# Selection of hardware platform for CBM Common Readout Interface

W.M.Zabołotny[a], G.H.Kasprowicz[a], A.P.Byszuk[a],
D.Emschermann[b], M.Gumiński[a], K.T.Poźniak[a], R.Romaniuk[a]

[a]Institute of Electronic Systems
Warsaw University of Technology
[b]GSI-Helmholtzzentrum für Schwerionenforschung GmbH

Presenter: Wojciech M. Zabołotny

# Original concept of the CBM readout

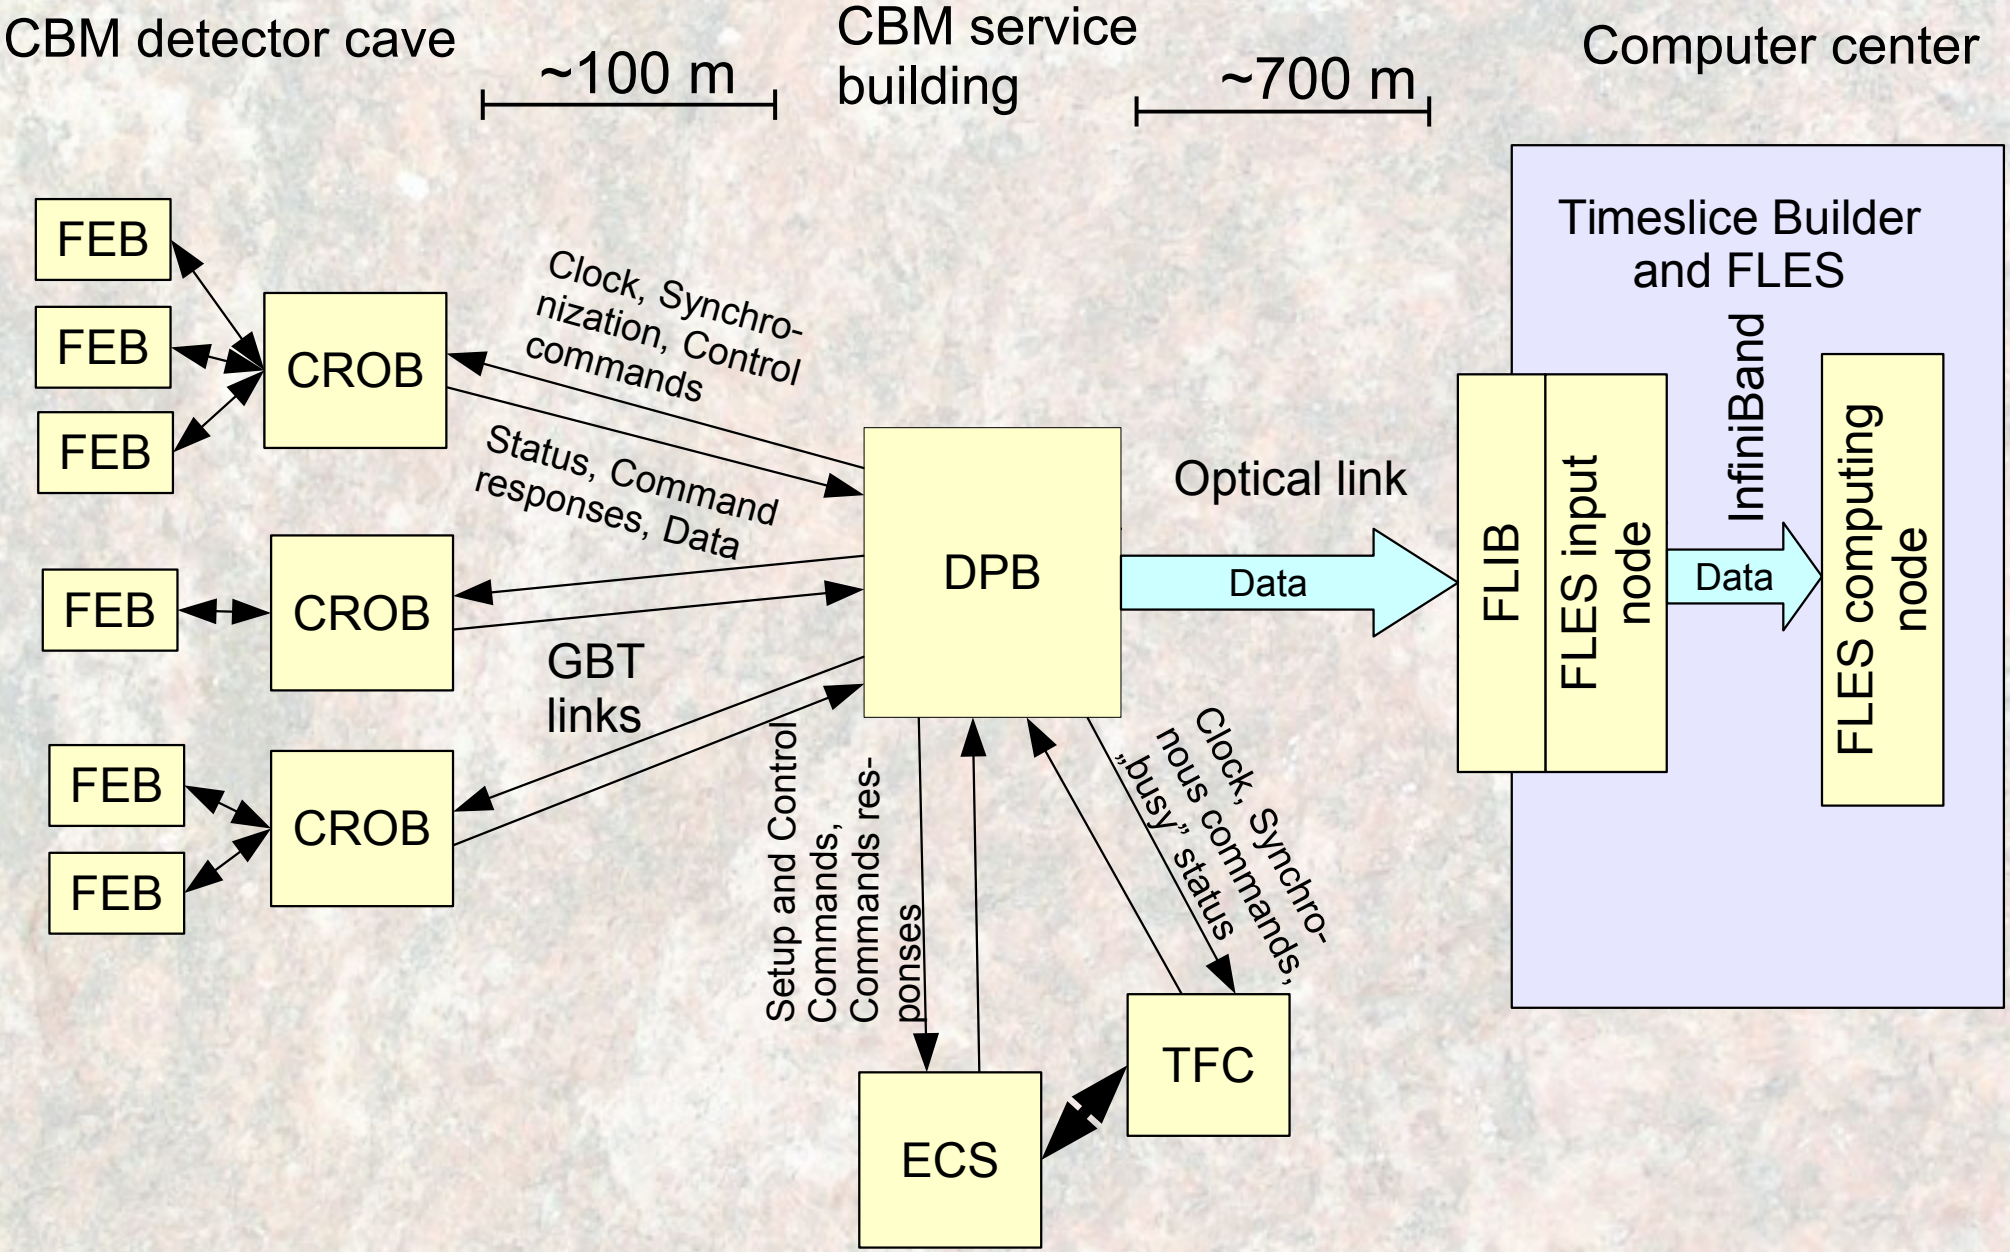

# STS-XYTER tester

The DPB firmware blocks were tested in the STS-XYTER tester. The tester uses the black-box model of the GBTx and E-Links and IPbus-based readout with limited bandwidth.

