST output driver
  - 4 levels, 300mVpp

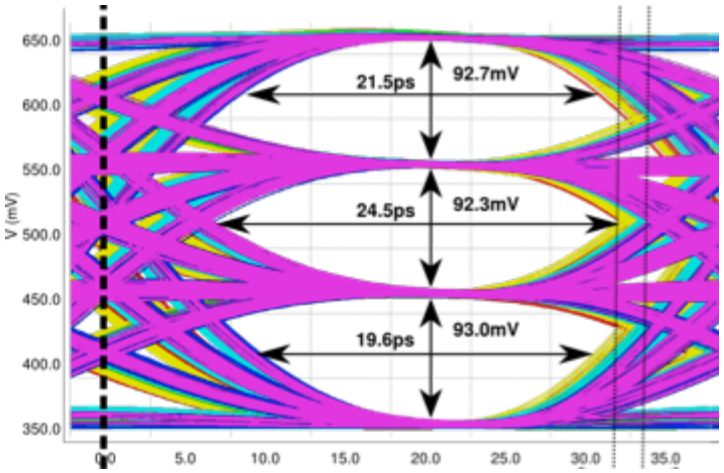


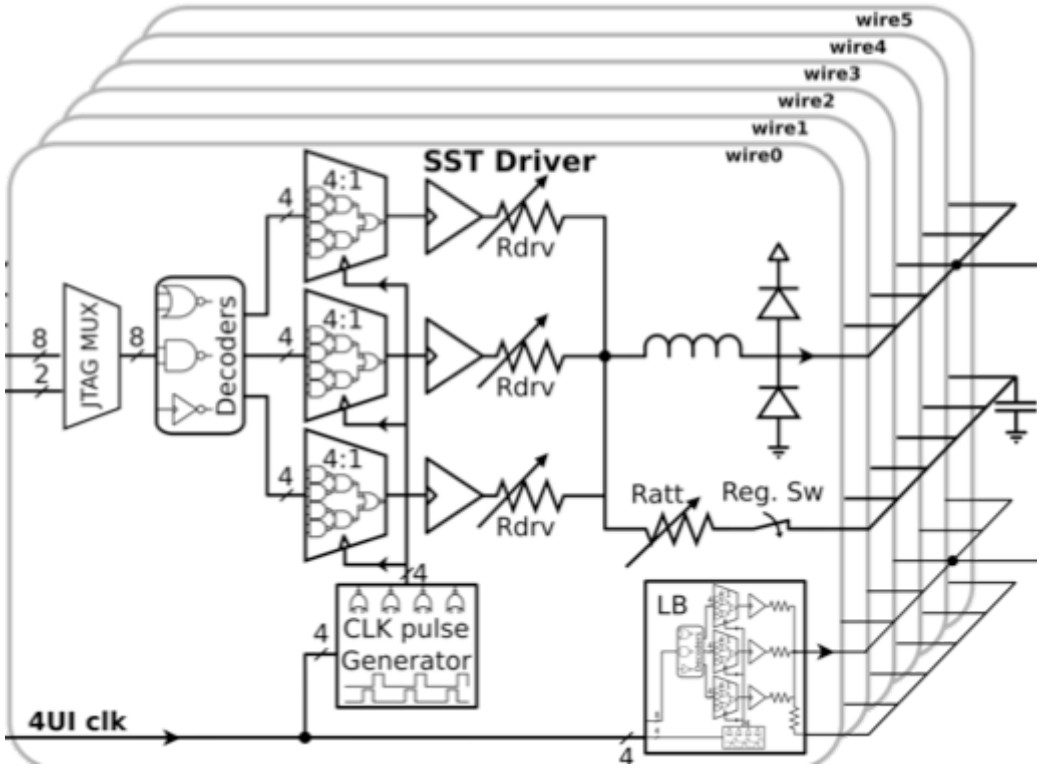

# Transmitter Blocks

- SST output driver
  - 4 levels, 300mVpp
  - Balanced, SSO free

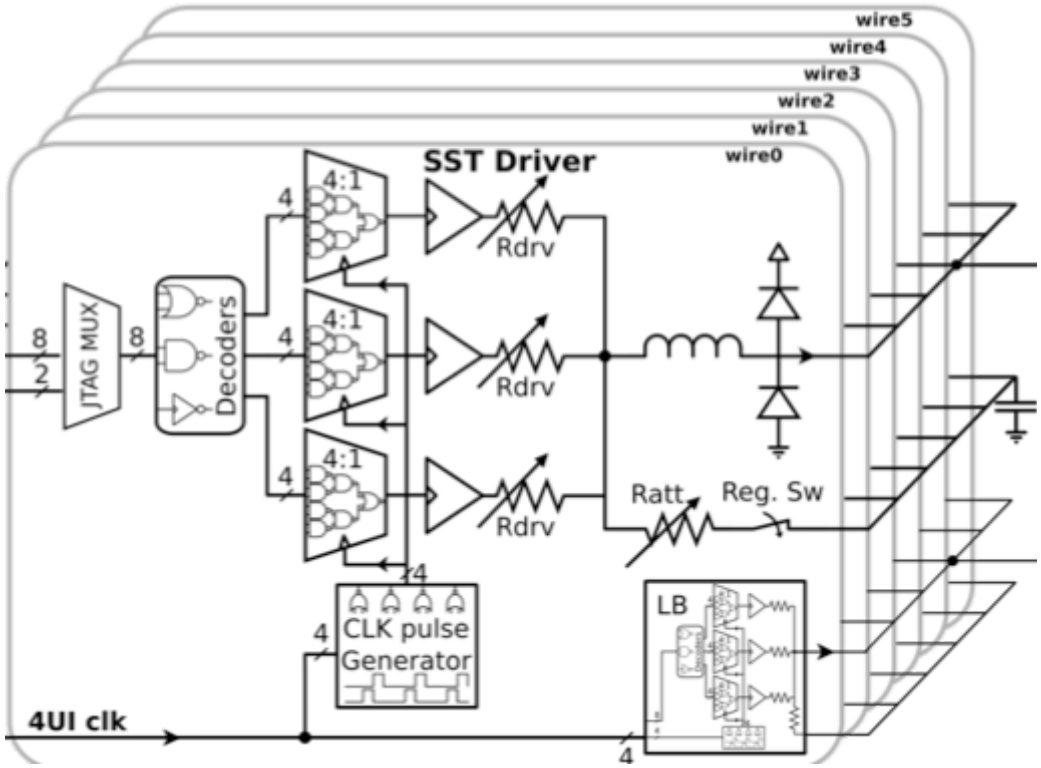

# Transmitter Blocks

- SST output driver
  - 4 levels, 300mVpp
  - Balanced, SSO free
  - 75 ohms
    - 125Ω || 190Ω
  - Vcm=Vdd/2
  - ¼ rate arch

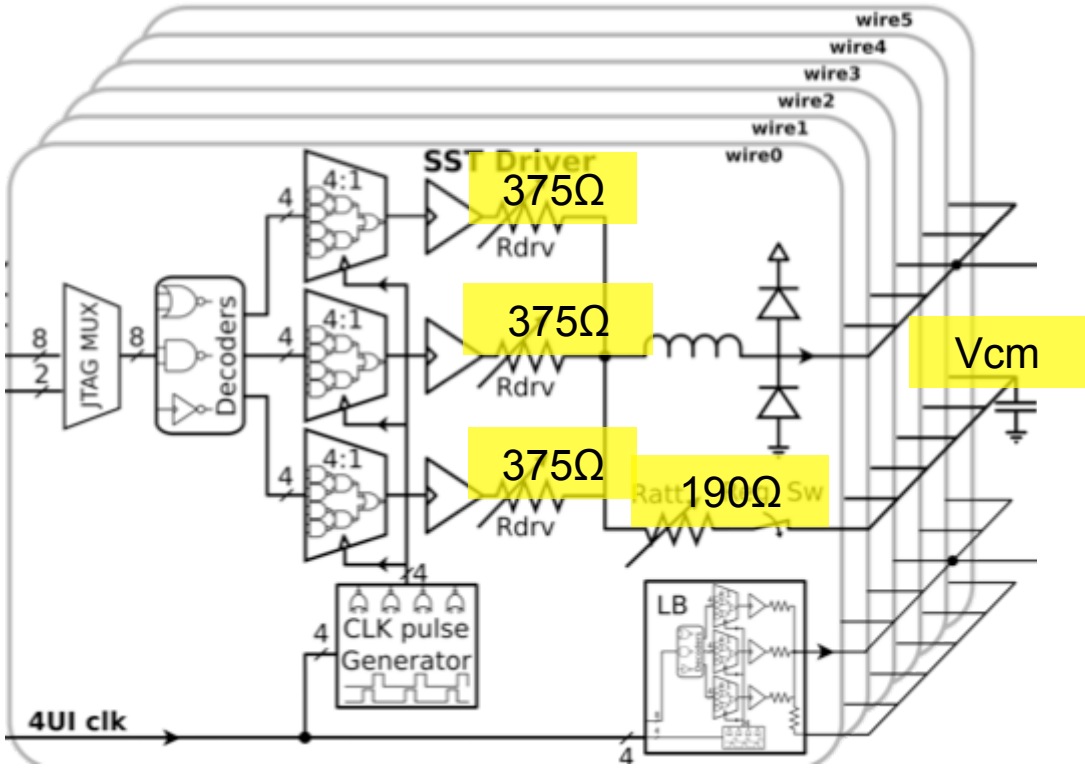

# Transmitter Blocks

- SST output driver
  - 4 levels, 300mVpp
  - Balanced, SSO free
  - 75 ohms
    - 125Ω || 190Ω
  - Vcm=Vdd/2
  - ¼ rate arch
- CMOS output in JTAG/test mode

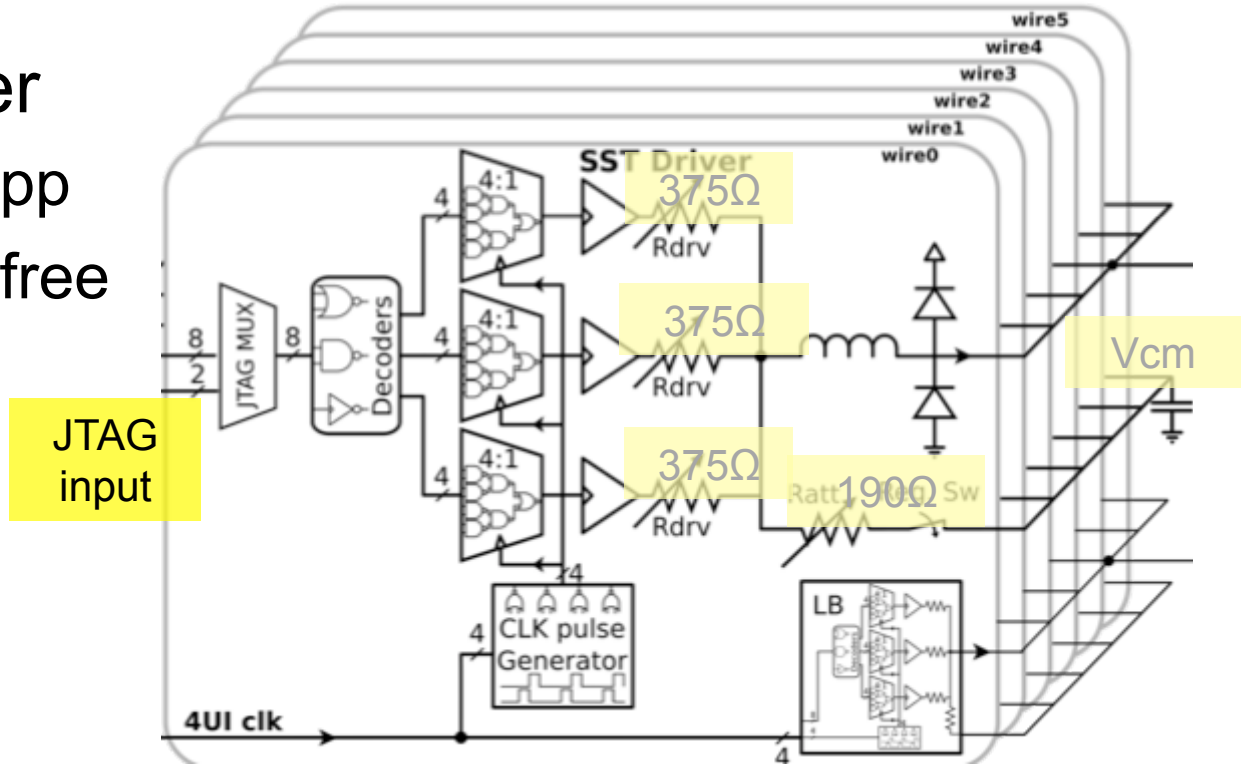

# Receiver

# Receiver

- Continuous Time Front End:
  - DC coupled, with level shifter
  - T-Coils for passive equalization
  - Multi-input gain stage (decoder) that performs linear combination of incoming signal

# Receiver

- Continuous Time Front End:
  - DC coupled, with level shifter
  - T-Coils for passive equalization
  - Multi-input gain stage (decoder) that performs linear combination of incoming signal
- Discrete Time Front End:
  - 4-ph data sampling system, followed by 4:32 demux

# Receiver

- Continuous Time Front End:
  - DC coupled, with level shifter
  - T-Coils for passive equalization
  - Multi-input gain stage (decoder) that performs linear combination of incoming signal
- Discrete Time Front End:
  - 4-ph data sampling system, followed by 4:32 demux
- Clock data alignment block (CDA) is used to align sampling clock edge to center of data eye

# Receiver

- Continuous Time Front End:
  - DC coupled, with level shifter
  - T-Coils for passive equalization
  - Multi-input gain stage (decoder) that performs linear combination of incoming signal
- Discrete Time Front End:
  - 4-ph data sampling system, followed by 4:32 demux
- Clock data alignment block (CDA) is used to align sampling clock edge to center of data eye
- Wide-band PLL to track correlated jitter (data/clk)
  - Uses 8UI fwd clock as ref clock

