The ATLAS Level-0 Muon Trigger in the barrel

HEP2022 - 39th Conference on Recent Developments in High Energy Physics and Cosmology, Thessaloniki, Greece

Riccardo Vari - INFN Roma
on behalf of the Level-0 Barrel Muon Trigger group

The ATLAS Trigger and DAQ schema for Run4

- Single level Level-0 hardware trigger with an output rate of 1 MHz, Level-0 readout latency is 10 μs
- Calorimeters and muons front-end full granularity readout at 40 MHz
- New Global Event processor replaces the current L1Topo and integrates topological functions with additional selection algorithms using information from muons and calorimeters
- Readout based on FELIX system for all detectors
- FPGA-based boards off-detector, on-detector where possible
- Possible hardware accelerator system for tracking at the Event Filter
- Goal of better e, γ, τ, jet identification and measurement, at hardware and software trigger levels and offline
- Event Filter output increases to 10 kHz

ATLAS Phase-II TDAQ Upgrade TDR

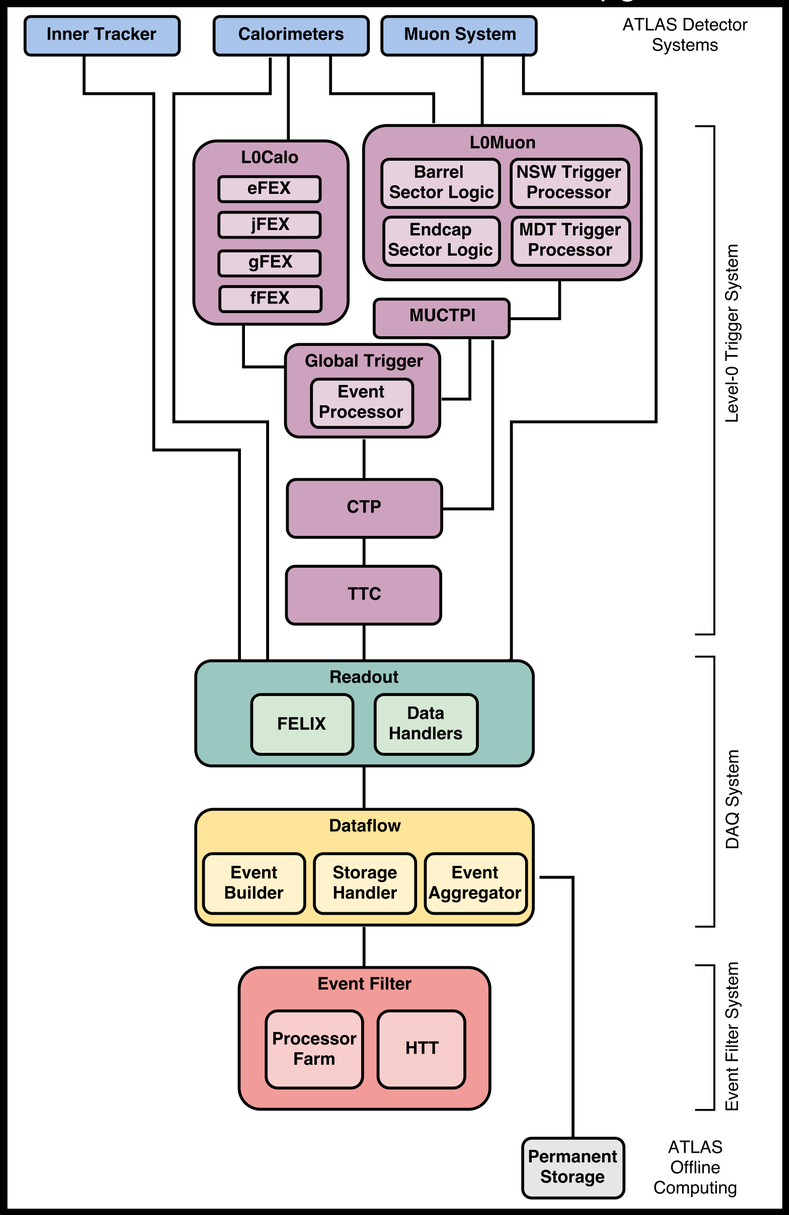

ATLAS muon detectors upgrade for Run4

- New RPC and sMDT detectors in the inner region of the barrel:
  - current BIS MDT replaced by new (sMDT + RPC)
  - new RPC triplets installed on top of the existing BIL MDT
- New sTGC triplets in the end-cap inner region EIL4
- The new detectors allow to:
  - reduce the trigger fake rate in barrel and end-cap regions
  - increase the trigger performances
  - increase the geometrical coverage in the barrel
- Detectors and front-end prototypes built, validation ongoing