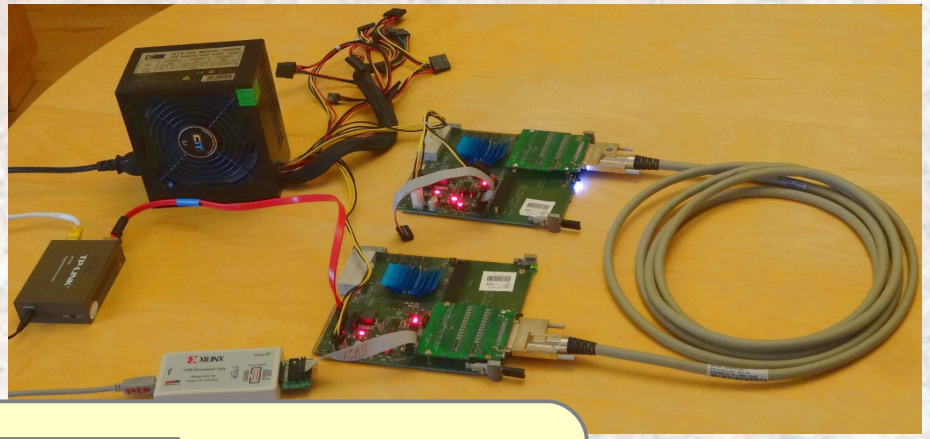

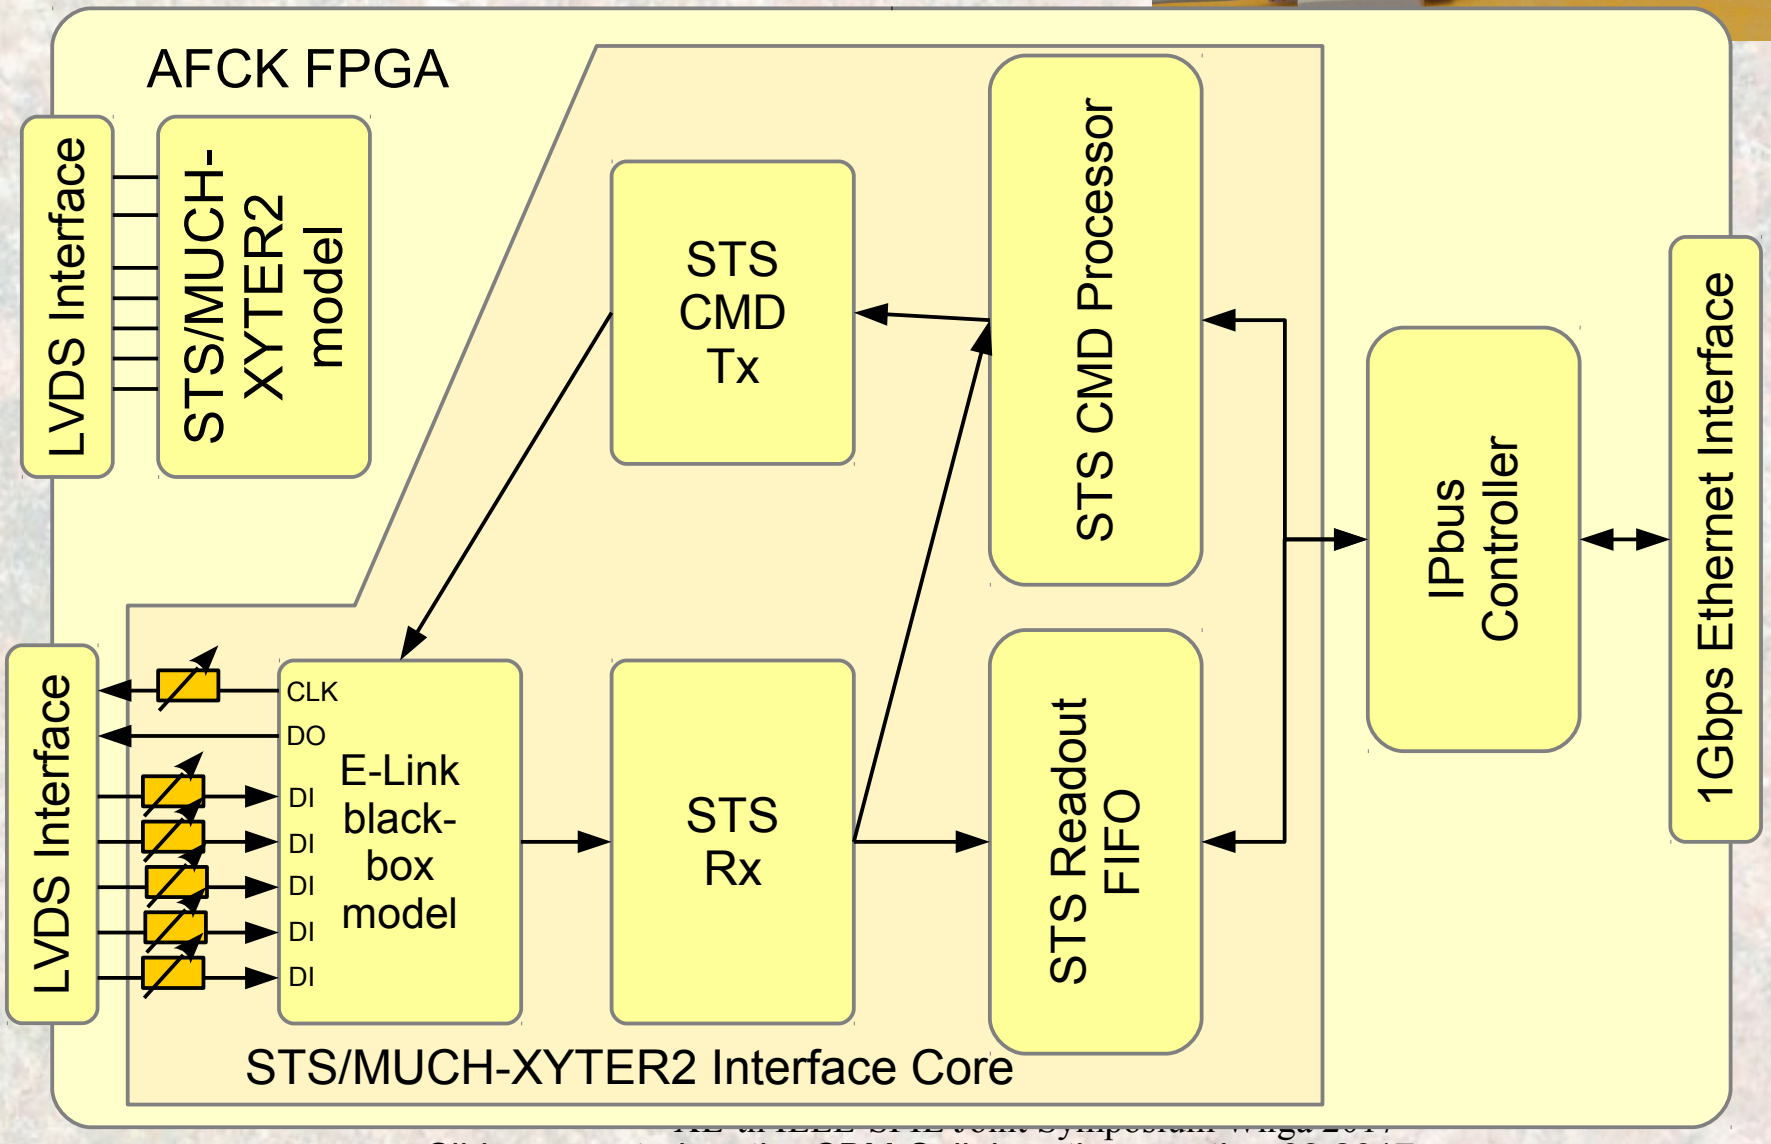




Slide presented on the CBM Collaboration meeting 03.2017

# Integration of the tester with GBT

- To test the operation of the GBTX ASIC in modes needed in the STS/MUCH readout (160 Mb/s „GBT frame” for downlink and 320 Mb/s „wide bus” for uplink), the dedicated AFCK-GBT-tester has been created.
- The tester firmware uses the GBT-FPGA IP Core from CERN (currently version 5.0.0)
- Testing using both – interactive mode and complex routines was possible due to the addition of Forth CPU – J1B with SwapForth to the firmware.