# Receiver Block Diagrams

# Receiver Block Diagrams

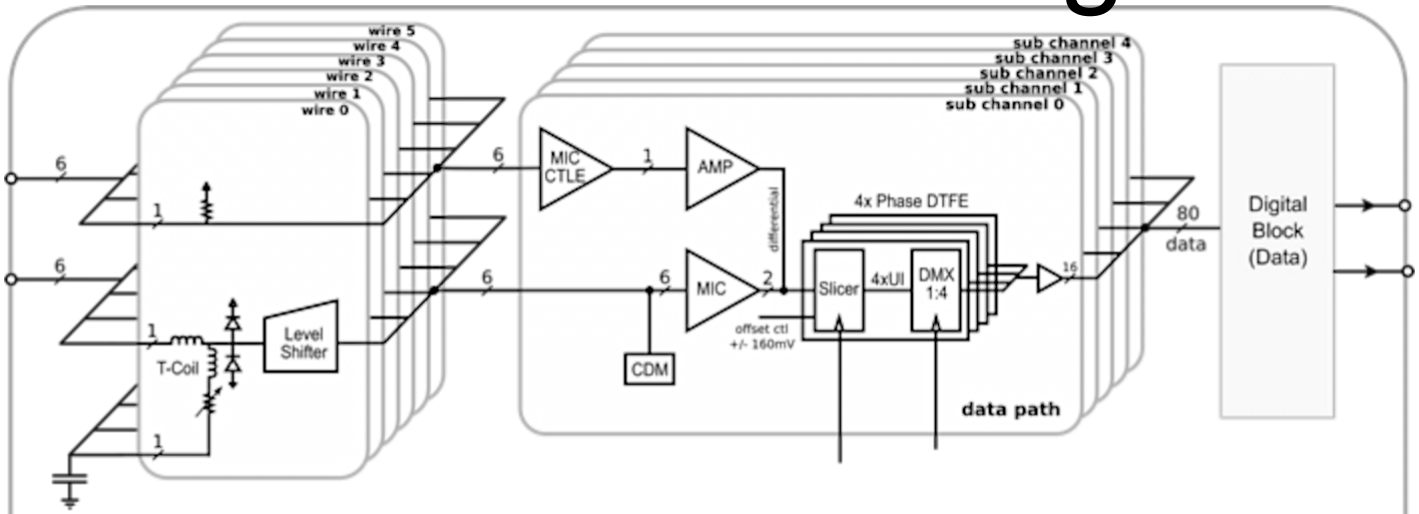

# Receiver Block Diagrams

Rx input
6 wires, 5b
@25Gbaud

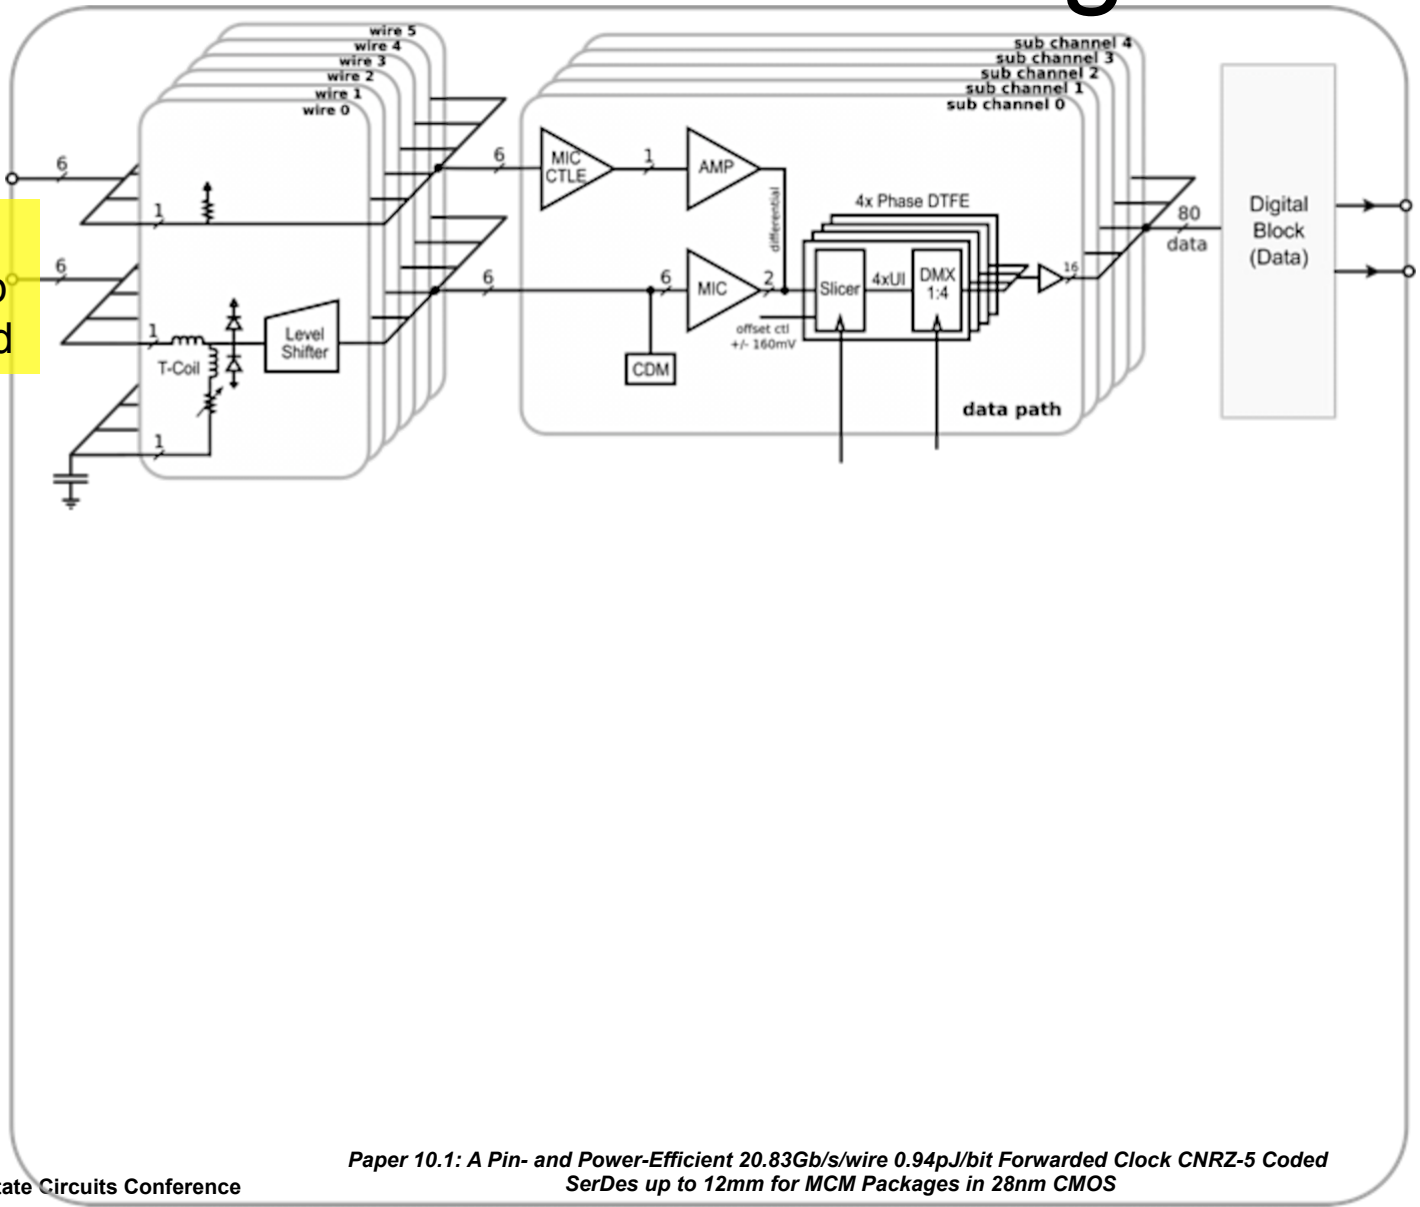


160b data
32UI clk

# Receiver Block Diagrams

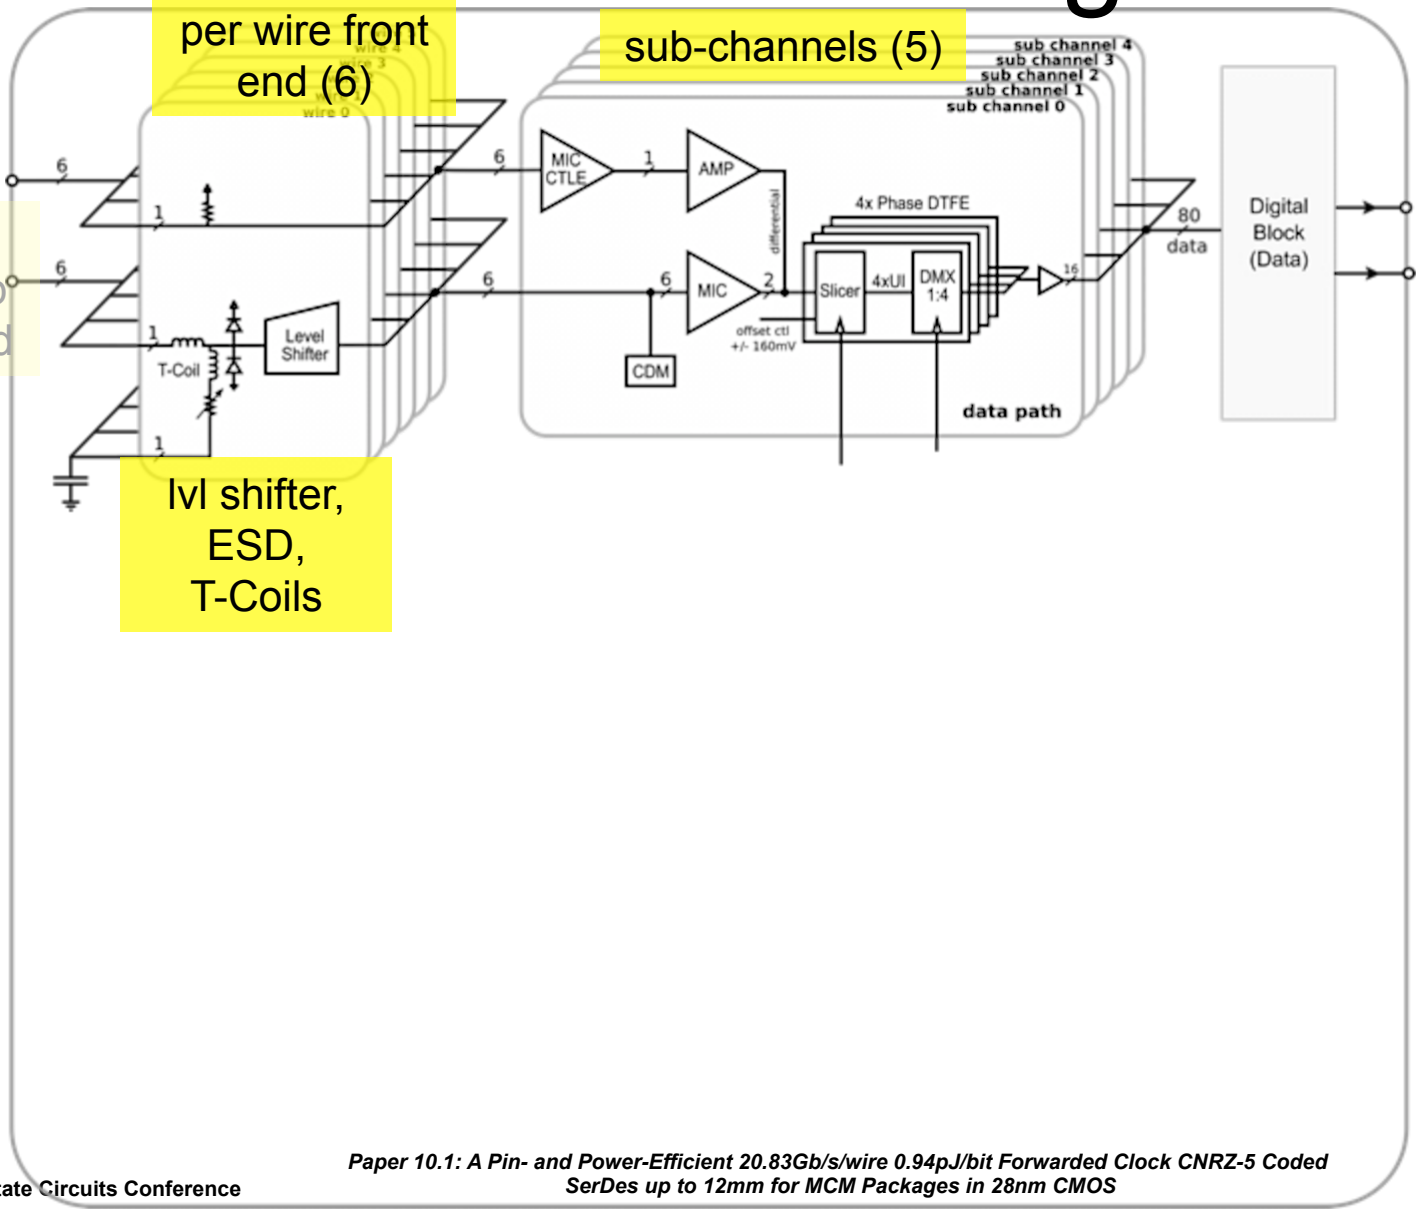

# Receiver Block Diagrams

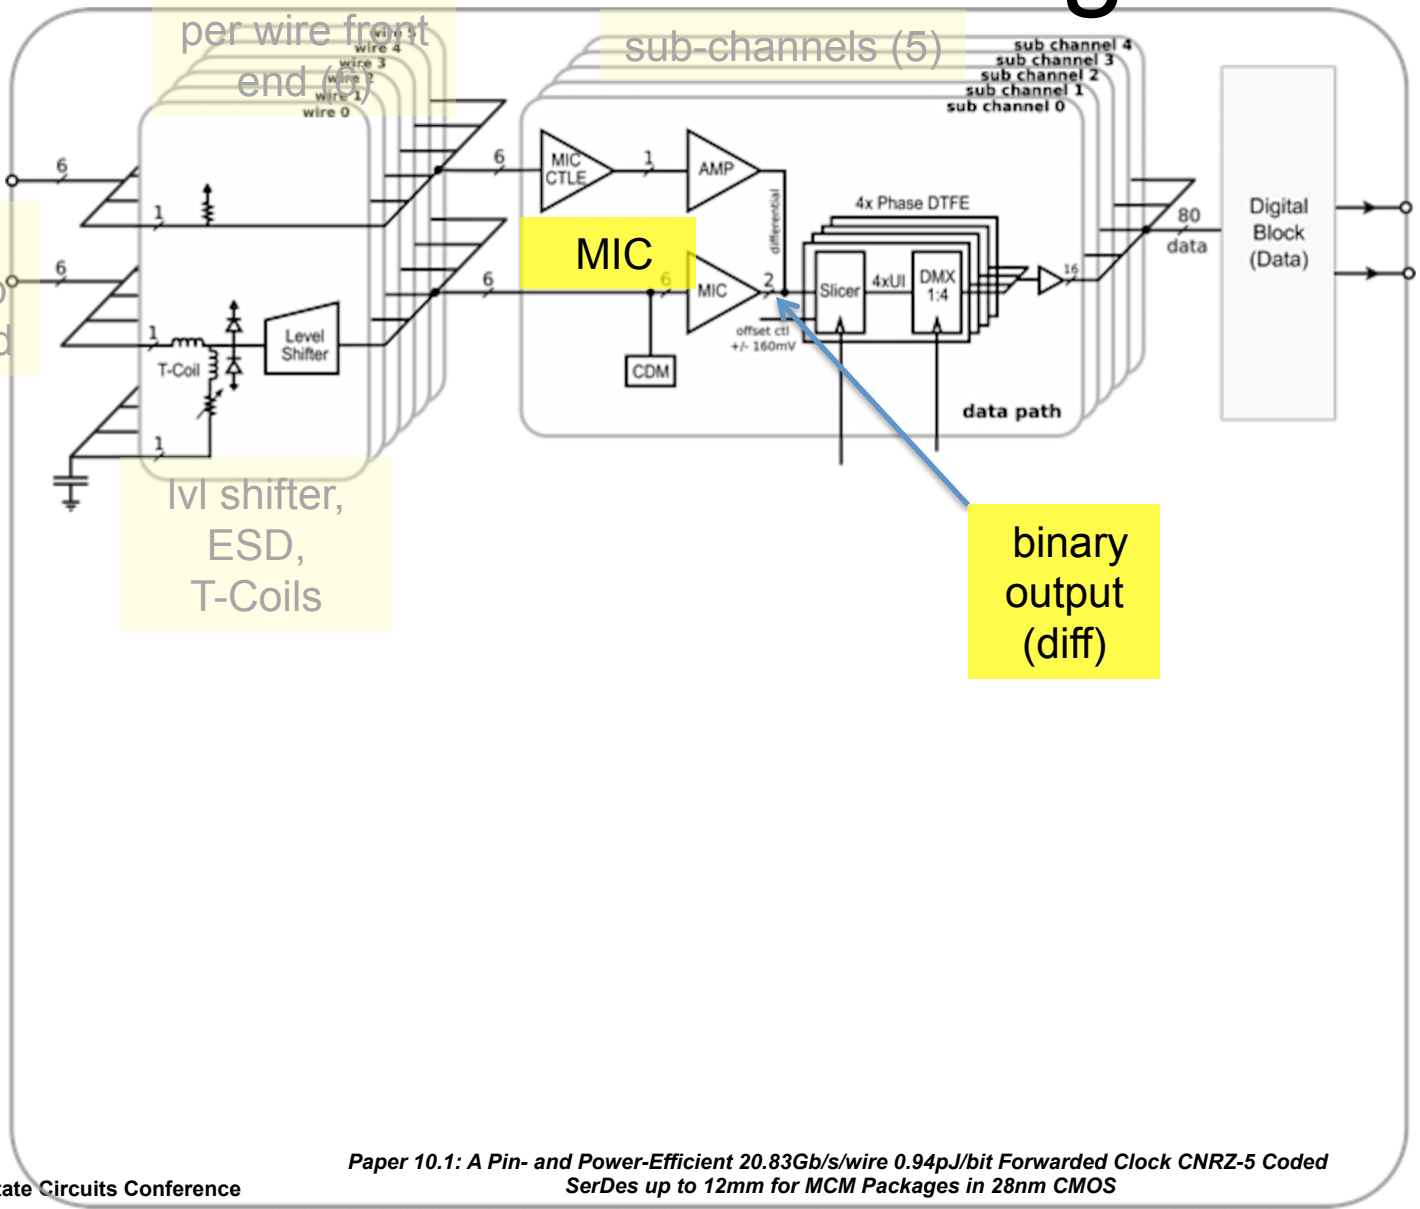

# Receiver Block Diagrams

Rx input
6 wires, 5b
@25Gbaud

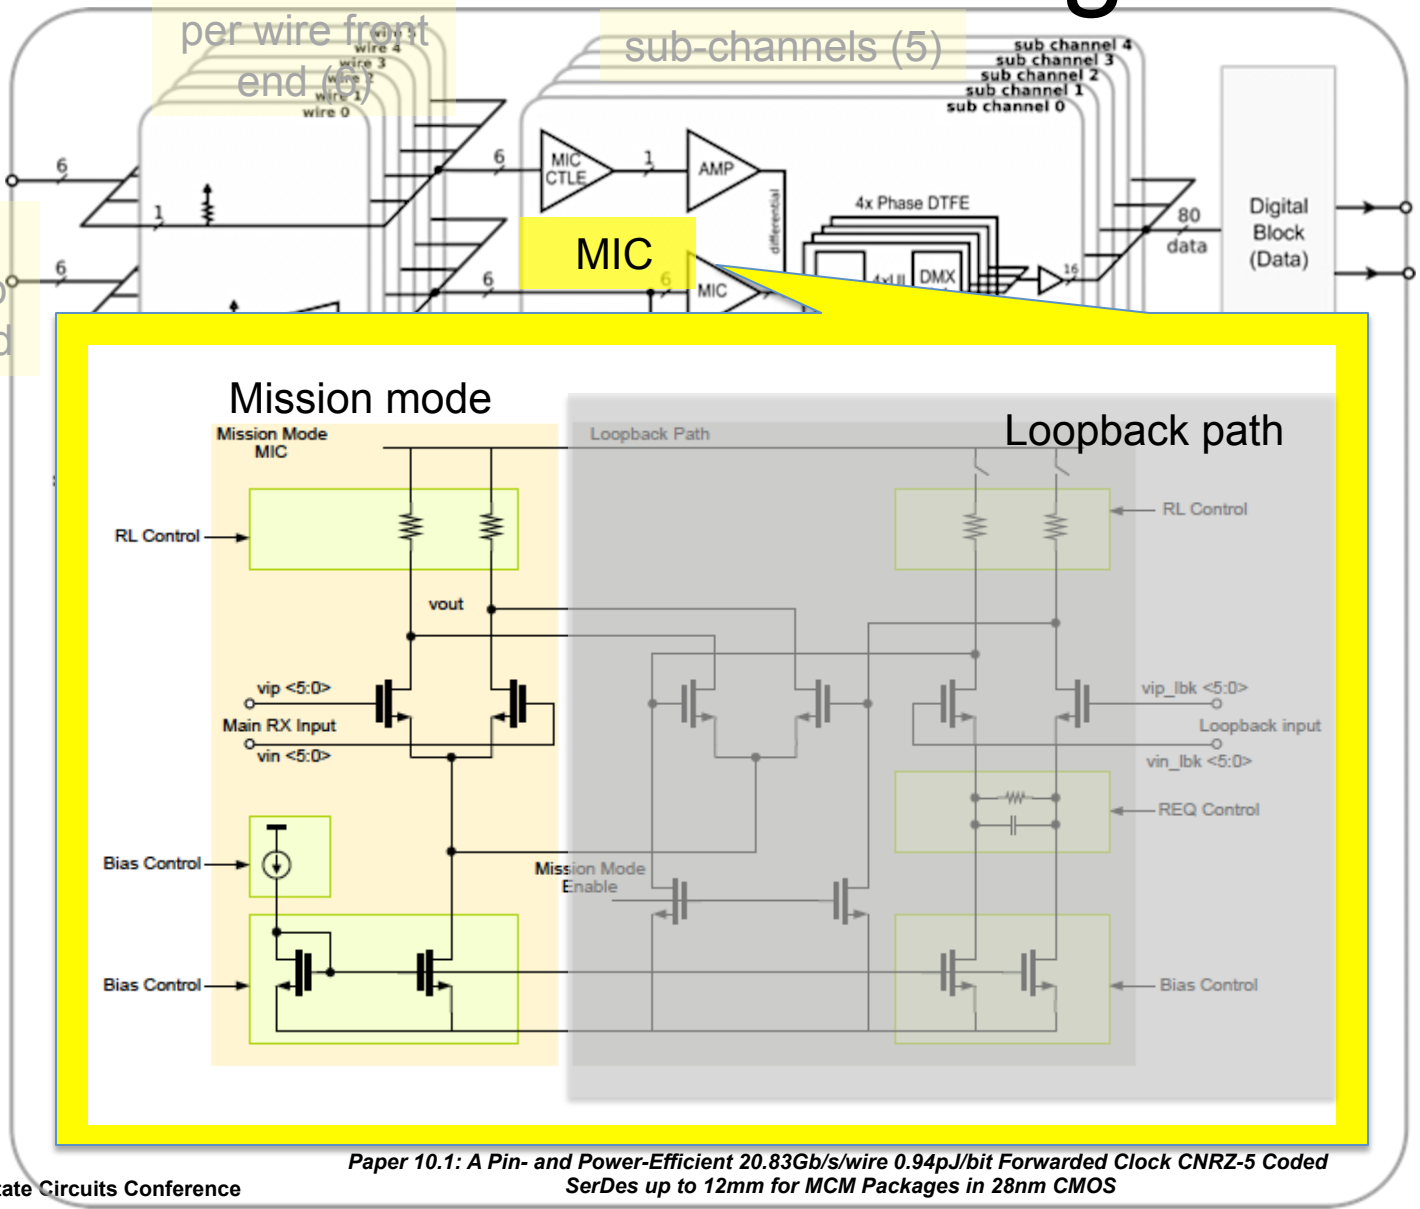


160b data
32UI clk

# Receiver Block Diagrams

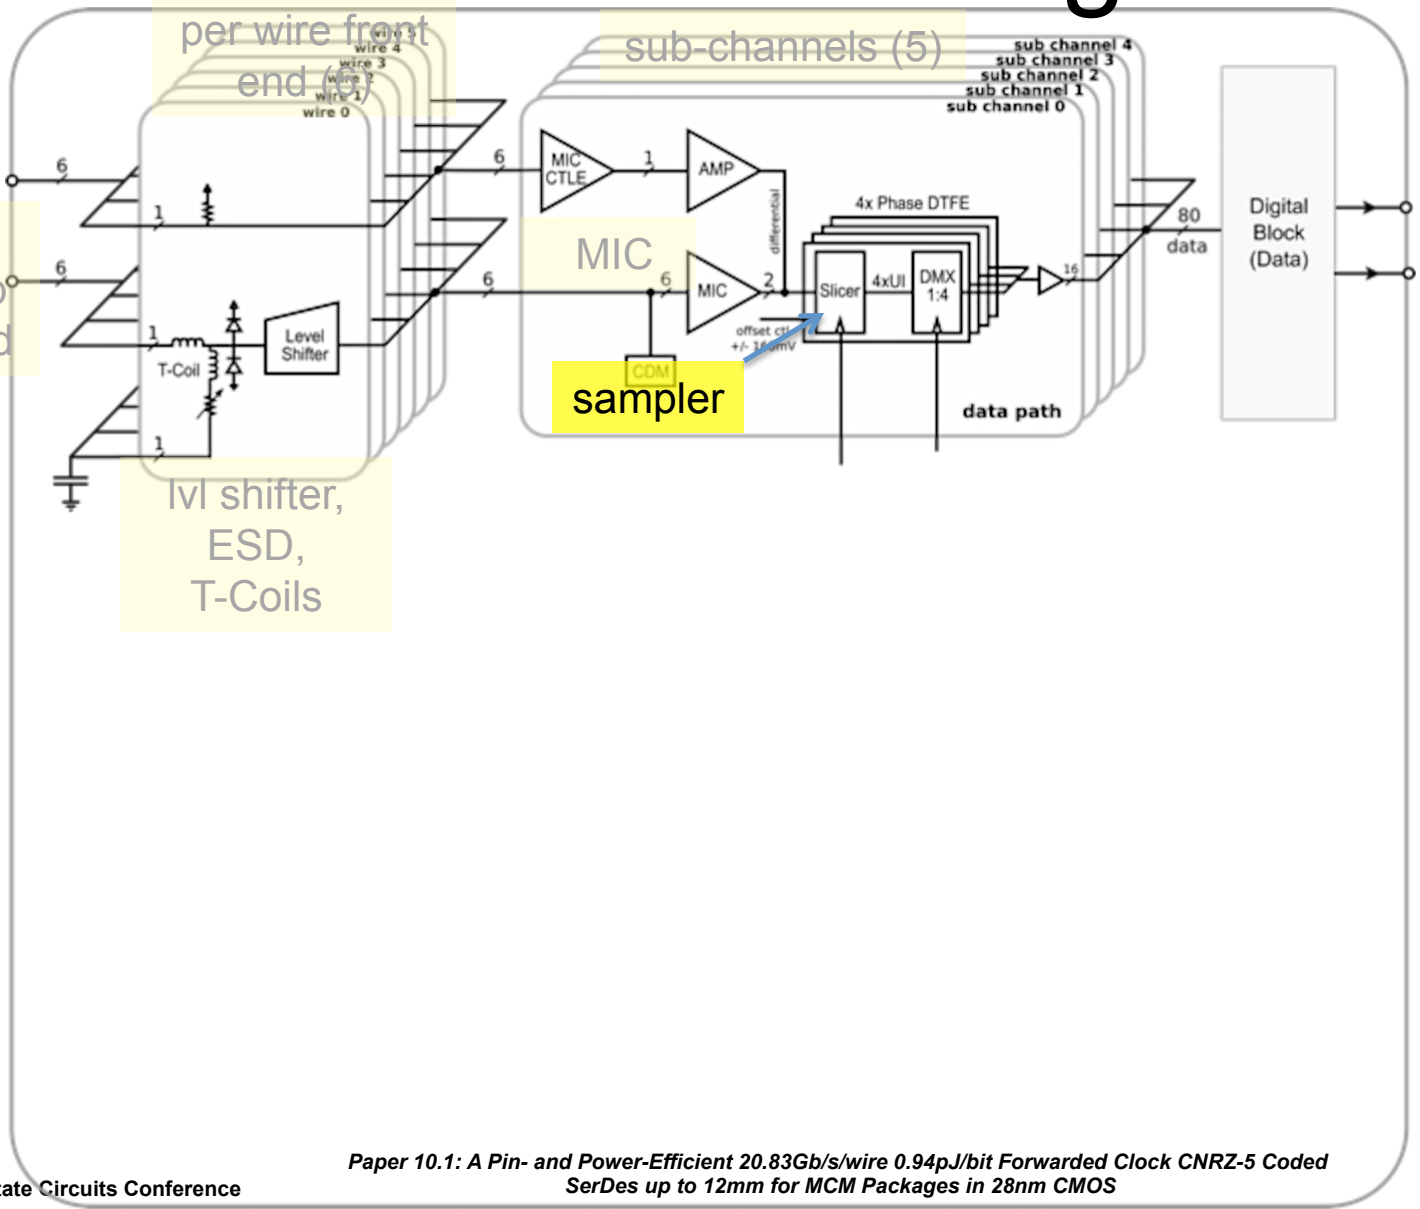

# Receiver Block Diagrams

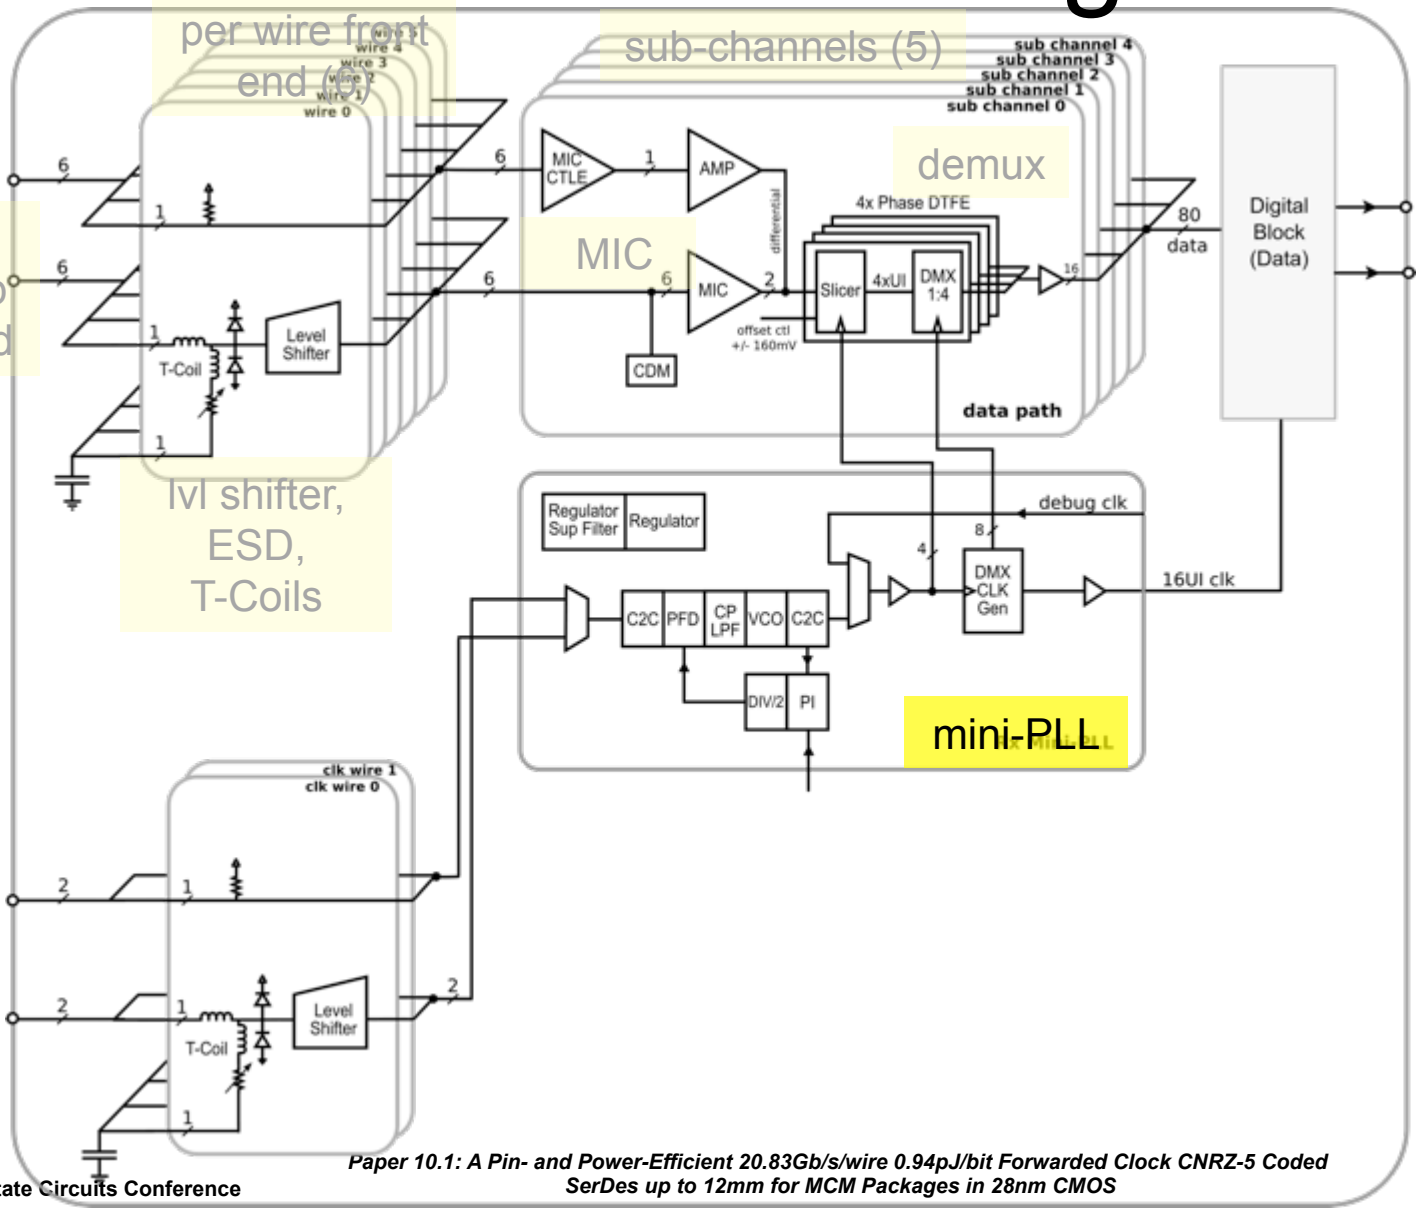