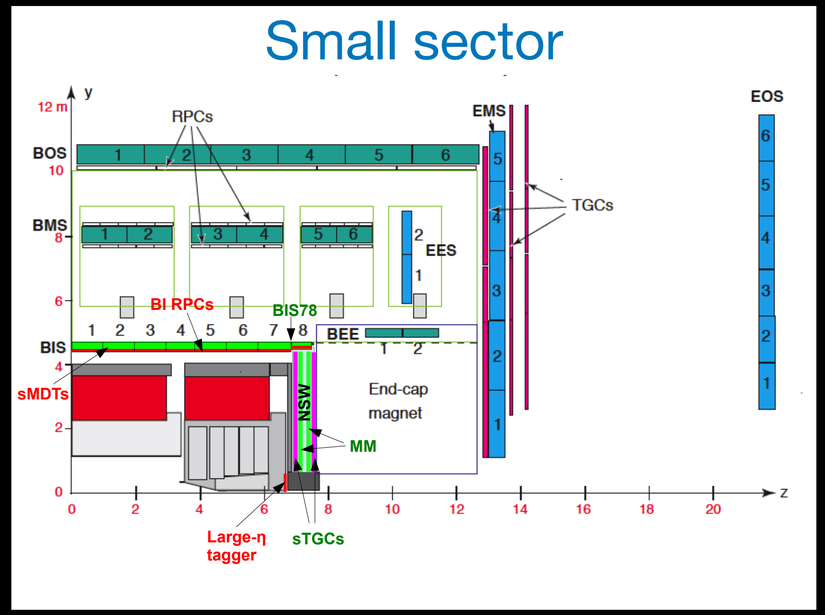

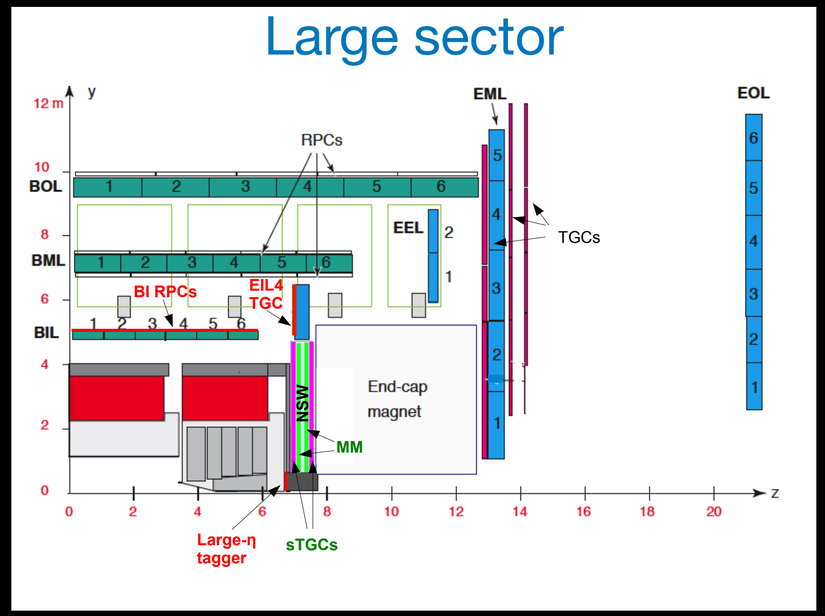

ATLAS Level-0 muon trigger

- The data from the RPC, TGC, and NSW detectors used in the Run3 system will be complemented with BI RPC, Tile calorimeter and MDT
- Increased selection efficiency and reduced fake trigger rate
- New MDT trigger sharpens turn-on curve and increases the rejection power
- Possibility to loose RPC trigger selection to increase the geometrical acceptance in the barrel, from ~70% to ~95%
- Rate suppression of ~50% for muons with p_T < 20 GeV
- New on-detector electronics full digital readout to off-detector @ 40 MHz
- Barrel and end-cap new off-detector Sector Logic trigger boards perform the coincidence trigger algorithm and send the seed to the MDT Trigger Processors
- New MDT Trigger Processors match the MDT hits with the RPC/TGC seed vectors in space and time
- Large use of FPGAs on and off-detector
- Board prototypes (front-end, SL, MDT-TP) available, currently under test