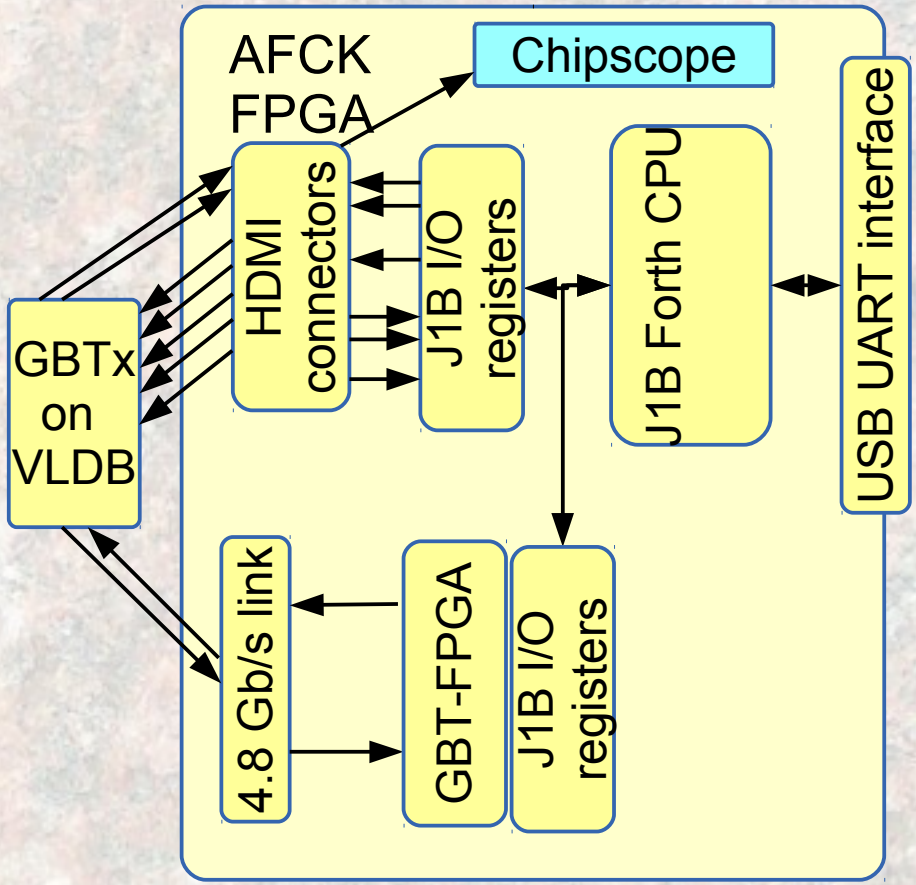


Slide presented on the CBM Collaboration meeting 03.2017

# Hardware used for GBT testing

- Test setup based on the AFCK board supplemented with the FM-S14 board and the „e-Link FMC”.
- Two HDMI cables with HDMI-micro to HDMI-mini adapters.
- Shifted clock routed via separate coaxial cable to the U.FL connector on the AFCK (required disabling of Si570 though!)