Rx input
6 wires, 5b
@25Gbaud

lvl shifter,
ESD,
T-Coils

160b data
32UI clk

# Receiver Block Diagrams

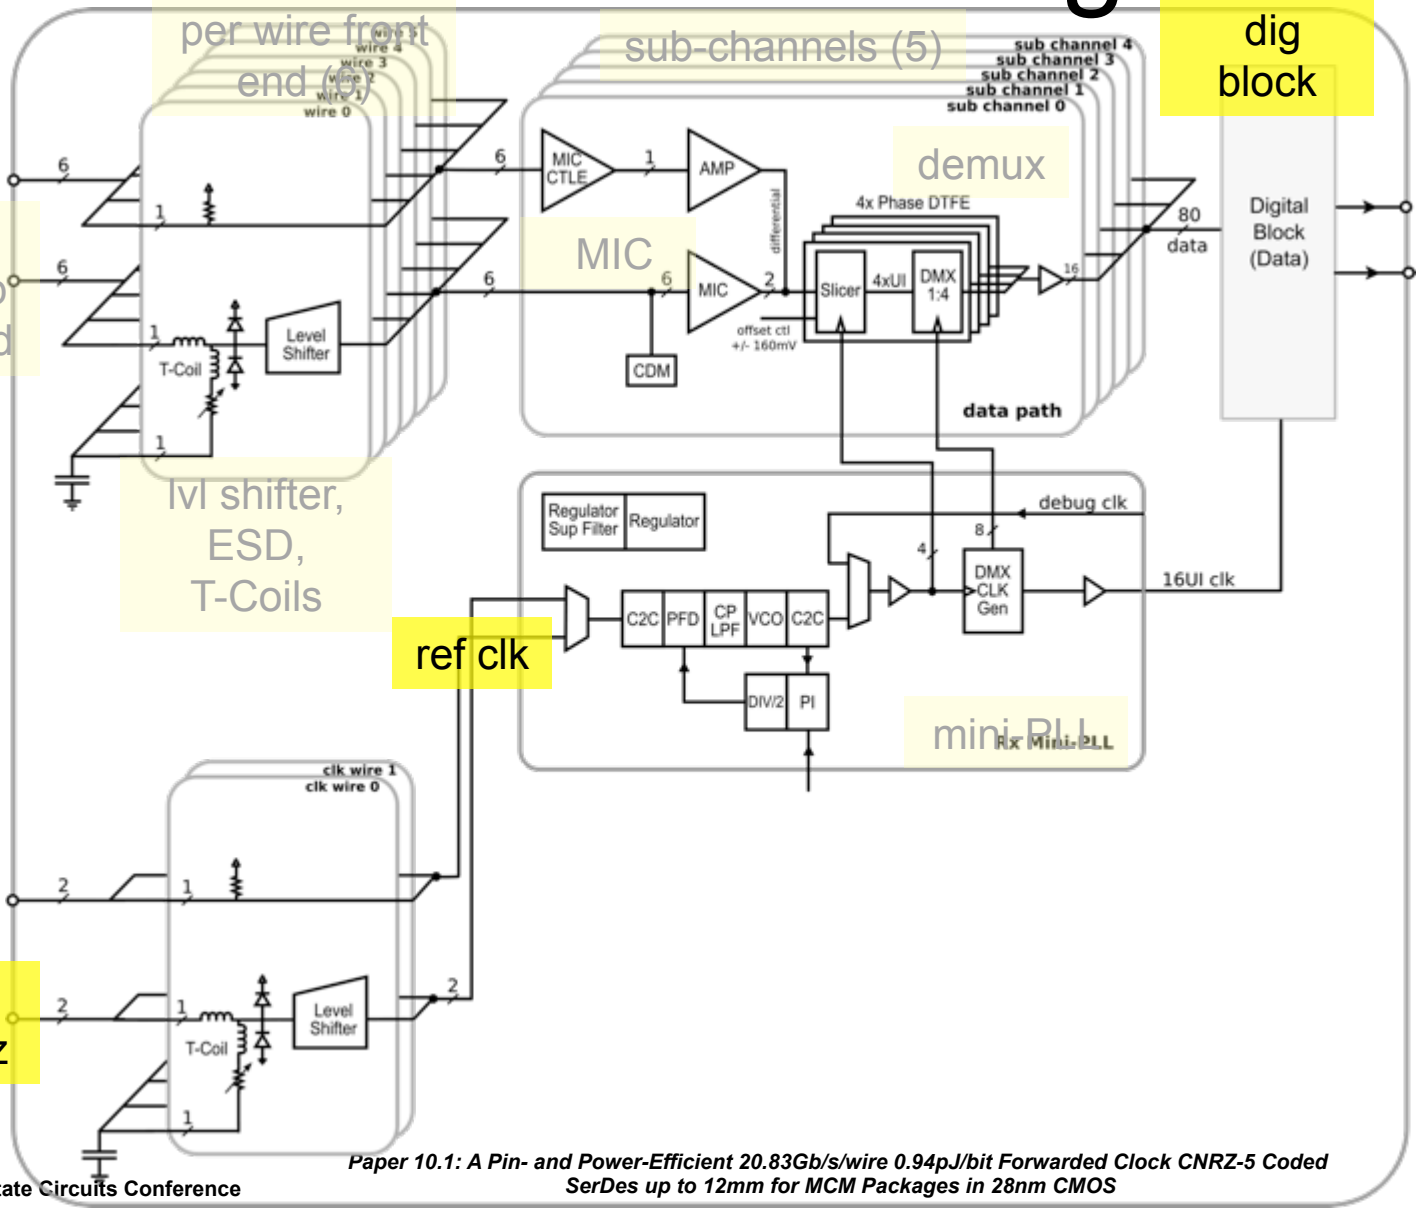

# Receiver Block Diagrams

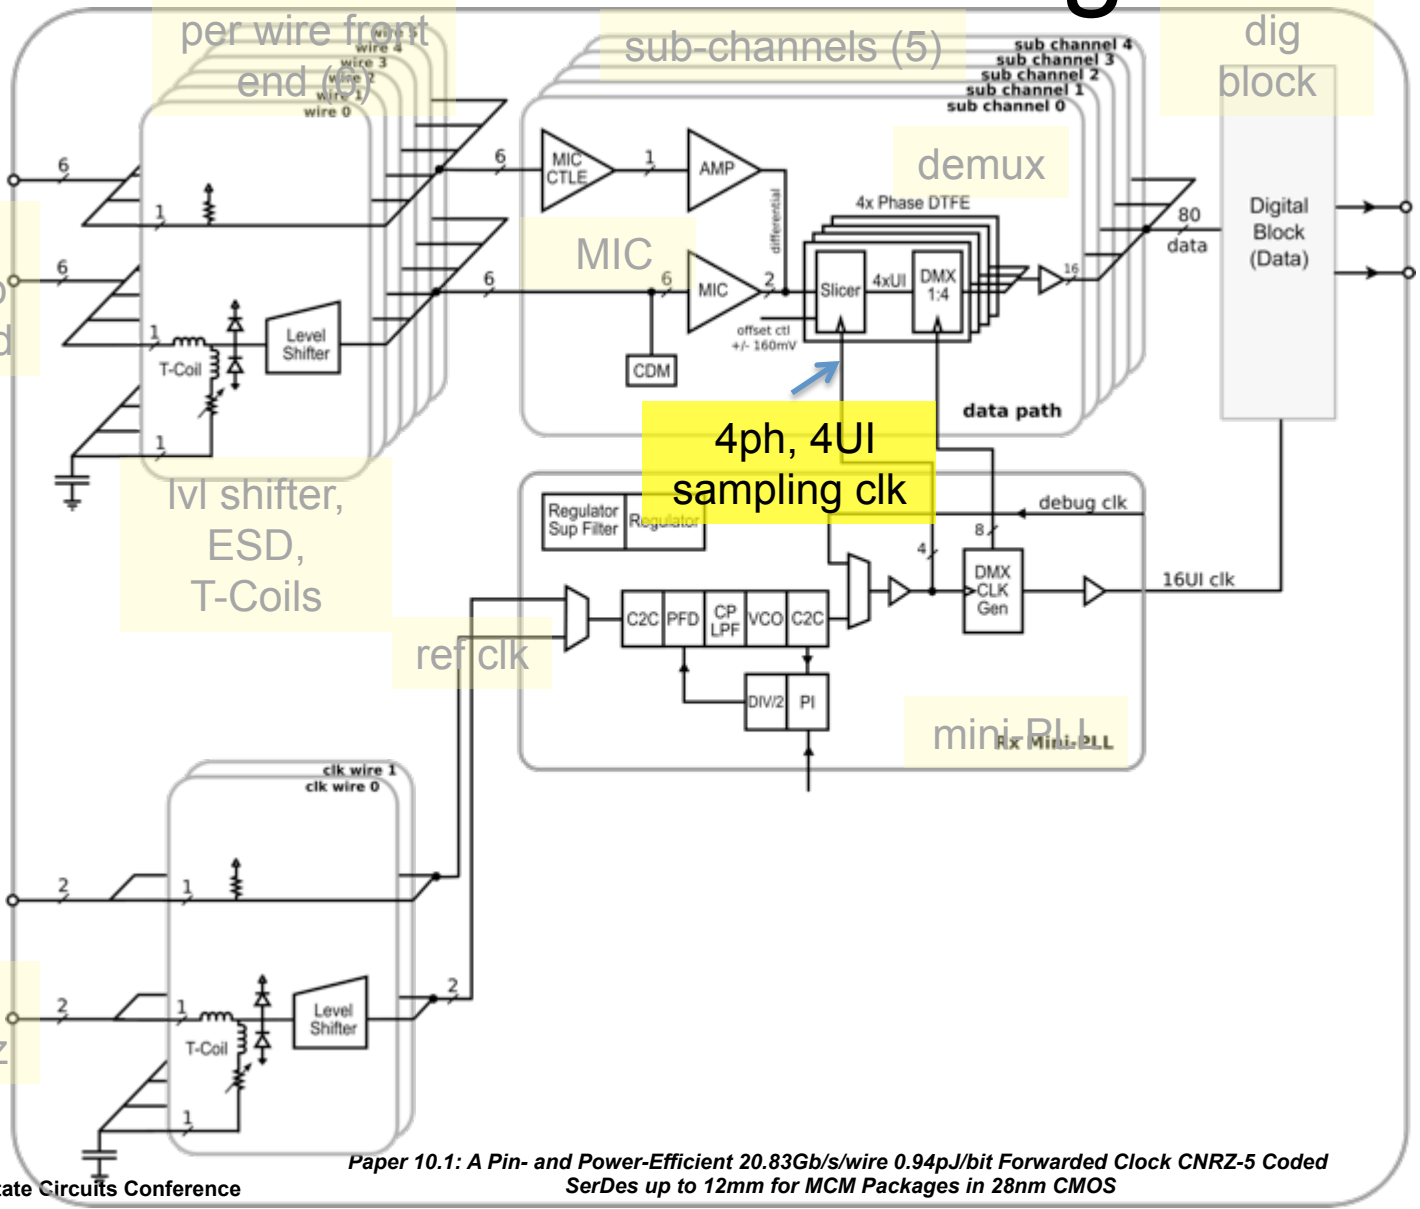

# Receiver Block Diagrams

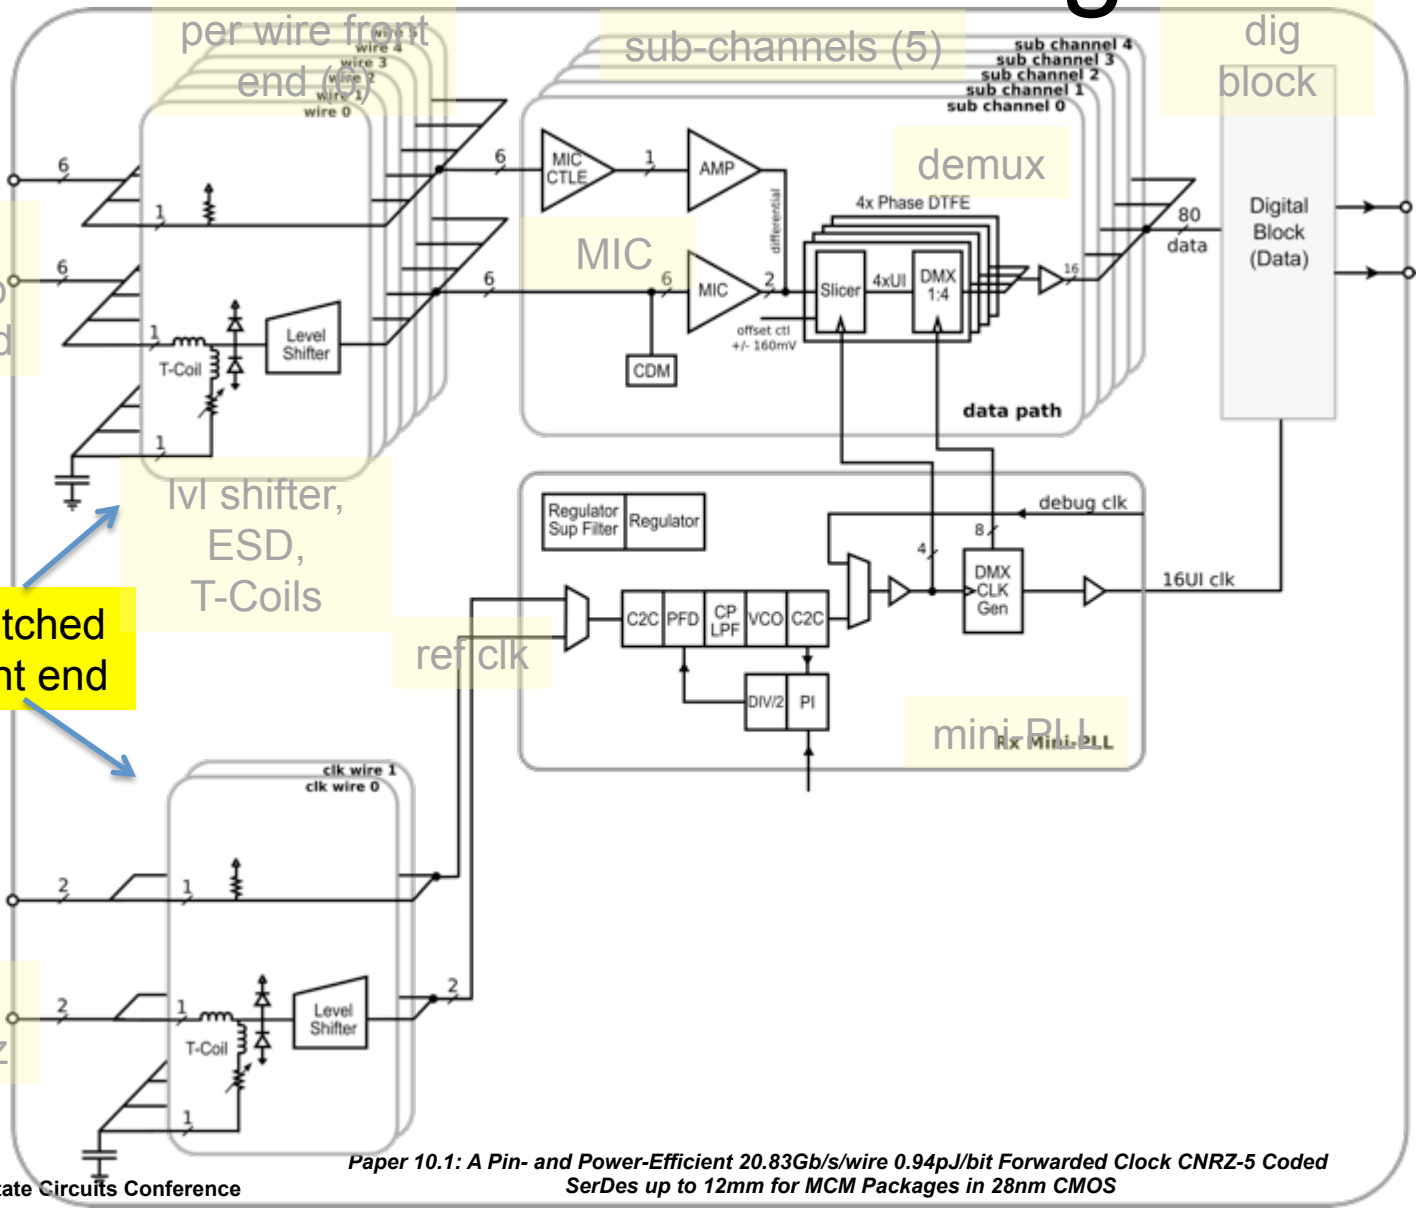


Rx input
6 wires, 5b
@25Gbaud

lvl shifter,
ESD,
T-Coils

matched
front end

ref clk

mini-PLL

fwd clk
3.125GHz

160b data
32UI clk

# Receiver Block Diagrams

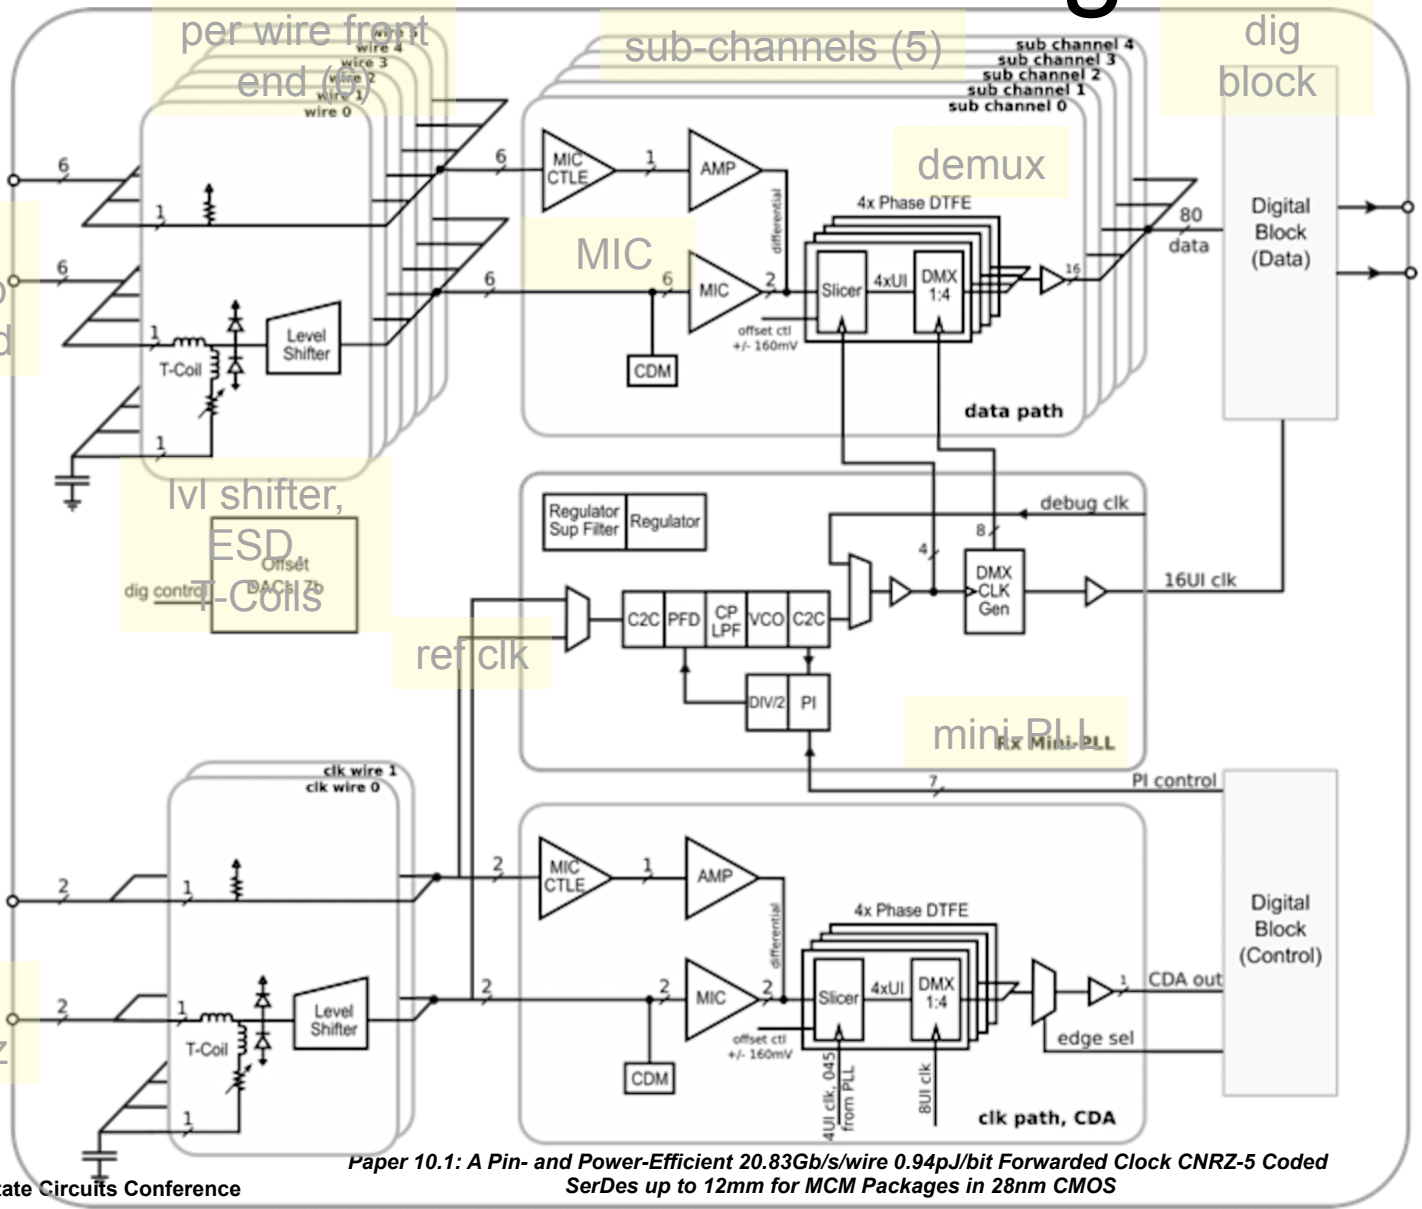

# Receiver Block Diagrams

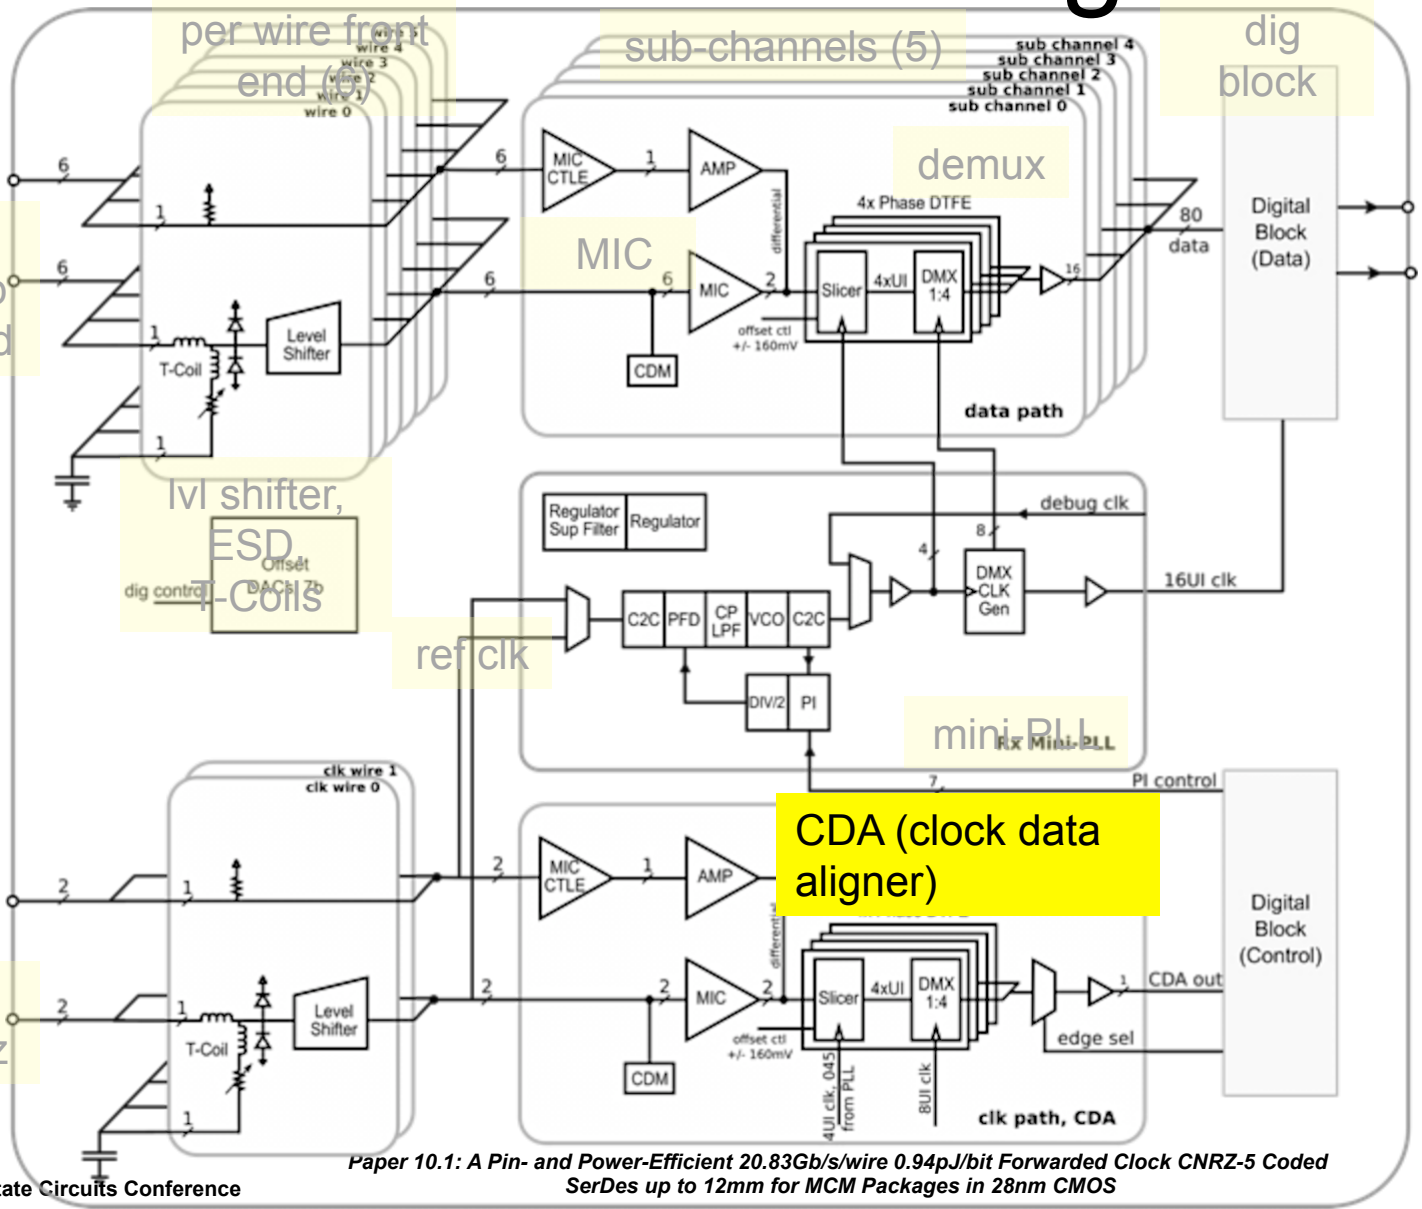