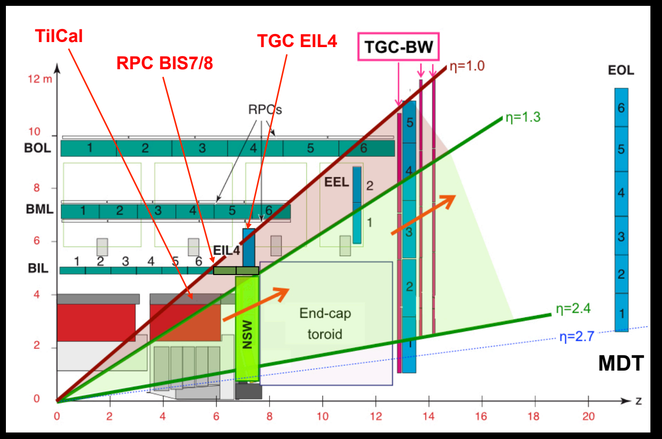

ATLAS Muon Spectrometer Phase-II TDR plots

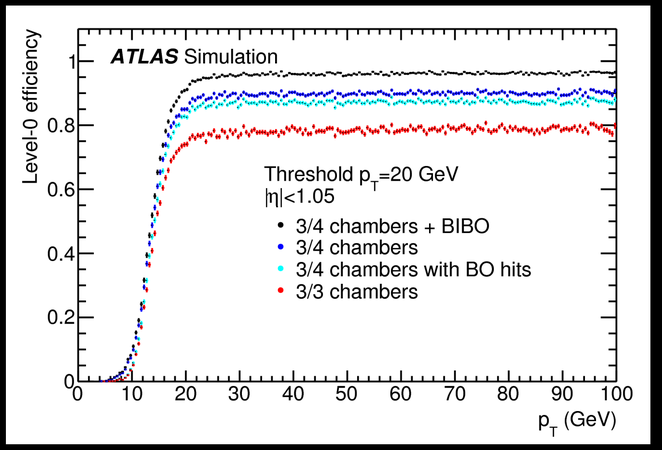

ATL-PHYS-PUB-2016-026

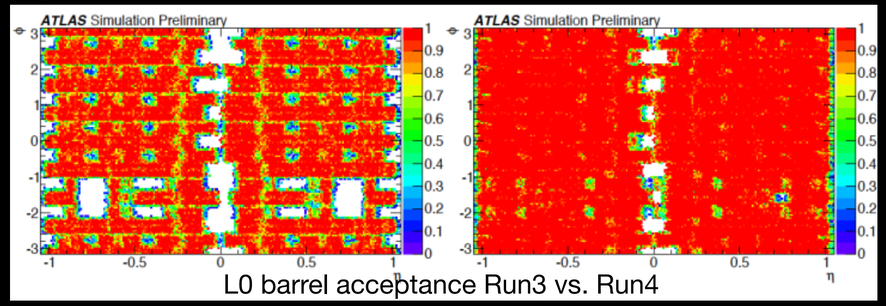

L0 barrel acceptance Run3 vs. Run4

RPC detector in the barrel

- The picture shows the RPC stations of one ATLAS sector (one side)
- 9 layers of concentric RPC
- Max number of RPC stations in one sector: 6 in BO region, 6/7 in BM, 10 in BI (BI final layout not fully defined yet)
- BM and BO are equipped with eta and phi strips, BI has eta only strips
- BM and BO eta strip length equal to half RPC station
- BM and BO sectors are currently treated as two independent trigger sectors
- BI chambers equipped with eta strips only (front-end electronics on both sides), phi coordinate to be derived by centroid calculation