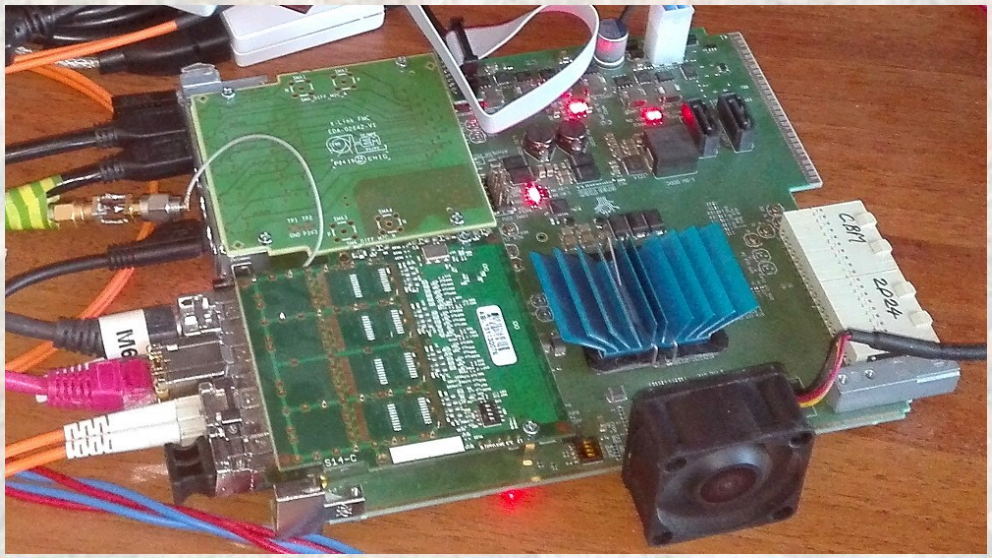

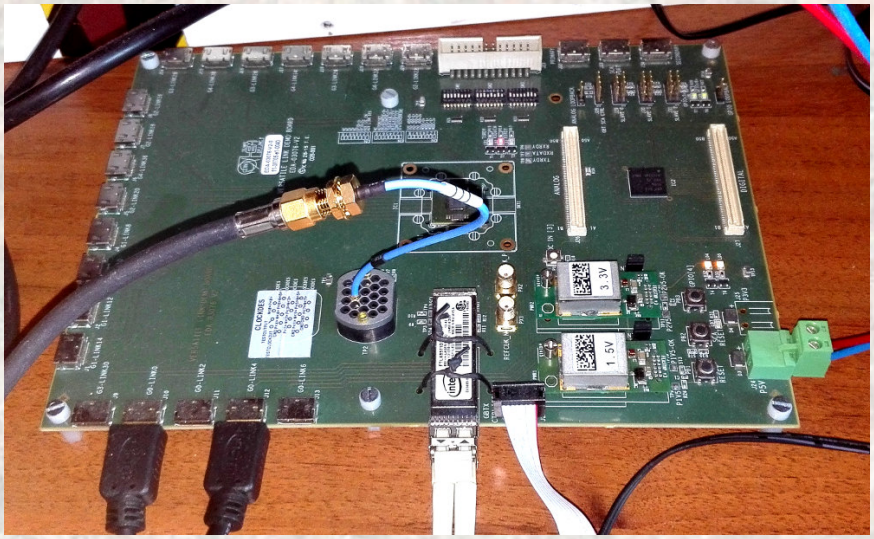

Slide presented on the CBM Collaboration meeting 03.2017

# Current state of the GBT-based STS-XYTER tester

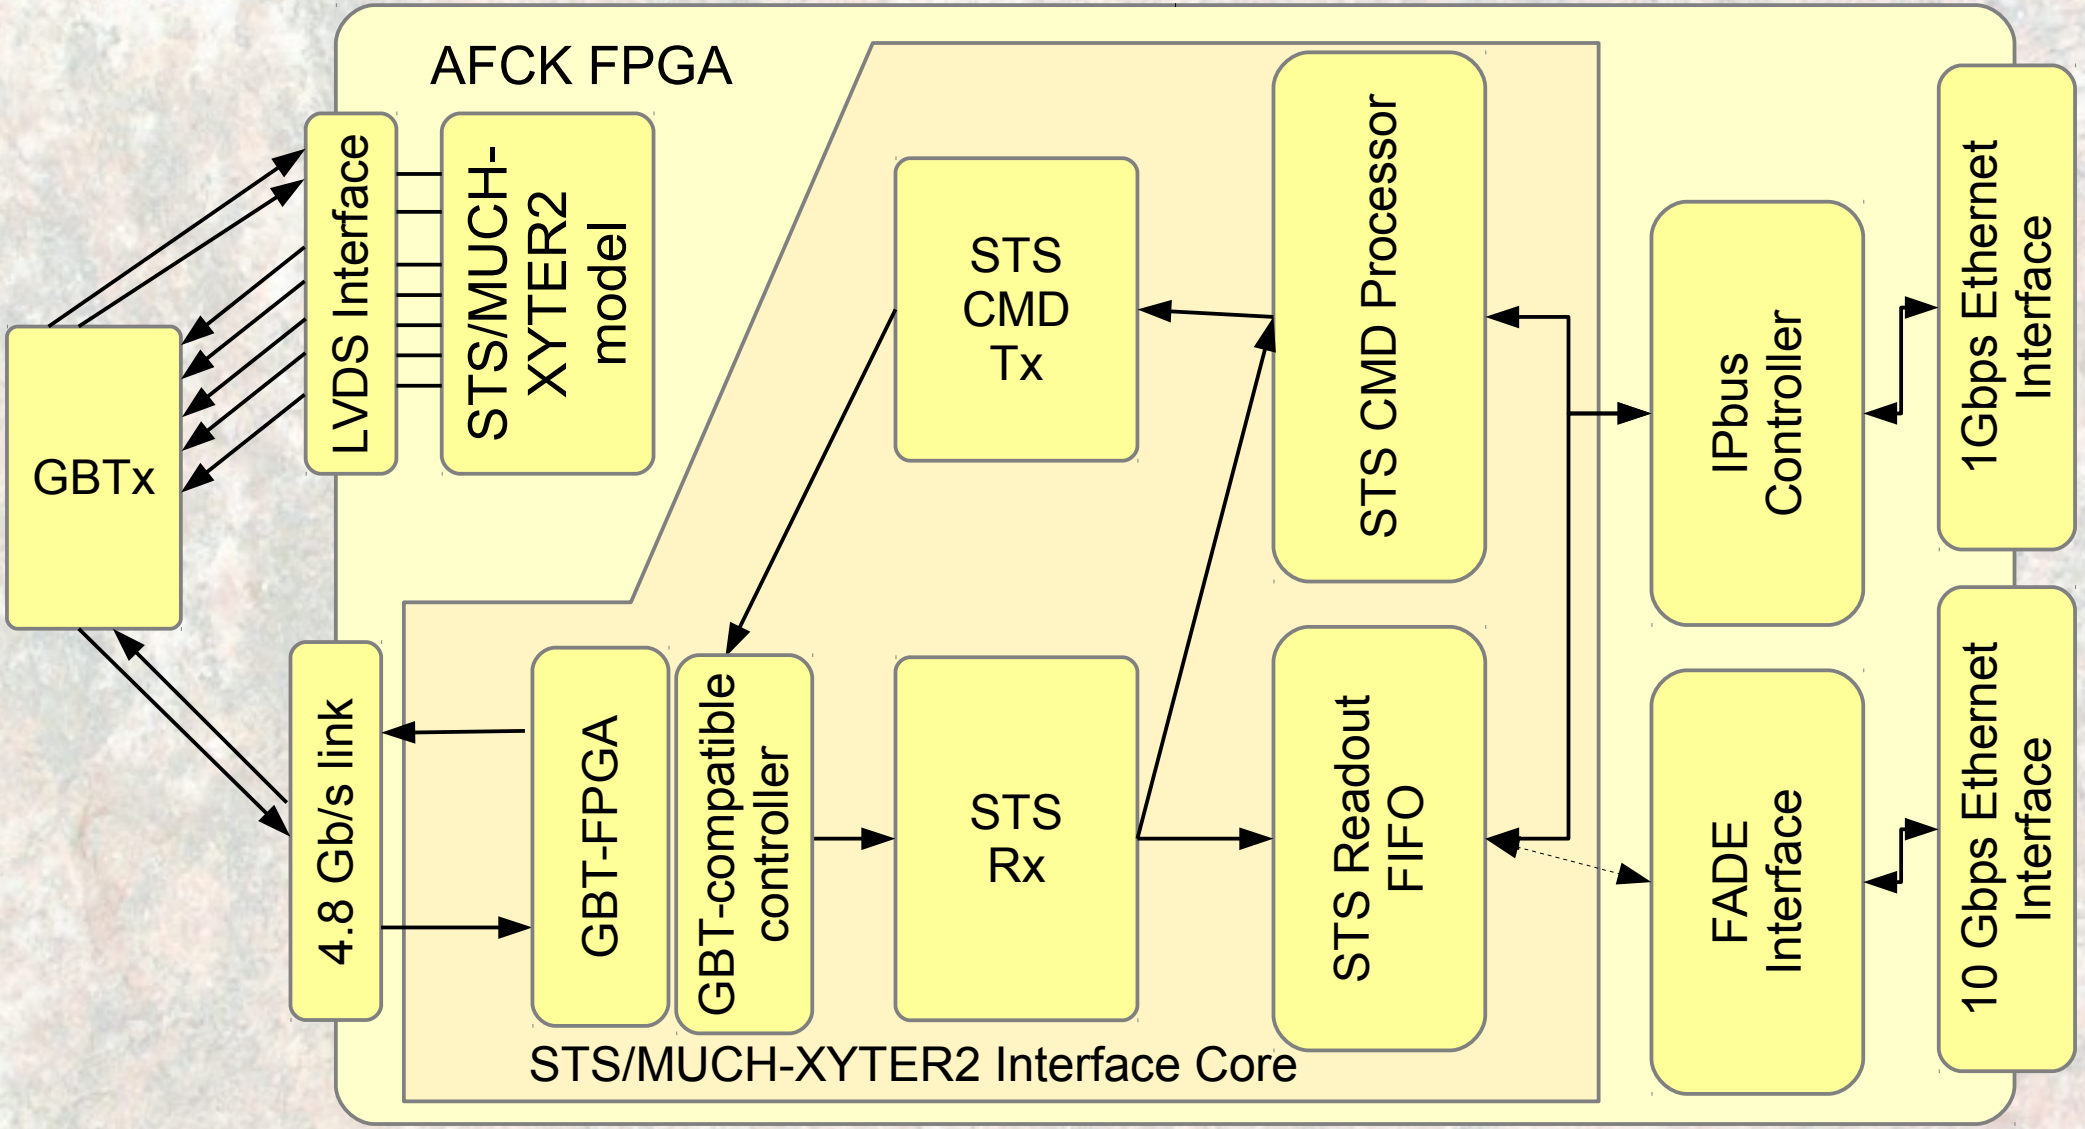

# Previous concept of the CBM readout

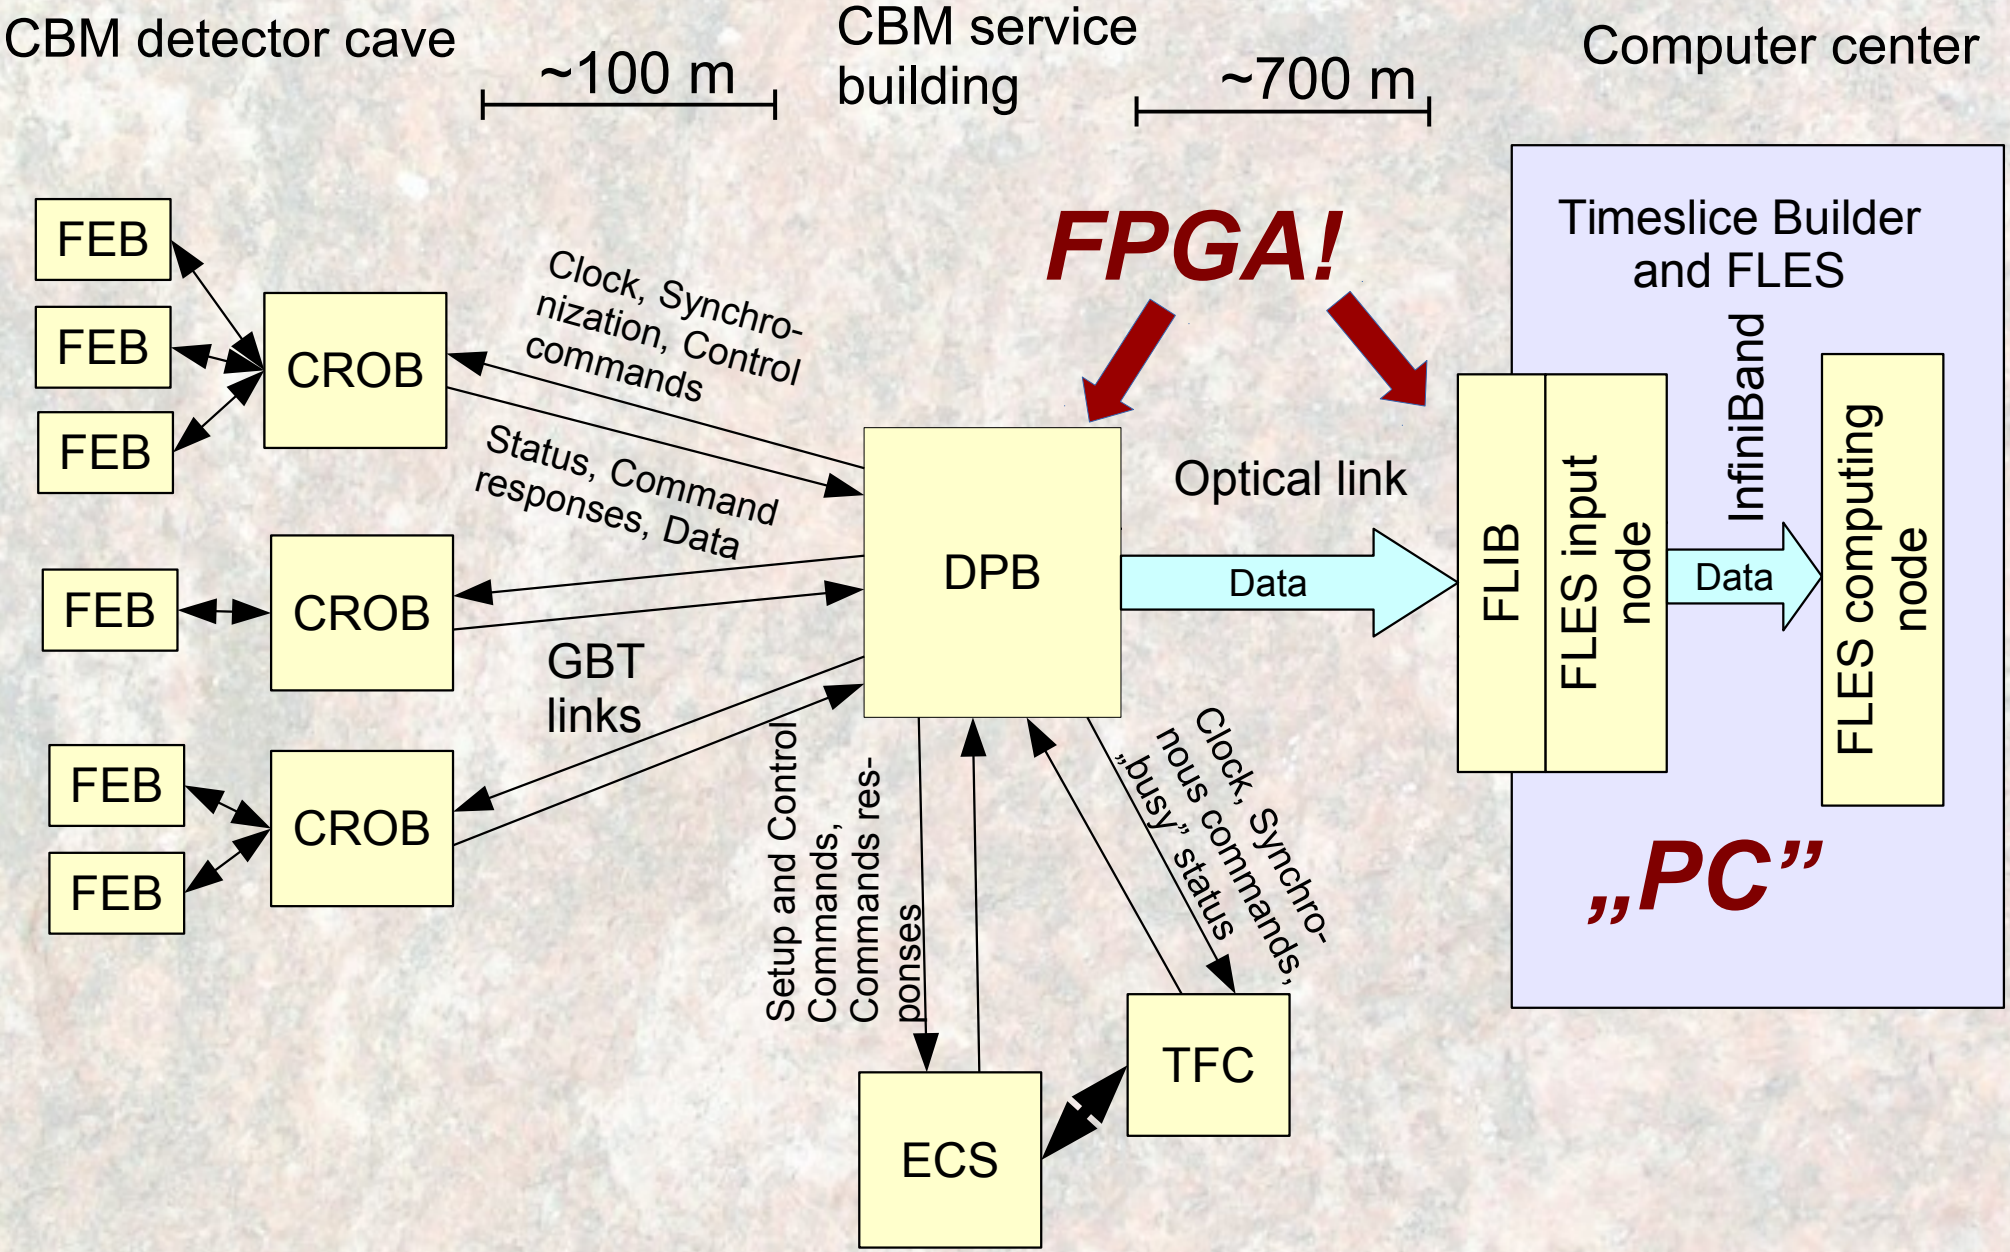