# Receiver Block Diagrams

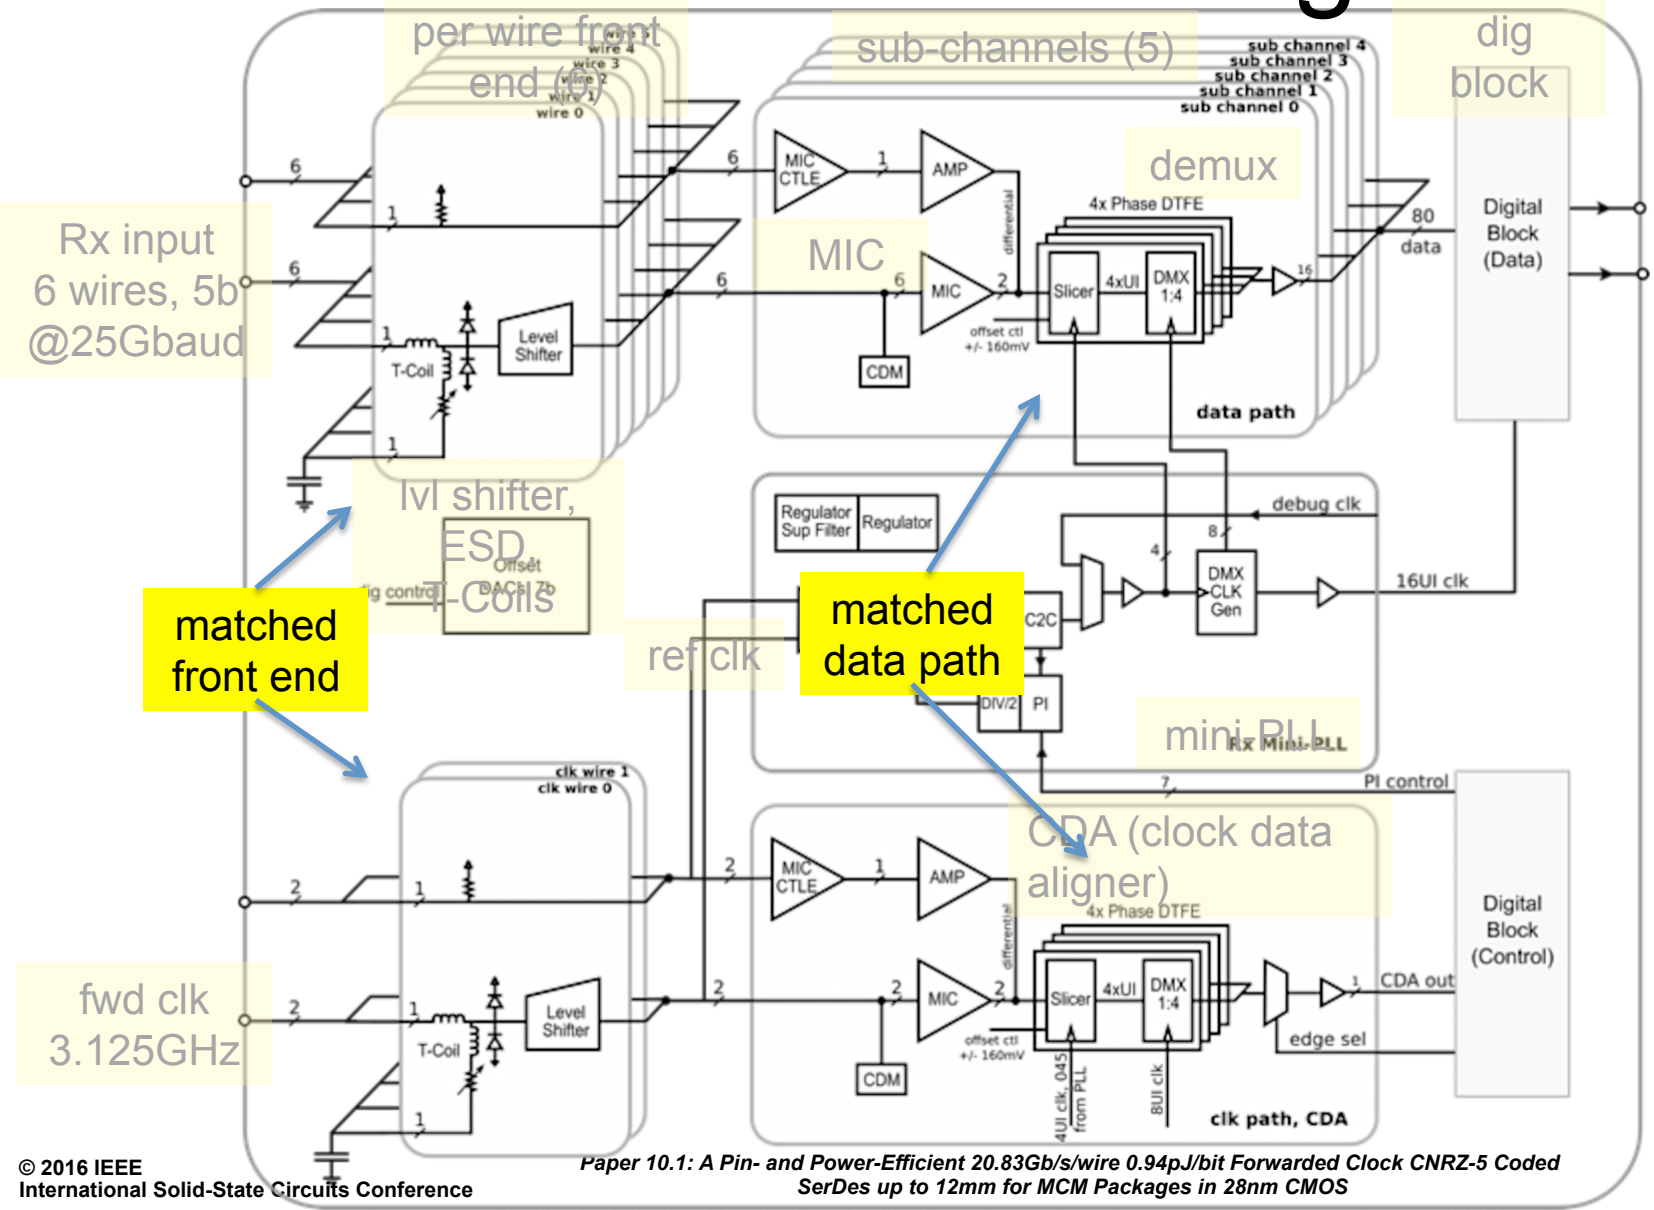

# Receiver Block Diagrams

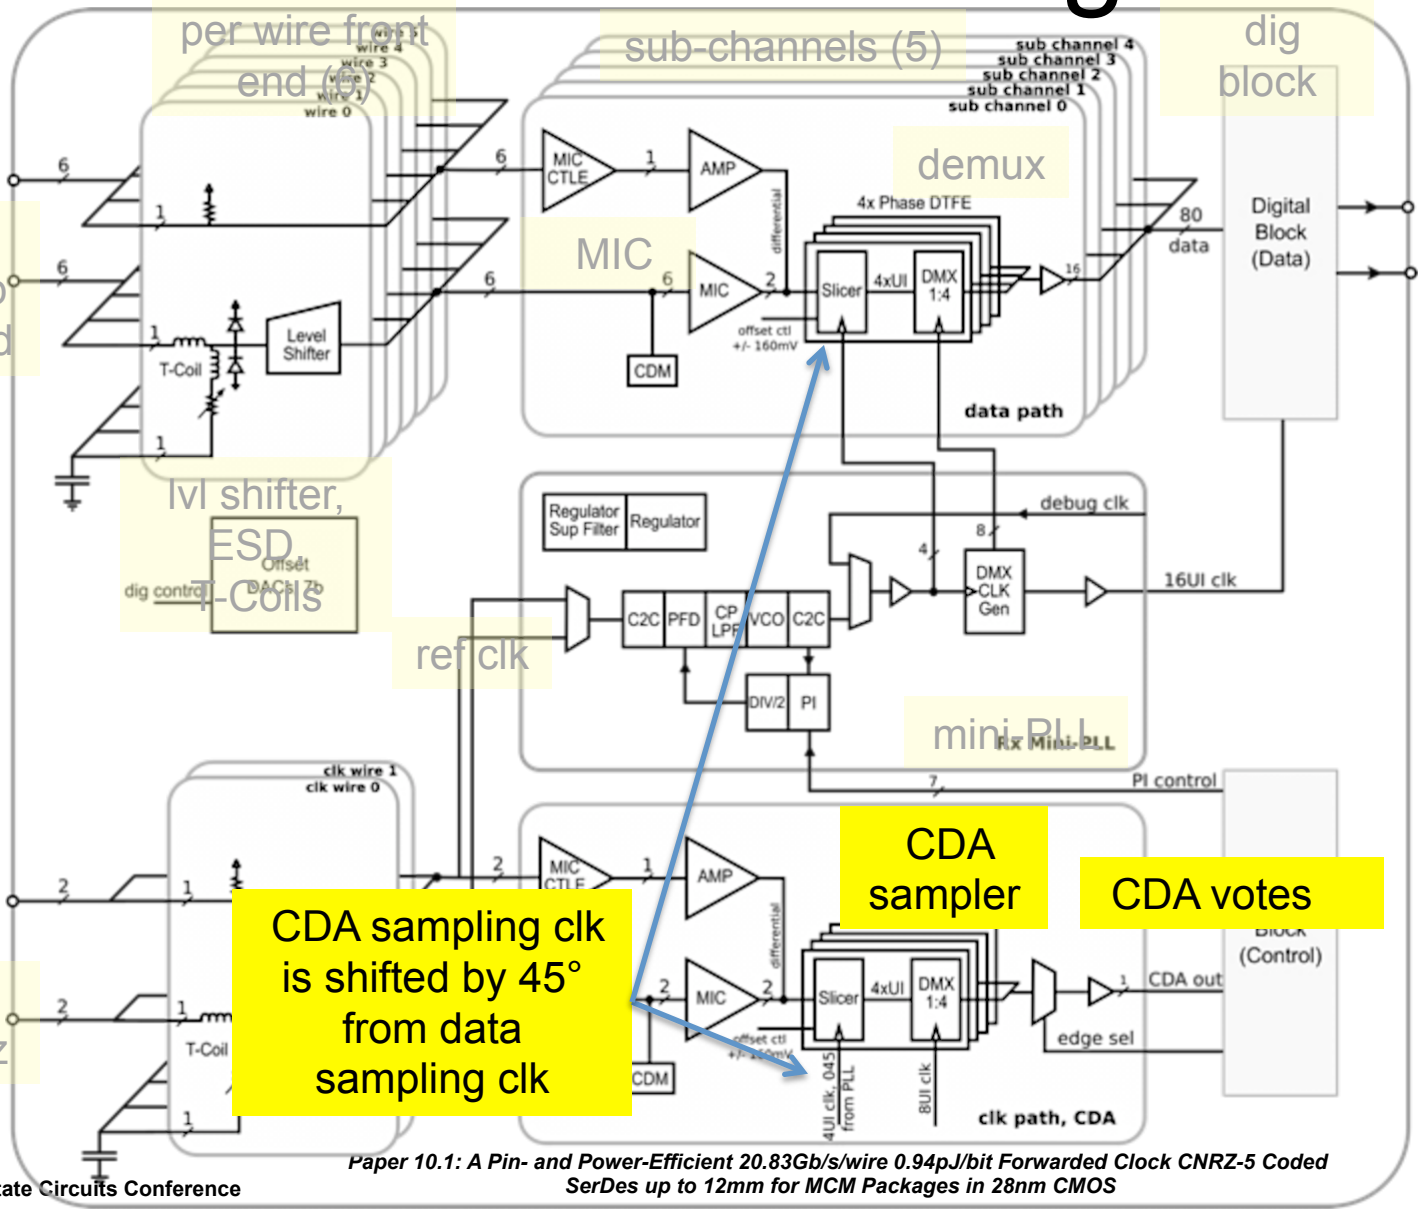

# Receiver Block Diagrams

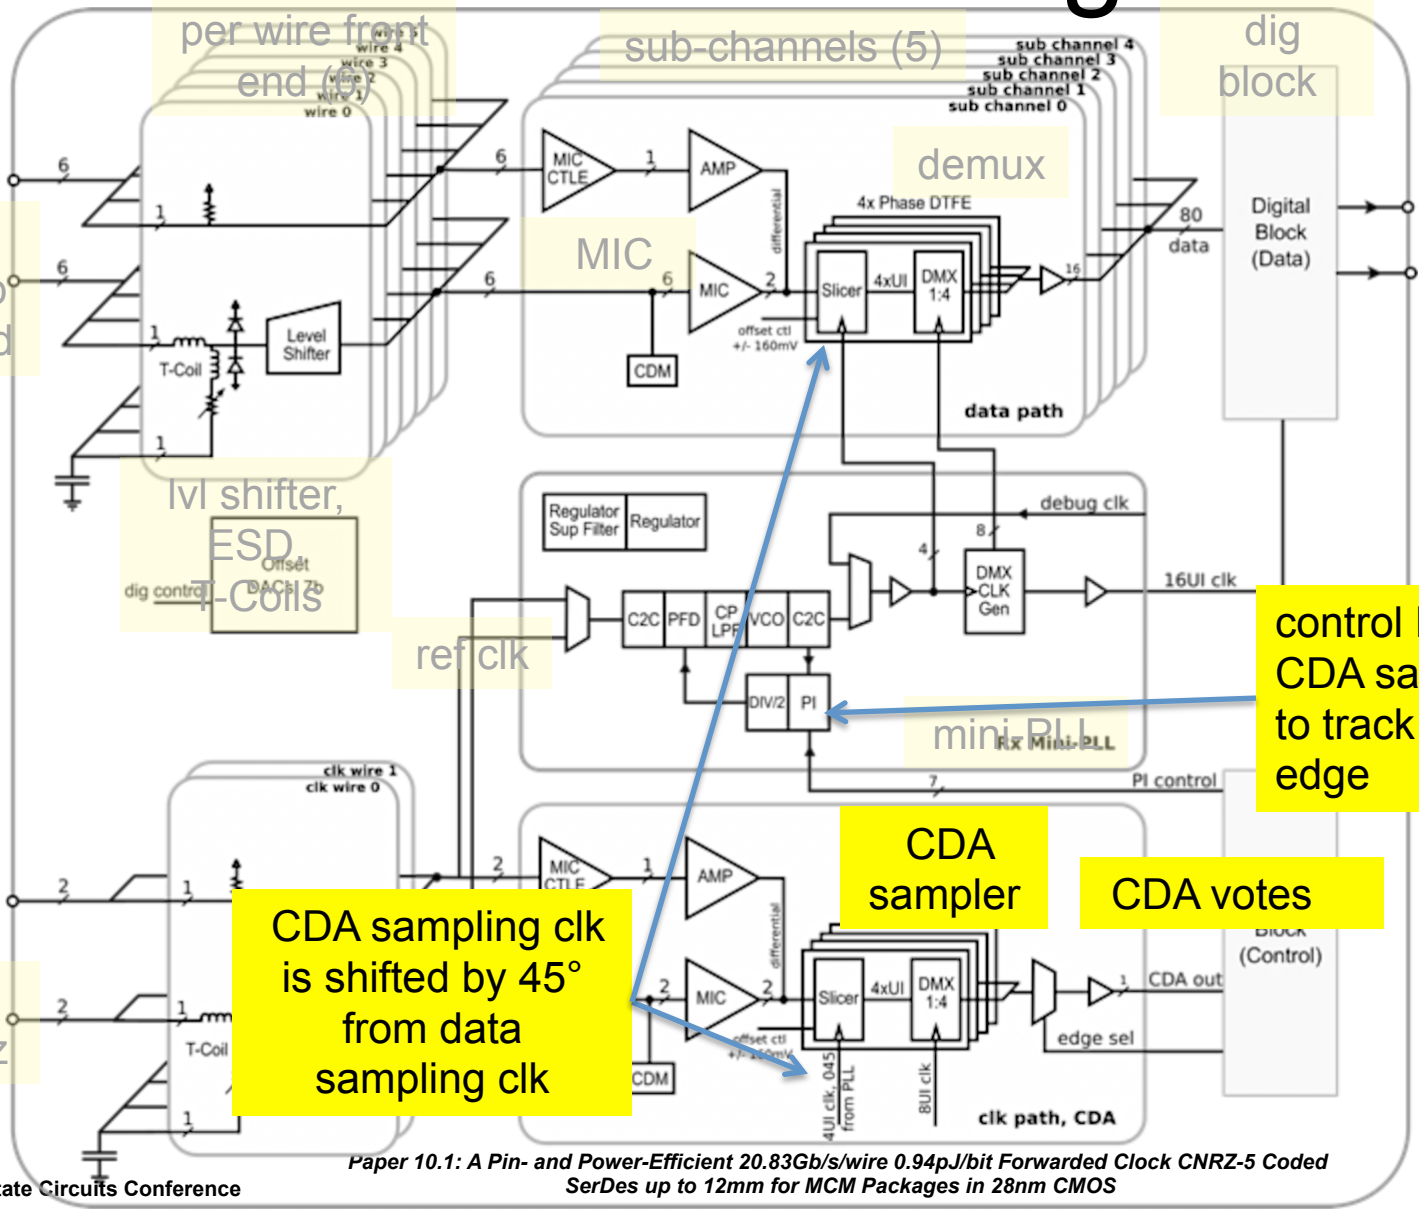

# Receiver Block Diagrams

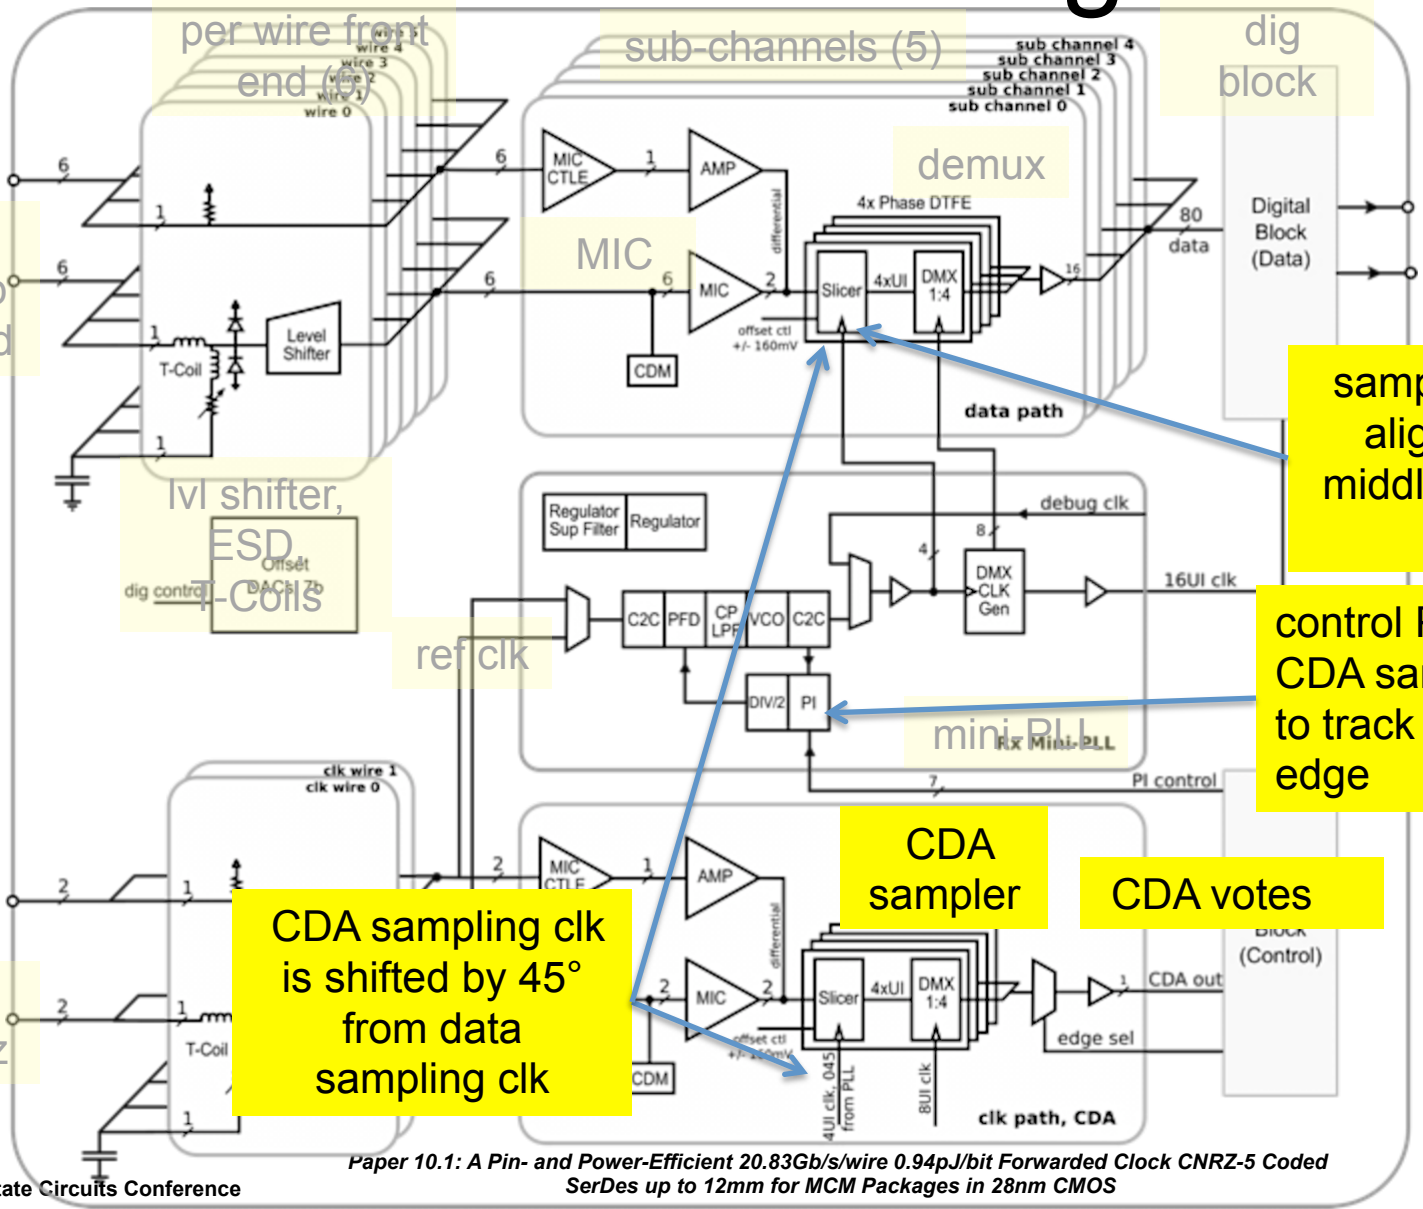

# Outline

- Introduction and motivation
- Signaling
- Macro architecture
  - Common block
  - Tx
  - Rx
- System Implementation
- Results
- Conclusion

# Bump Map and Technology

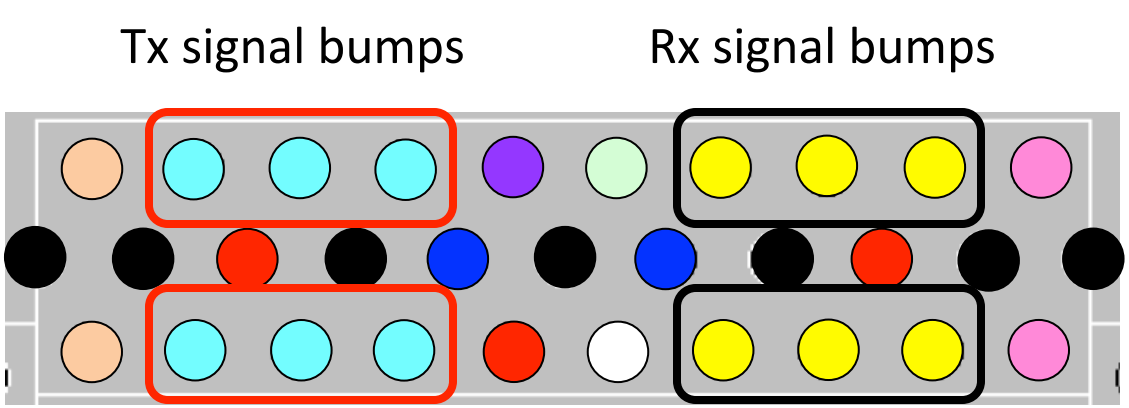


**Process**:

- TSMC 28nm HPM
- Metal stack is 9 layer: 6X2R + 1 UTALRDL
- DGO process (dual gate oxide, 1.0V and 1.8V devices).