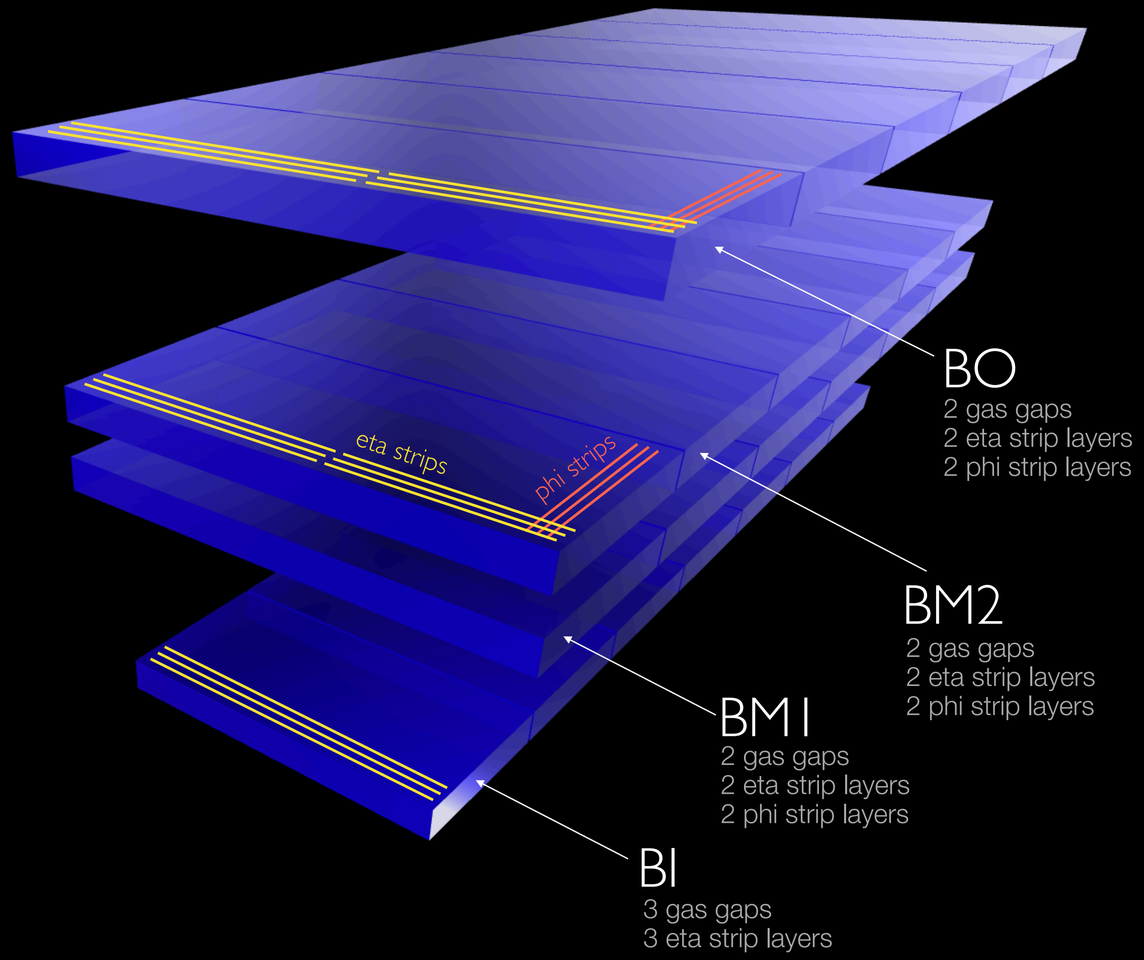

L0 muon trigger in the barrel region

- About 1600 DCT (Data Collector and Transmitter) boards in the barrel (up to 50 DCT per ATLAS sector)
- Max 288 RPC-FE inputs per DCT
- The DCT digitises the RPC-FE signals and sends the elaborated data out to the Sector Logic
- The off-detector Sector Logic board receives DCT data and executes the trigger algorithm and the readout logic
- BI-DCT receives digital signals from the FE-TDC (rising and falling edge time) and sends RPC-hit time data to the SL
- BMBO-DCT receives analog pulses from the FE, measures the rising and falling edge time with its FPGA-TDC (1.2 Gb/s sampling rate) and sends RPC-hit time data to the SL