# New concept of the CBM readout

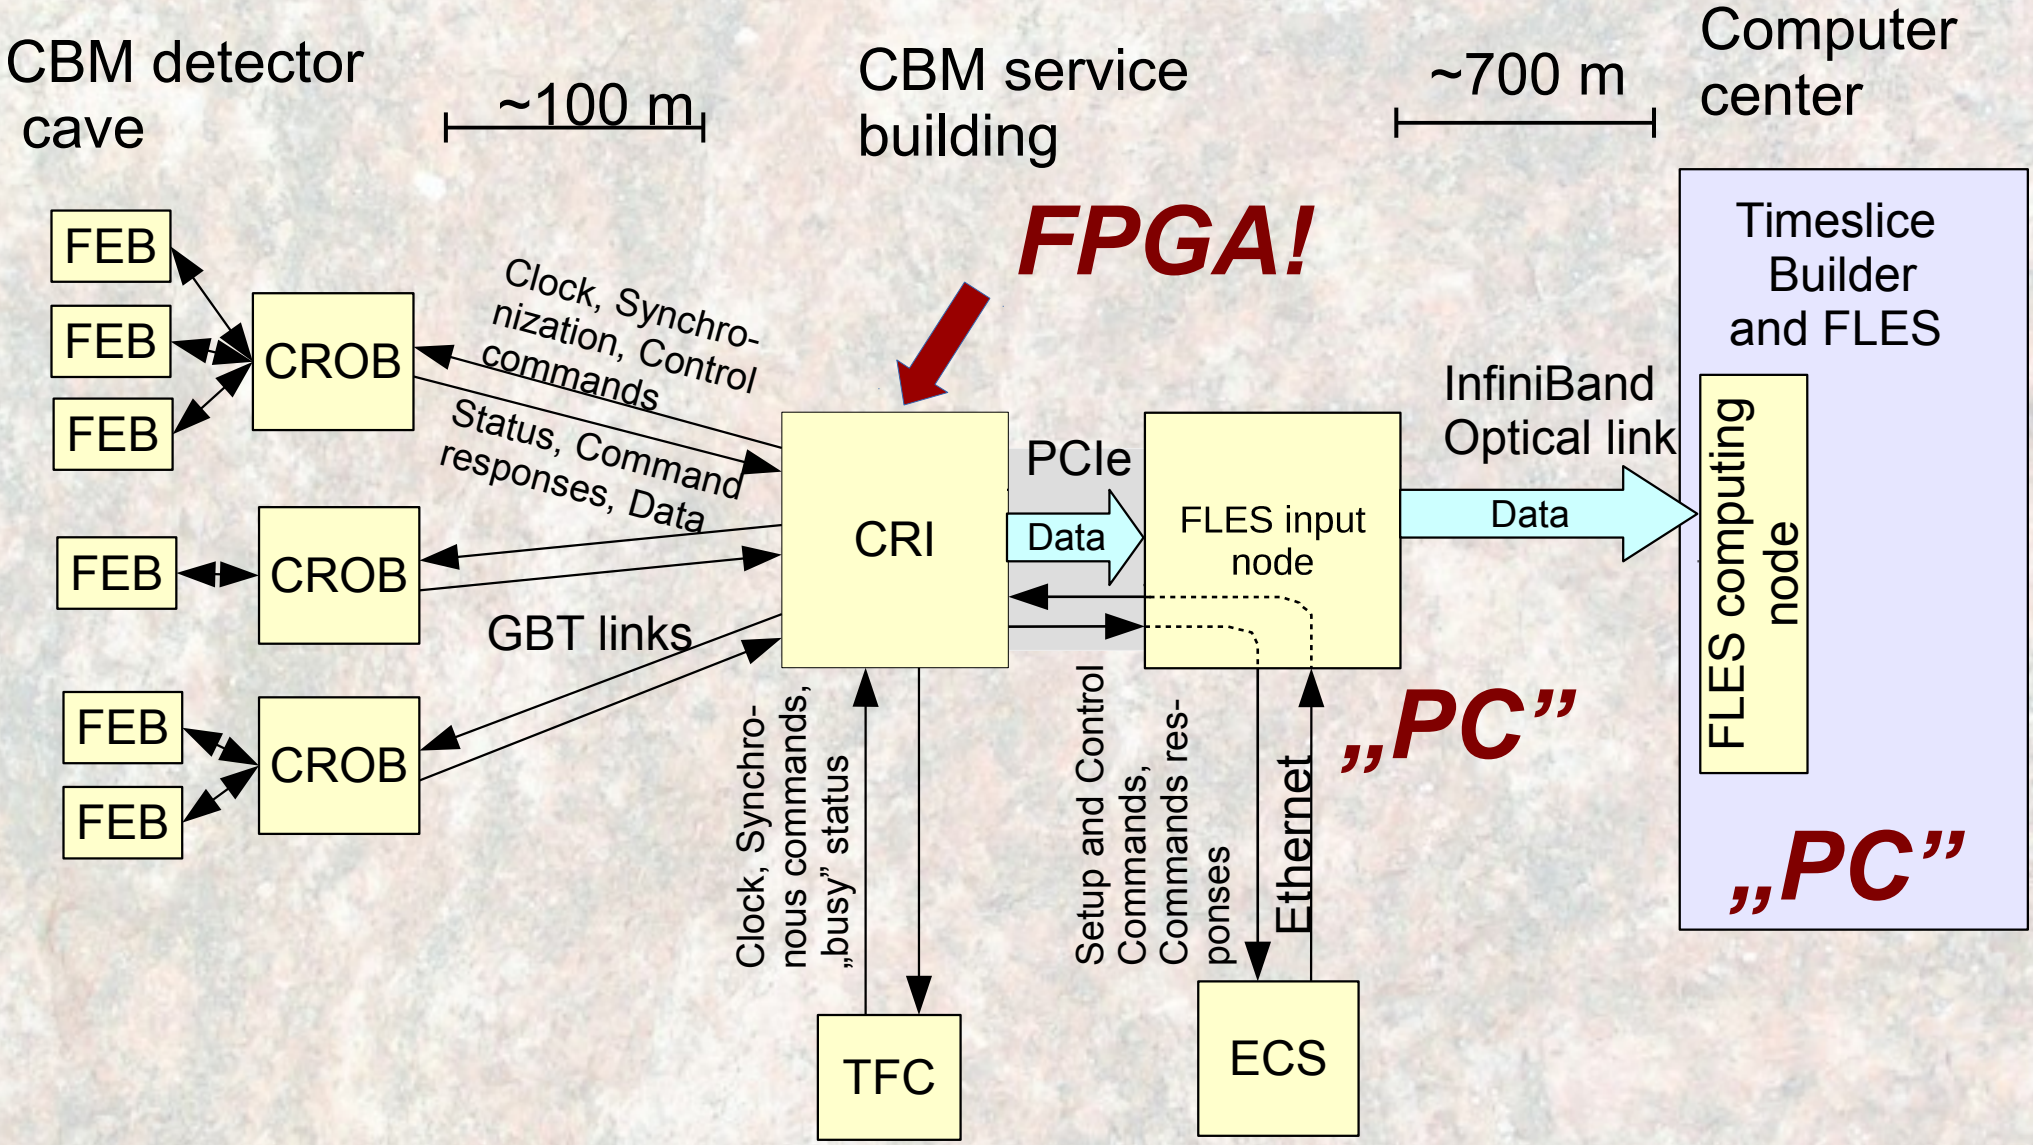

# Requirements to the CRI board

- Handling of the required number of GBT links to receive the data from the CROBs
- Providing the sufficient PCIe bandwidth to send the received data to the FLES input node