- Devices used: HVT, LVT and nominal VT (3 types)

# Bump Map and Technology

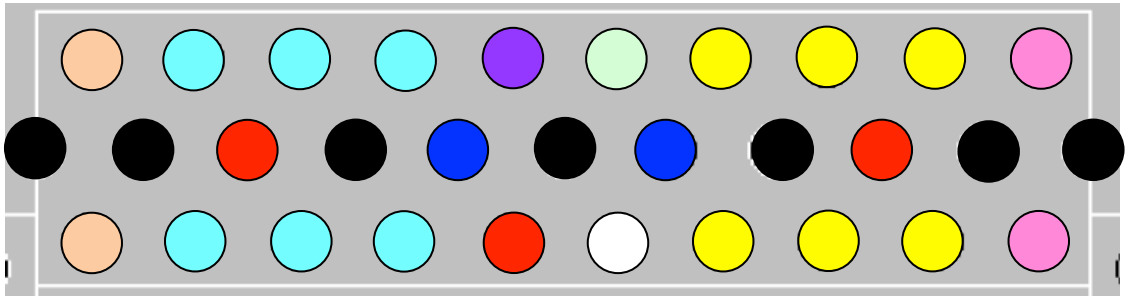

- vss – ground
- vdda - Rx/PLL analog power (1.0 V)
- vddh – PLL power supply (1.5V-1.8V)
- vddd – macro digital power supply (1.0V)
- Anabus signal
- Tx FCLK
- Tx signals
- Rx FCLK
- Rx signals
- Unused

**Process**:

- TSMC 28nm HPM
- Metal stack is 9 layer: 6X2R + 1 UTALRDL
- DGO process (dual gate oxide, 1.0V and 1.8V devices).
- Devices used: HVT, LVT and nominal VT (3 types)

# Testchip

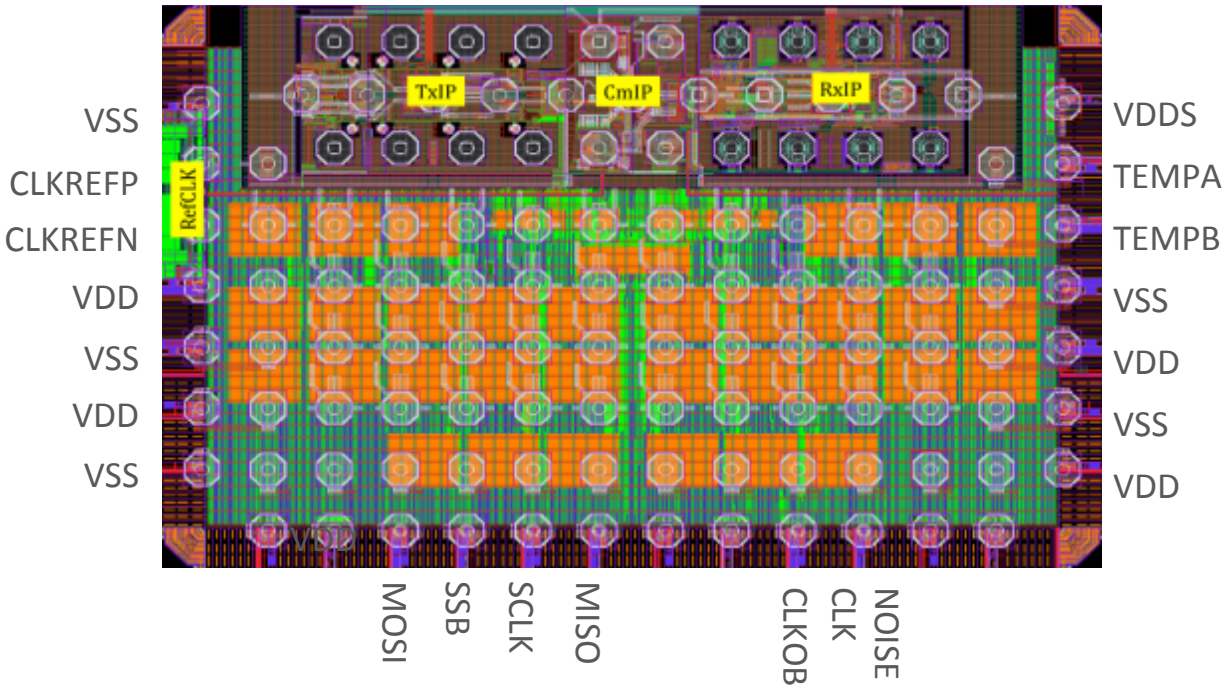


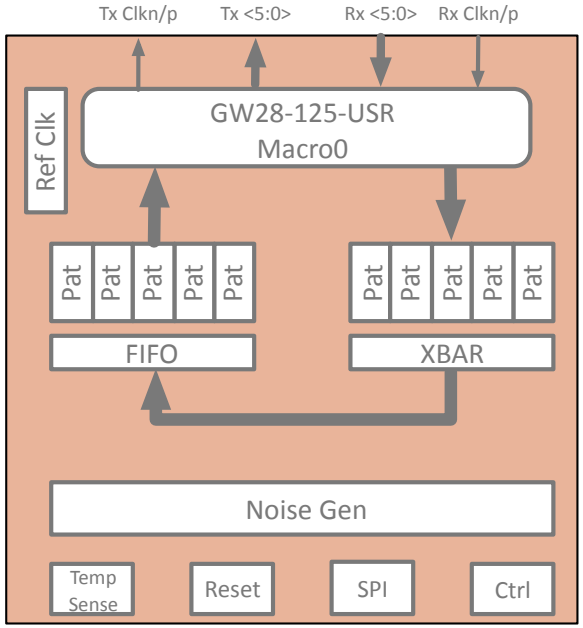