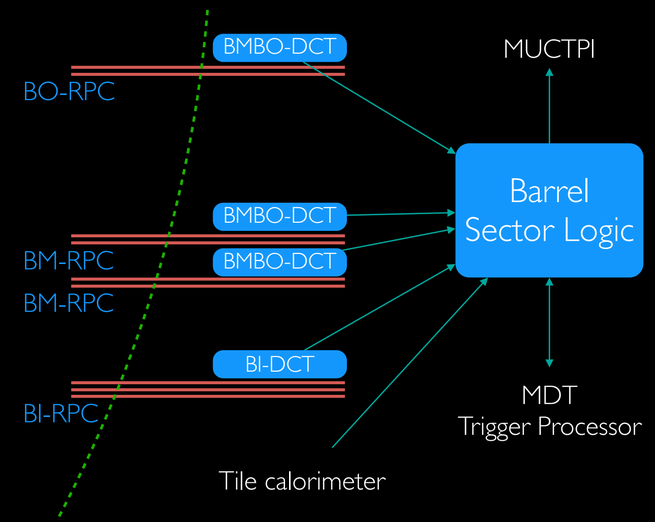

L0 Muon Barrel TDAQ schema

- 32 SL board in the L0 Muon Barrel system (one SL per ATLAS sector)
- Up to 50 DCT send RPC time hit data from one sector to the SL
- SL performs trigger algorithm based on hit coincidences on up to 9 RPC layers
- Up to 3 candidates sent to the MDT-TP, up to 4 candidates sent to the MUCTPI every BC
- RPC hit data stored into local buffers during L0 latency for readout
- BIS78 trigger algorithm performed independently and result sent to endcap SL
- Interface to L0 Barrel TDAQ server and DCS clients devoted to Zynq based mezzanine (Mercury XU5-5EV)