- Providing the sufficient amount of logical resources to concentrate the data

# CROB board

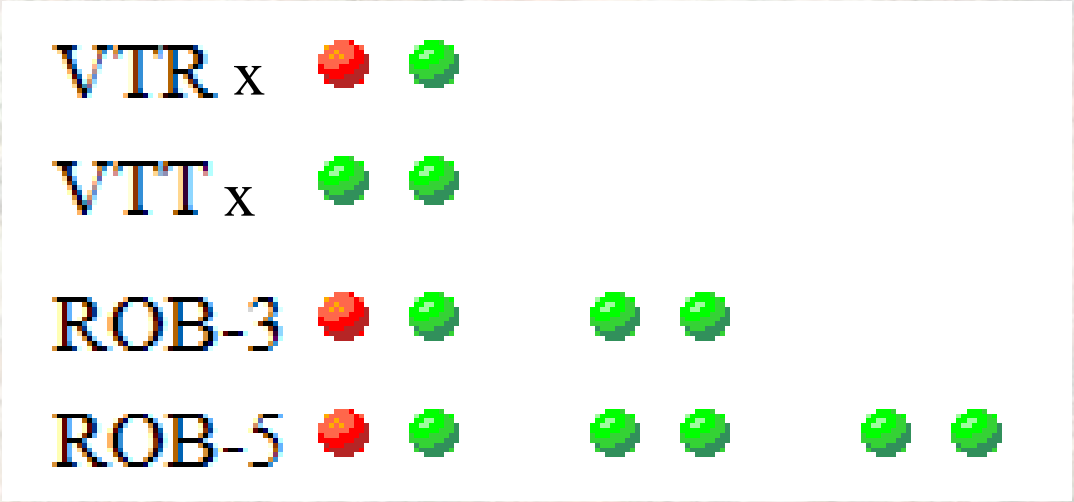


(STS)

(TRD)

ROB-3

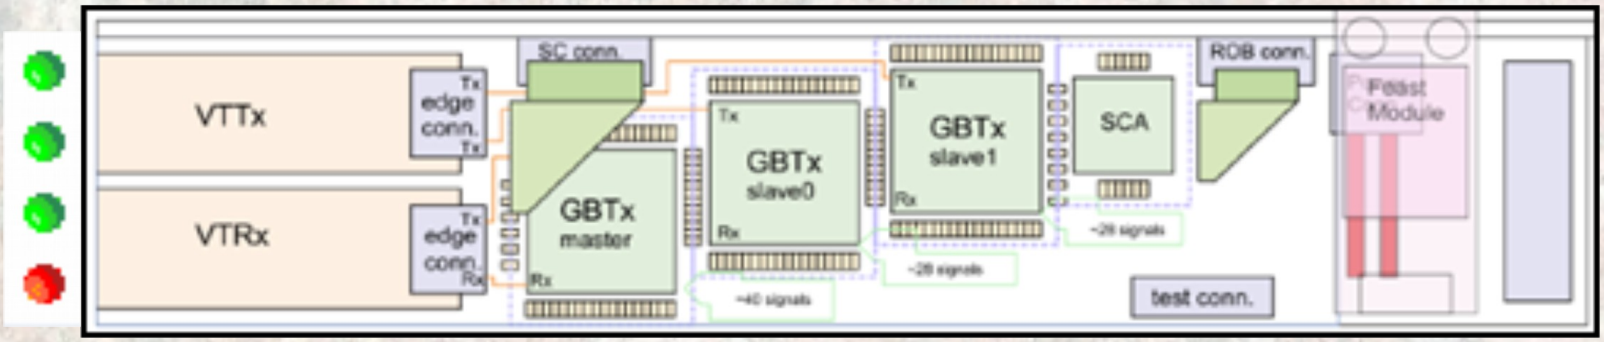


A ROB-3 will interface to 3 MGTs on the DPB.