- Architecture:
  - One instance of CNRZ-5 IP (Tx + Rx)
  - One common block
  - 62.5 Gb/s and 125 Gb/s over six wires
- Die size (actual without scribe) 2138.4µm x 1386.9µm (2.96sqmm)

# Chip Micrograph

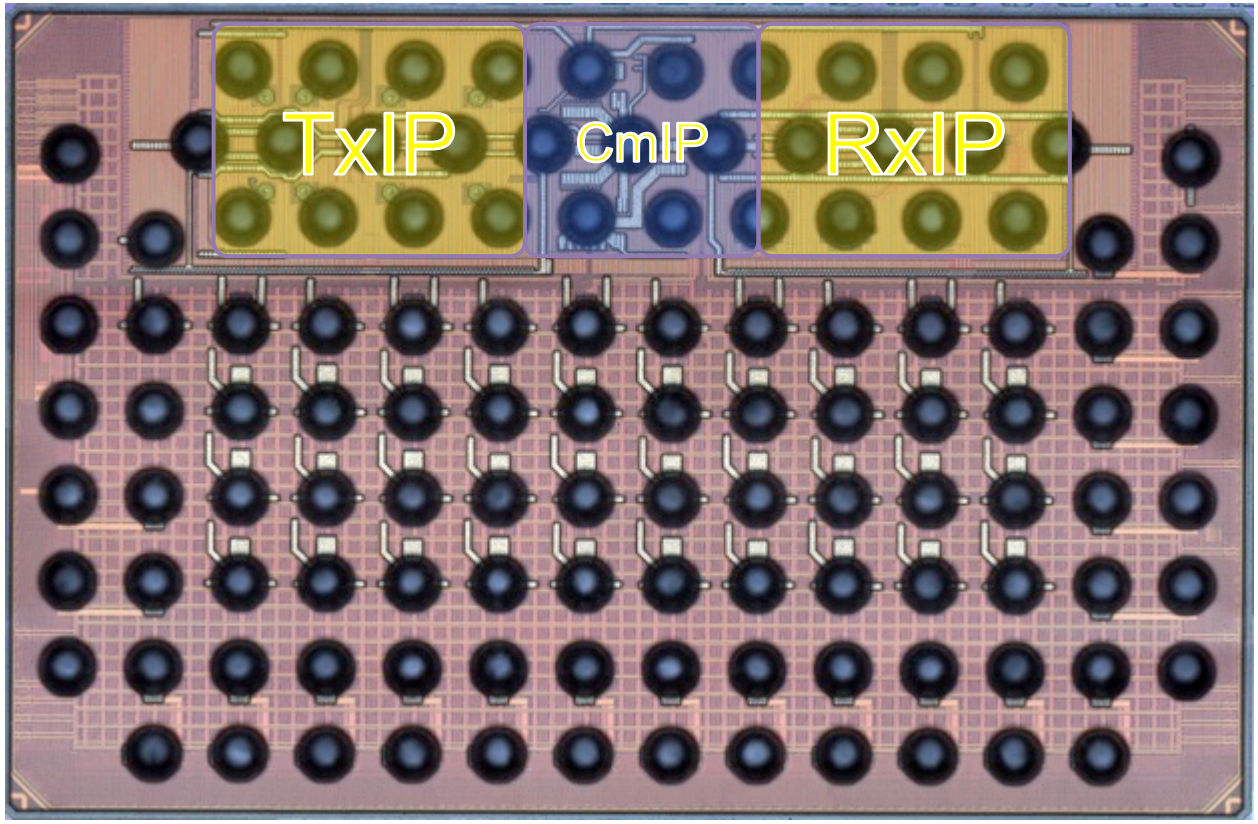

# Features

| | |
|---|---|
| Technology | 28nm CMOS HPM, VDD=1.0V, 9M, DGO |
| MCM Channels | Losses (s21) of ~0.6dB, ~1.25dB and ~2.5dB |
| Data Rate | 10.44-20.83 Gb/s/wire (12.5-25 Gbaud) |
| Power and Energy Efficiency | 117.5 mW at 125 Gbps |
| BER | < 1e-15 at 25 Gbaud |