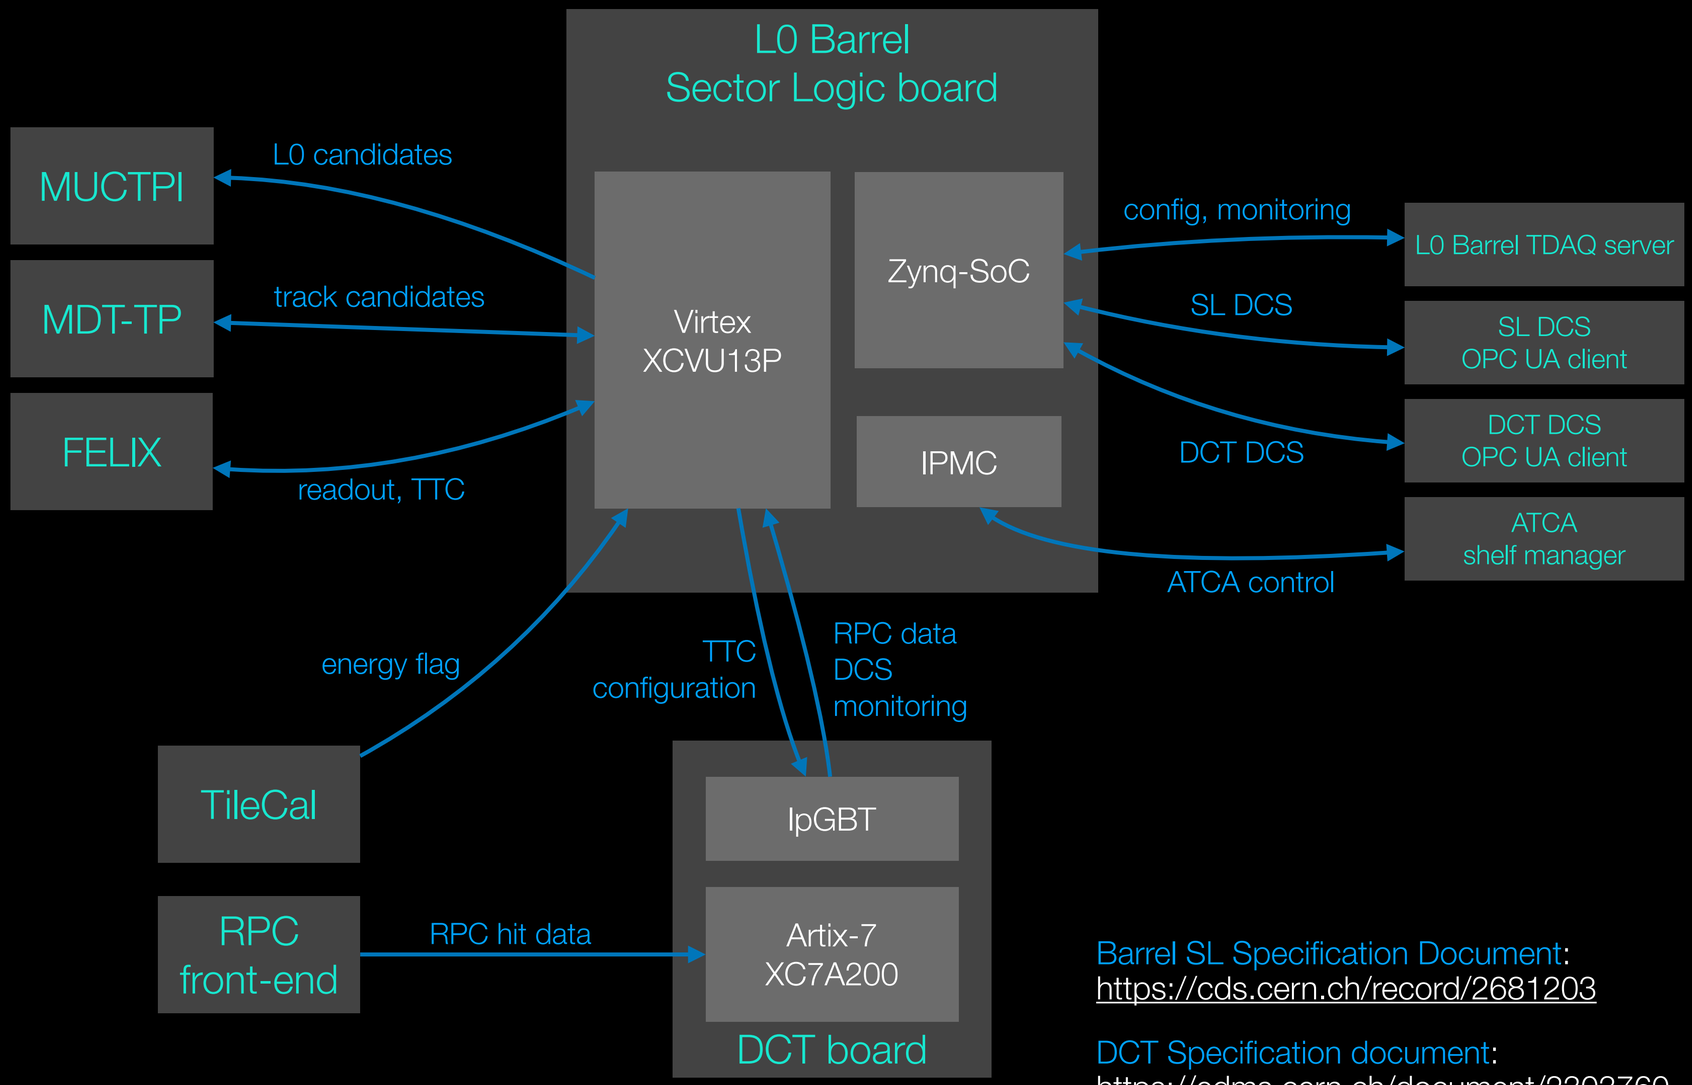

Barrel SL Specification Document:
https://cds.cern.ch/record/2681203

DCT Specification document:
https://edms.cern.ch/document/2303760

The DCT board

- FPGA based (Xilinx Artix-7 XC7A200T FFG1156 -3)
- 288 x 1.2 Gb/s input FPGA embedded deserialisers used for TDC logic (BMBO) or digital receivers (BI)
- BMBO-DCT and BI-DCT differences:
  - RPC-FE cabling, connectors and signal type are different
  - DCT FE receivers: LVDS for BMBO-DCT, CML for BI-DCT
  - BI-DCT only: 40 MHz clock sent from BI-DCT to FE boards
- 28 E-links @ 320 Mb/s used to send time and DCS data from the FPGA to the lpGBT