**Based on the slide presented by David Emschermann** @ CBM ASIC workshop in Darjeeling 22.02.2017

# Available optical transceivers

Source: https://www.finisar.com/optical-transceivers/ftl410qd2c

QSFP+: Finisar FTL410QD2C 22 with approximate price of $300 for 4 duplex channels.

Source: Foxconn

12-channel - Foxconn MiniPOD 23 modules with the approximate price of $360 for simplex channels (either Tx or Rx).

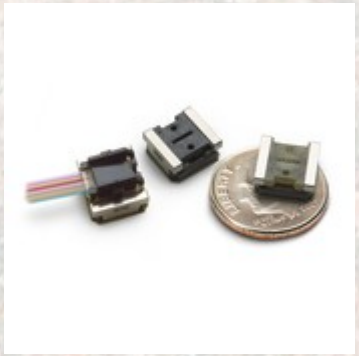

Source: Foxconn

12-channel - Foxconn MicroPOD 24 modules with the approximate price of $130 for 12 Tx channels and $280 for 12 Rx channels.

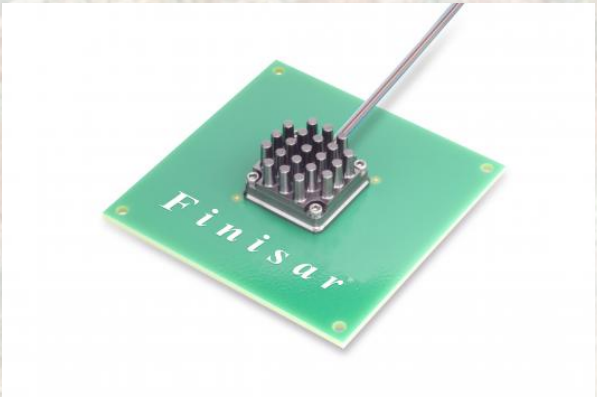


Source: Foxconn

24-channel - Finisar FBOTD10SL1C00 25 with the approximate price of $860 for 24 duplex channels. The RX only FBRTP10SL1C00 modules are also available.

# Optical input possibilities

- It seems that the best cost per channel may be achieved with the 12-channel MicroPOD modules. The smallest configurations with a reasonable utilization of optical channels are the following:
  - With two 12-channel receivers and one 12-channel transmitter, we can connect 8 CROBs to a single CRI. In this case, we waste 4 transmit channels (may be used as spares).
  - With three 12-channel receivers and one 12-channel transmitter, we can connect 12 CROBs to a single CRI. In this case, we fully utilize optical channels.

# Input data bandwidth analysis

- The number of handled CROBs defines also the required PCIe bandwidth
  - Worst case analysis: 4.48 Gbps/GBT link, 13.44 Gbps/CROB, **108 Gbps for 8 CROBs, 162 Gbps for 12 CROBs**
  - Optimistic analysis: 9.41 Mhit/s/SMX2 chip, 20 bits of non-aggregated data/hit, 40 SMX2 chips/CROB, 7.53 Gbps/CROB, **60.22 Gb/s for 8 CROBs, 90.34 Gb/s for 12 CROBs.**

# PCIe bandwidth analysis

- PCIe 3.0 or 3.1: 7.877 Gbps/lane, UltraScale+ devices may use up to 16 lanes.
  - For pessimistic assessment:
    - we need 16 lanes for 8 CROBs (bandwidth 126 Gb/s - margin of 14%)
    - we need two links each with 16 lanes for 12 CROBs (bandwidth of 252 Gb/s - margin of 36%).
  - For optimistic assessment:
    - a single link with 8 lanes should be sufficient for 8 CROBs (bandwidth of 63 Gb/s) but the margin of 4% is probably too low, and therefore the 16-lanes link will be needed (margin over 50%).
    - The same link will be needed in case of 12 CROBs (margin of 28%).
- The PCIe 4.0: 15.752 Gbps/lane, UltraScale+ devices may use up to 8 lanes.
  - For pessimistic assessment:
    - one 8-lanes link for 8 CROBs
    - two 8-lanes links for 12 CROBs
  - For optimistic assessment
    - one 8-lanes link both for 8 and 12 CROBs

# Required MGT resources

| #CROBs | #GBTs | #lanes PCIe 3 | | #lanes PCIe 4 | |
|---|---|---|---|---|---|
| | | opt. | pess. | opt. | pess. |
| 8 CROBs | 24 | 16 | 16 | 8 | 8 |
| 12 CROBs | 36 | 16 | 32 | 8 | 16 |

# Total number of required MGT resources (+1 for TFC)

| #CROBs | For PCIe 3 | | For PCIe 4 | |
|---|---|---|---|---|
| | opt. | pess. | opt. | pess. |
| 8 CROBs | 41 | 41 | 33 | 33 |
| 12 CROBs | 53 | 69 | 45 | 53 |

# Available FPGAs

| | KU11P | KU15P | | ZU11EG | ZU17EG | | ZU19EG | |
|---|---|---|---|---|---|---|---|---|
| System Logic Cells | 653,100 | 1,143,450 | | 653,100 | 926,124 | | 1,143,450 | |
| Block RAM Blocks | 600 | 984 | | 600 | 796 | | 984 | |
| UltraRAM Blocks | 80 | 128 | | 80 | 102 | | 128 | |
| Enclosure | FFVE 1517 | FFVE 1517 | FFVA 1760 | FFVC 1760 | FFVC 1760 | FFVD 1760 | FFVC 1760 | FFVD 1760 |
| Number of transceivers GTH+GTY | 32+20 | 32+24 | 44+32 | 32+16 | 32+16 | 44+28 | 32+16 | 44+28 |
| Total number of transceivers | 52 | 56 | 76 | 48 | 48 | 72 | 48 | 72 |
| Example price | ca. $3500 | ca. $4600 | ca. $7900 $5000? | ca. $3900 | ca. $4400 | ca. $4900 | ca. $4900 | ca. $5400 |

# Cost analysis

- The estimated cost per CROB (assuming $130 for MicroPOD Tx and $280 for MicroPOD Rx):
  - (2 · $280 + $130 + $3500) /8CROBS ≈ $524 in case of 8-CROB version with XCKU11P in FFVE1517 enclosure.
  - (3 · $280 + $130 + $4600) /12CROBS ≈ $464 in case of 12-CROB version with XCKU15P in FFVA1517 enclosure (optimistic assessment).
- For ZynqMP chips:
  - (2 · $280 + $130 + $3900) /8CROBS ≈ $574 in case of 8-CROB version with ZU11EG in FFVC1760 enclosure.
  - (3 · $280 + $130 + $4900) /12CROBS ≈ $489 in case of 12-CROB version with ZU17EG in FFVD1760 enclosure.

# Available FPGAs

| | KU11P | KU15P | | ZU11EG | ZU17EG | | ZU19EG | |
|---|---|---|---|---|---|---|---|---|
| System Logic Cells | 653,100 | 1,143,450 | | 653,100 | 926,124 | | 1,143,450 | |
| Block RAM Blocks | 600 | 984 | | 600 | 796 | | 984 | |
| UltraRAM Blocks | 80 | 128 | | 80 | 102 | | 128 | |
| Enclosure | FFVE 1517 | FFVE 1517 | FFVA 1760 | FFVC 1760 | FFVC 1760 | FFVD 1760 | FFVC 1760 | FFVD 1760 |
| Number of transceivers GTH+GTY | 32+20 | 32+24 | 44+32 | 32+16 | 32+16 | 44+28 | 32+16 | 44+28 |
| Total number of transceivers | 52 | 56 | 76 | 48 | 48 | 72 | 48 | 72 |
| Example price | ca. $3500 | ca. $4600 | ca. $7900 $5000? | ca. $3900 | ca. $4400 | ca. $4900 | ca. $4900 | ca. $5400 |

# Availability of resources

- Preliminary analysis of resources defines number of BRAMs (for data sorting): 320 BRAMs for 8 CROBs, 480 BRAMs for 12 CROBs
- Consumption of other resources will be assessed after the first CRI firmwares are created.
- Upgrade possibilities (in case of unsufficient resources, in the same package):
  - XCKU11P → XCKU15P
  - ZU11EG → ZU17EG
  - ZU17EG → ZU19EG