| Testability | **IP:** Internal loopback; Rx Eyescope; PRBS31 Pattgen and Verification; Analog test bus.<br>**Testchip:** Pattern generators; Noise Generators. |

# Demo MCM

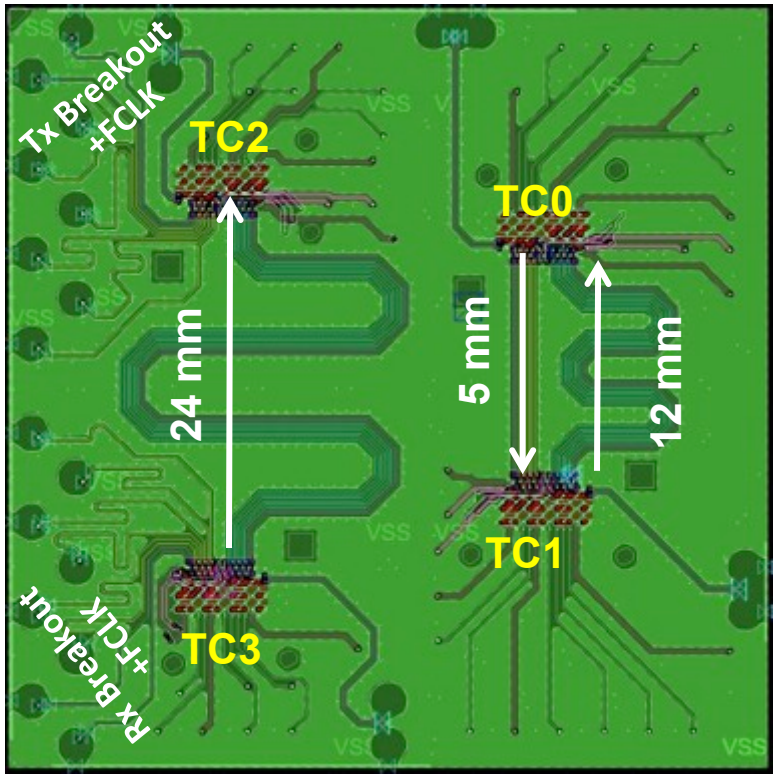


- Organic GX13 substrate, 222 stack up.
- 19mm square, 18 x 18 BGA, 1mm pitch.
- Populate up to 4 GW die interconnected in pairs.
- Provides up to 3 channels of 5mm, 12mm and 24mm.
- Channel trace length matched to ~1µm.
- Channel losses (s21) of ~0.6dB, ~1.25dB and ~2.5dB.
- CNRZ-5 channels, 50Ohms:
  - Trace width = 19.5µm
  - Space = 55.5µm
- Two die provide break out to Rx and Tx 6-wire interfaces plus FCLK respectively.