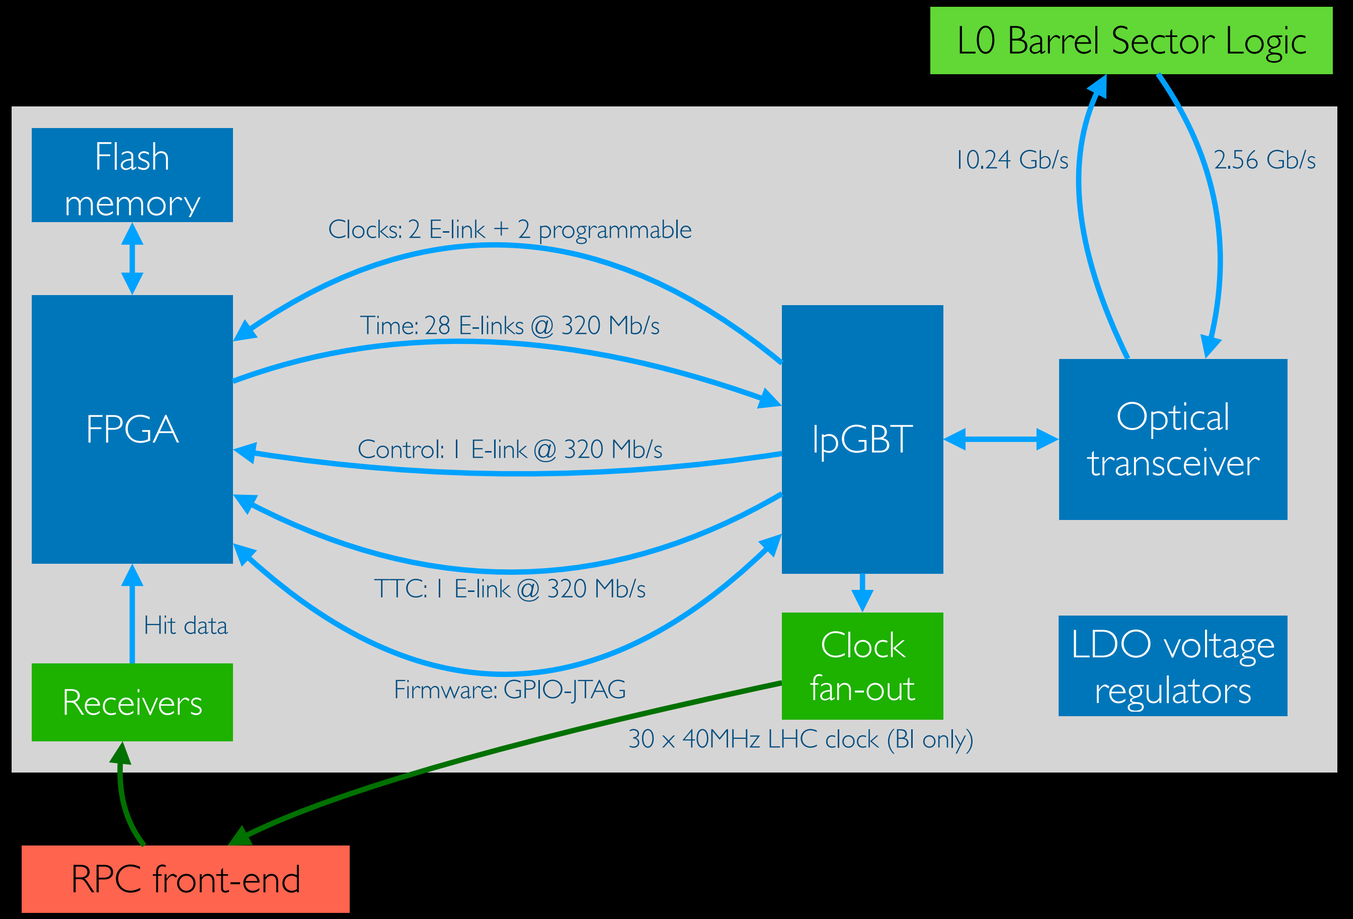

BMBO-DCT time data

27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
9-bit strip ID									9-bit BCID									5-bit rising time					5-bit falling time				

BI-DCT time data

27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
9-bit strip ID									5-bit BCID					14-bit (rising + falling) time													

BMBO and BI monitoring data

27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
9-bit reserved ID									19-bit monitoring word																		

BMBO-DCT prototype test

- Many tests done so far, in lab and using a real RPC
- No issues observed, moving towards final prototype design
- See Ioannis Mesolongitis' talk at ATLAS Muon Week: "Tests of RPC DCT BMBO prototype"