# Distribution of time and fast control

- The CROB modules (and further FEB boards) require precise synchronization
  - Time synchronization (PPS)
  - Reference clock signal
  - synchronous control commands.
- Possible solutions:
  - Support for White Rabbit or similar protocol in every CRI board. The clock is received from the optical signal from the SFP transceiver.
  - Dedicated receiver board use of a standard WR node in the form of a PCIe card. The clock and timing signals distributed via pear-to-pear or daisychain connectors (e.g. miniSAS).
  - Delivery of both clock and control messages using dedicated optical or copper links to each CRI board from external distribution unit (e.g. by the TTC-PON system developed at CERN)
- Selected FPGAs may implement all those options, but additional hardware (e.g. the jitter cleaner must be provided by the board)

# Is it reasonable to use ZynqMP?

- We get quad-core ARM almost for free!
- DDR4 RAM must be mounted – 2400MT/s throughput not sufficient to pass the data through the RAM even for 64 bit data width
- Use of ARM core for control
  - Without the DDR4 RAM – in „bare metal” configuration with 1MB L2 memory
  - With slow 32-bit LPDDR memory – even under control of Linux OS
- Usability of ZynqMP chips must be evaluated

# Visualization of the possible CRI board

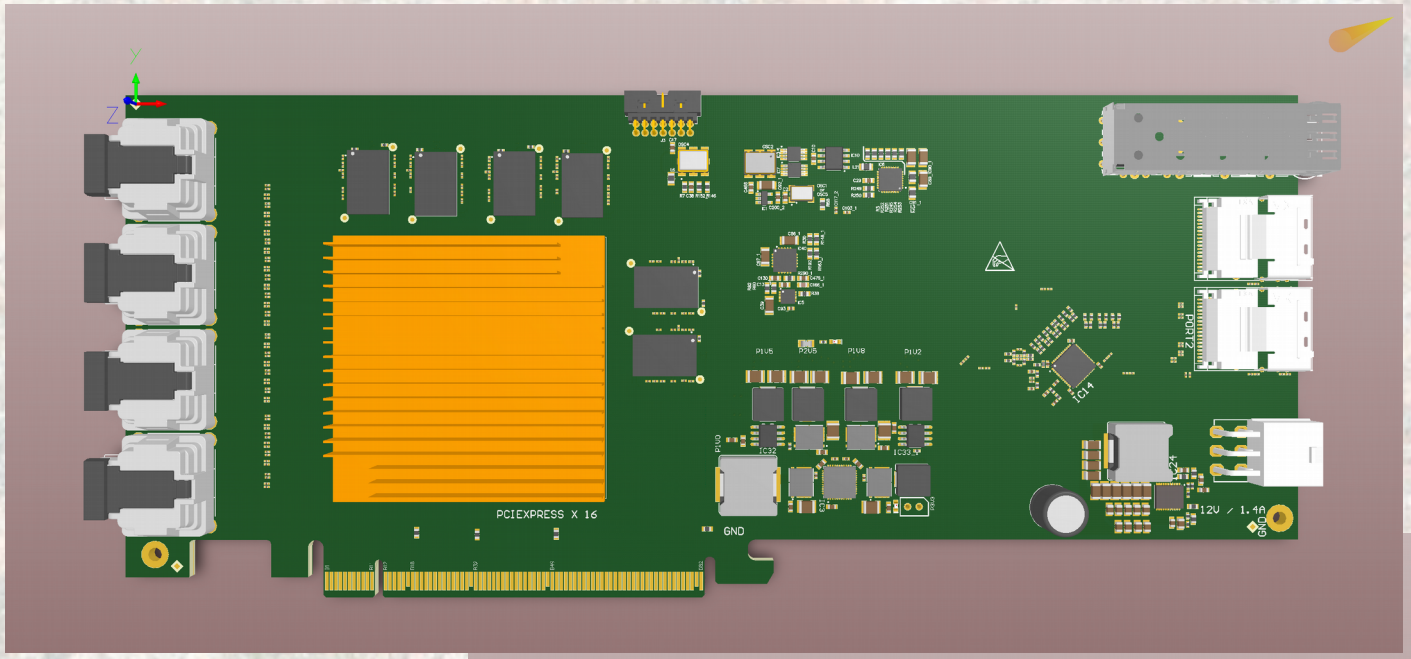


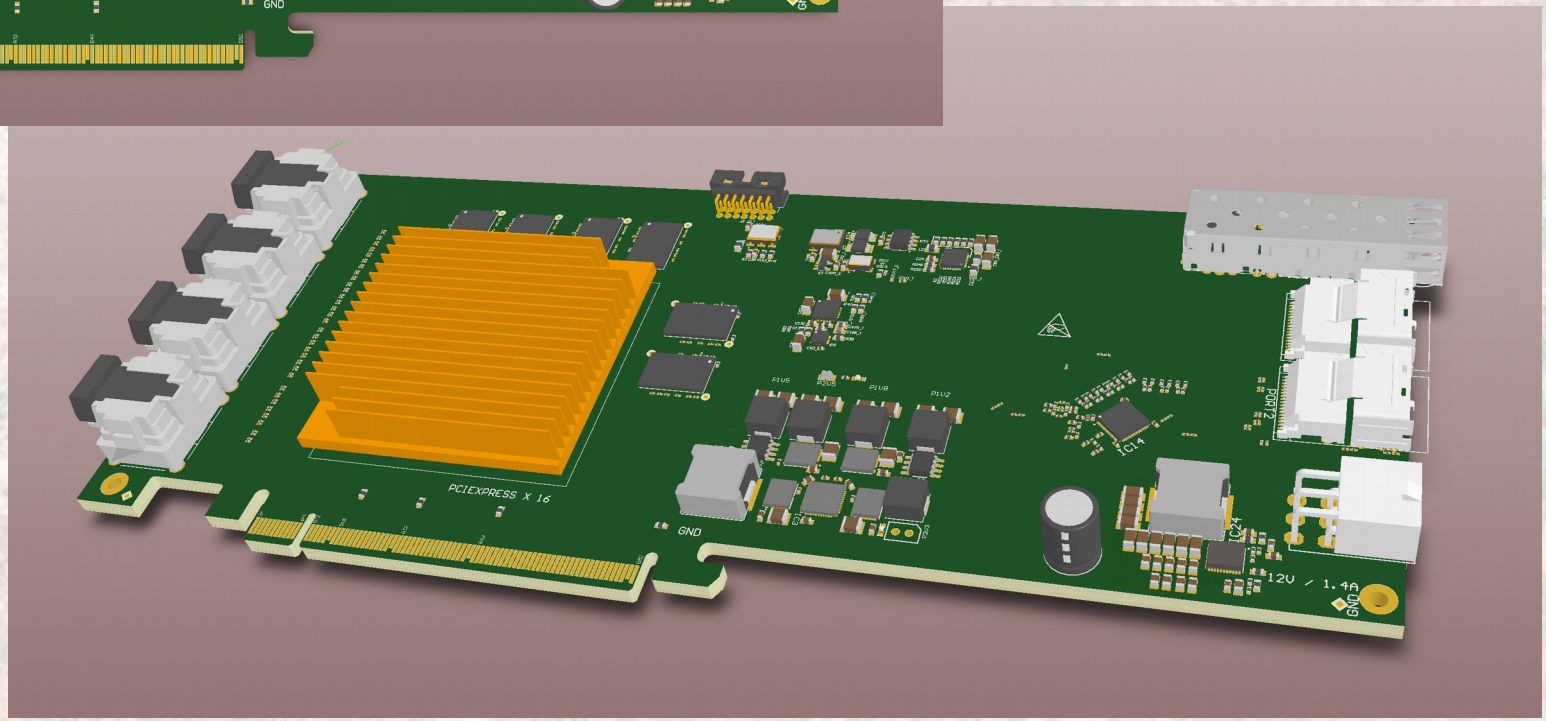

# Conclusions

- It is possible to create the CRI board supporting 8 or 12 CROBs in a PCIe card form factor, based on new UltraScale+ FPGAs from Xilinx
- The 12 CROB version, gives sligthtly lower cost/channel (12%) but  is more technologically challenging
- Information from the team responsible for FLES input nodes also suggest to use the 8 CROB version
- Final decision will be possible after the protype firmware is implemented.

Thank you for your attention!