# Evaluation Board

H+S Connectors.
Type H&S MXP

MCM is solder ball assembled onto a host PCB.
Solder ball array 1mm pitch fully populated.
Organic substrate.

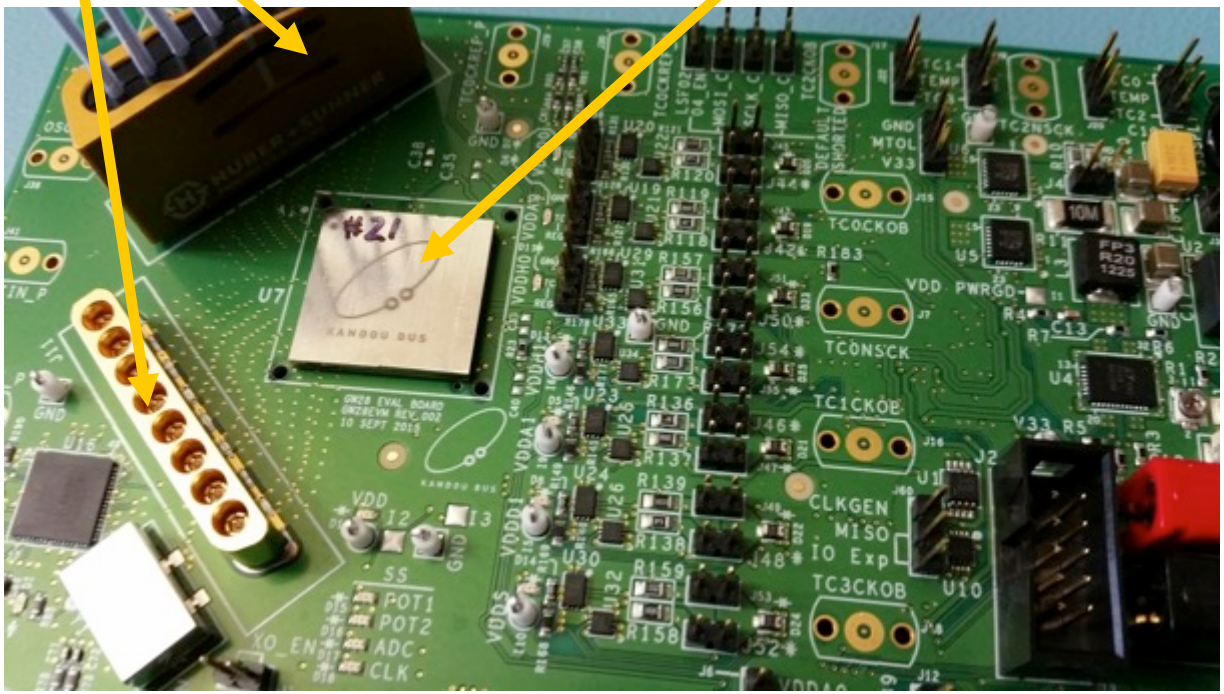


- MCM body size 19mm, 18 x 18 BGA
- 25 Gb/s HS signals off-MCM to PCB and H +S connectors from TC2 and TC3
- Single PCB design capable of accepting a socket and soldered MCMs

# Outline

- Introduction and motivation
- Signaling
- Macro architecture
  - Common block
  - Tx
  - Rx
- System Implementation
- Results
- Conclusion

# Results

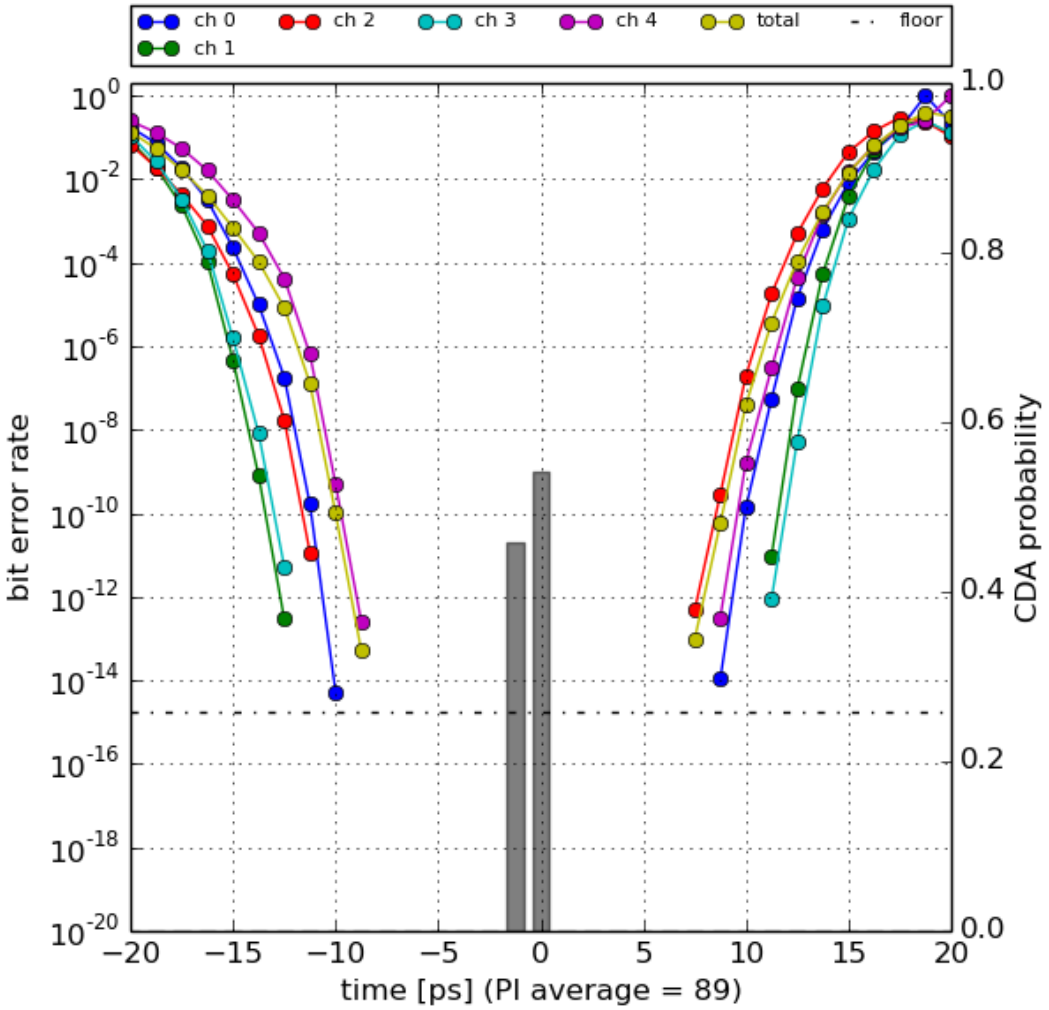


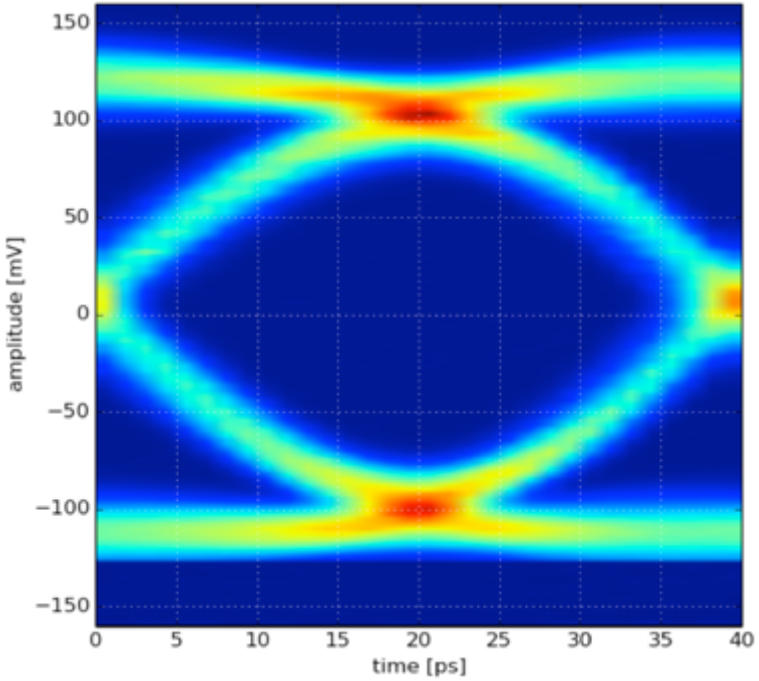

# Measured Power at 125 Gb/s

| | | V | mA | mW | | |
|---|---|---|---|---|---|---|
| CMIP | VDDA | 0.925 | 2.253 | **2.084** | *1.77%* | *8.48%* |
| | VDDH | 1.400 | 5.616 | **7.862** | *6.69%* | |
| | VDDD | 0.800 | 0.032 | **0.026** | *0.02%* | |
| TXIP | VDDA | 0.925 | 54.927 | **50.807** | *43.20%* | *54.17%* |
| | VDDH | 1.400 | 3.996 | **5.594** | *4.76%* | |
| | VDDD | 0.800 | 9.125 | **7.300** | *6.21%* | |
| RXIP | VDDA | 0.925 | 31.152 | **28.816** | *24.50%* | *37.35%* |
| | VDDH | 1.400 | 8.177 | **11.448** | *9.73%* | |
| | VDDD | 0.800 | 4.584 | **3.667** | *3.12%* | |
| | | | $P_{total}$ [mW] | **117.605** | | |
| | | | Rate [Gb/s] | 125 | | |
| | | | $E_{bit}$ [pJ/b] | **0.941** | | |

# Results

0.94 pJ/b

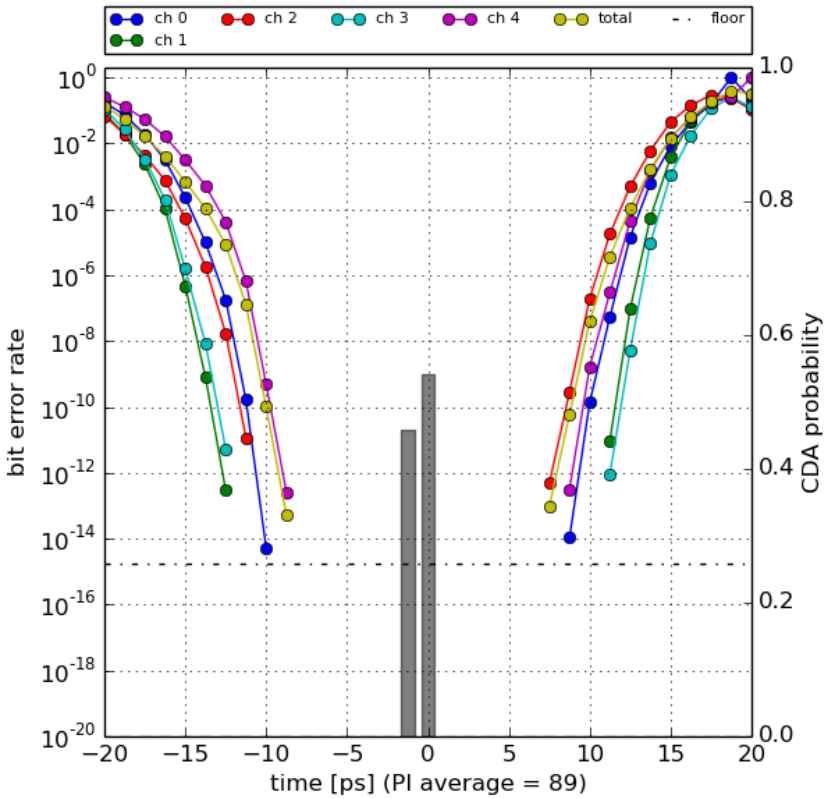


Vdda = 0.925V
Vddd = 0.8V
Vddh = 1.4V

1.04 pJ/b

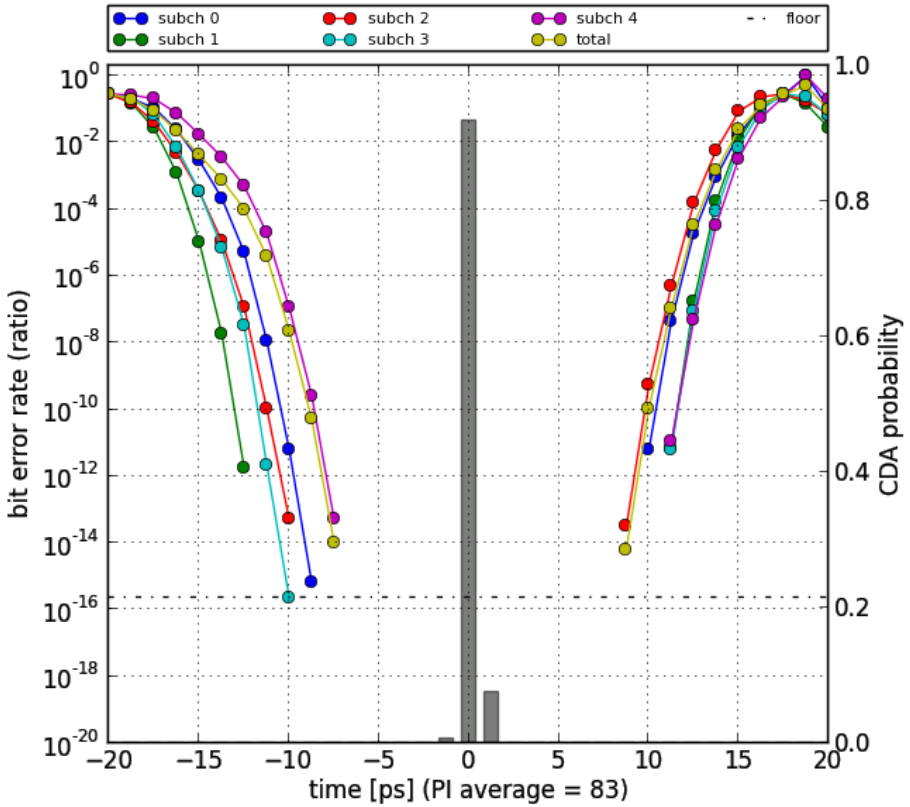


Vdda = 1V
Vddd = 1V
Vddh = 1.5V

# Results (Continued)

- Target BER of 1e-15 achieved at 25 Gbaud
  - Specified 12mm MCM channel across power supply corners and temperature stress
- Target BER of 1e-15 achieved at half data rate 12.5 Gbaud on 24 mm channel
- Target BER achieved under stress test (±5% supply tolerance, temp range from 0 to 100 C)
- Target BER achieved at 30 Gbaud on 12mm channel under nominal power supplies and temperature

# Conclusion

- Demonstrated an implementation of a SerDes based on CNRZ-5 coding, transmitting 5 bits on 6 correlated wires in every UI
- Coding has built-in resilience to various types of noise which helps lower the power consumption at high speeds
- Implementation uses a forwarded clock architecture and transmits up to 125 Gb/s on 6 wires at 0.94 pJ/bit

# References


[1] Kim et al., “A 10 Gb/s compact low-power serial I/O with DFE-IIR equalization in 65 nm CMOS”, IEEE Journal of Solid State Circuits, vol. 44, pp. 3526-3538, 2009.

[2] Poulton et al., “A 0.54pJ/b 20 Gb/s Ground-Referenced Single-Ended Short-Haul Serial Link in 28nm CMOS for Advanced Packaging Applications,” IEEE Journal of Solid State Circuits, vol. 48, pp. 3207-3218, 2013.

[3] Dickson et al., “A 1.4 pJ/bit, power scalable 16x12 Gb/s source-synchronous I/O with DFE receiver in 32 nm SOI CMOS technology,” in *Proc. IEEE Custom Integrated Circuits Conf.,* 2014, pp. 10-5.

[4] Hormati et al., “Method and Apparatus for Low-Power Chip-to-Chip Communications with Constrained ISI-Ratio”, U.S. Patent 9,100,232.

[5] A. Hormati and A. Shokrollahi, “ISI tolerant signaling: a comparative study of PAM4 and ENRZ,” DesignCon 2016.

# Thank you