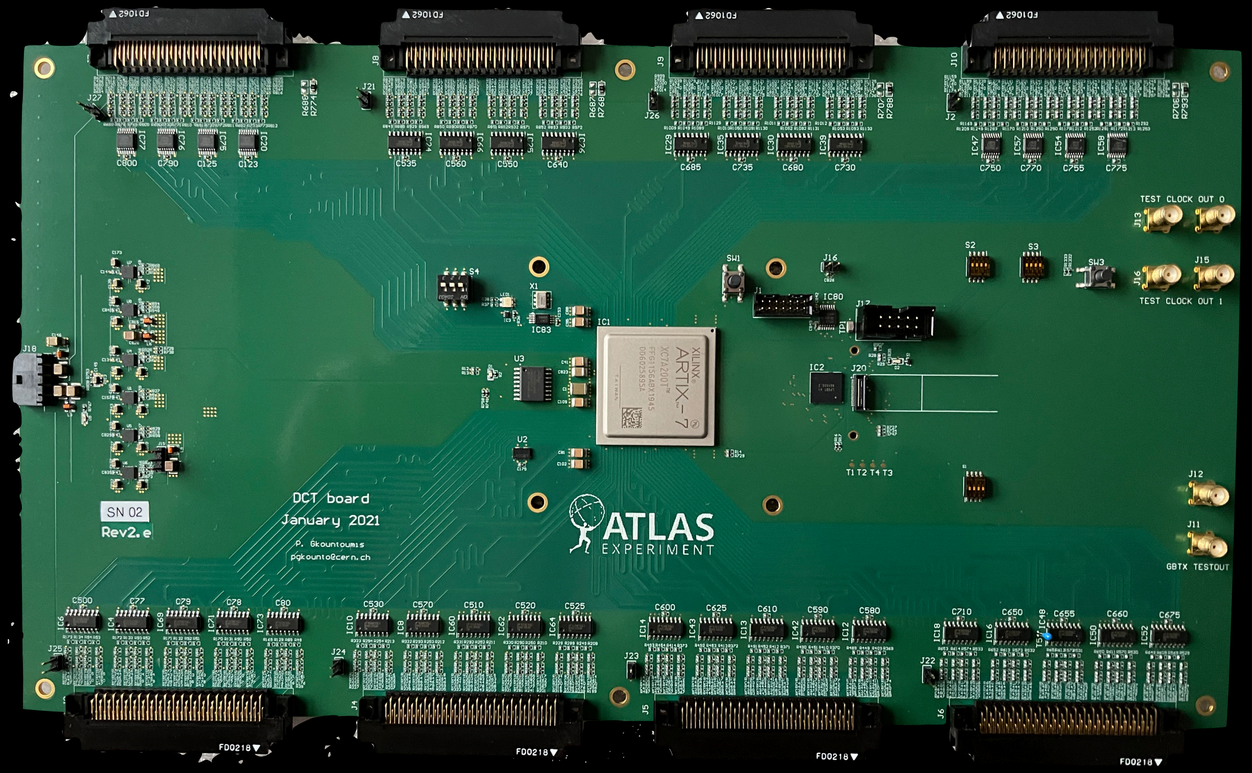

SFP+ eye diagram

VTRX+ eye diagram

IpGBT connectivity test

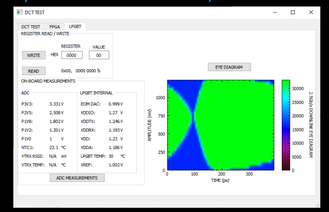

FPGA firmware upload through the IpGBT

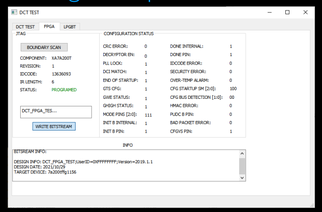

FPGA input

DCT connector input

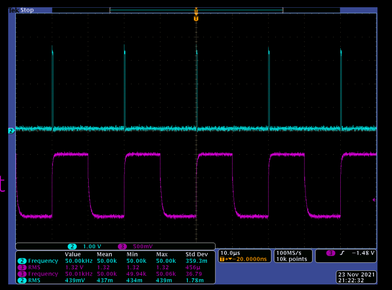

LVDS receiver output

Pulse to the front end

Front-end output

Latency requirements

- BMBO-DCT max allowed latency = 21 BC
- BI-DCT max allowed latency = 16 BC (BI-FE + BI-DCT max latency = 23 BC)

DCT latency: collision-to-last-bit sent to the Sector Logic

Contribution	Time [ns]	Description
ToF	60 (BO) 20 (BI) 32 (BIS8)	BO: 11 m for eta1, 18 m for eta6 (including curvature) BI: 6m for BIL1 (including curvature) BIS8: 9.5 m for BIS8 (including curvature)
RPC avalanche formation	14	From simulation
Propagation along the strip	12,5	2.5 m
FE amplifier + discriminator	0	negligible
FE TDC + serializer (BI only)	100-175	FE ASIC latency (zero for BM/BO)
Cable FE to DCT	50	10 m
DCT	375-525	DCT latency depends on the region and on the FE latency
fiber DCT to patch panel	25	5 m
fiber patch panel to USA15	400	80 m
TOTAL	1110	CBE value from TDAQ latency document

Latency simulation

- Simulation uses RPC hit real data from BML chambers, overlaying data from several BCs and several chambers to reproduce the expected phase-2 hit rates with a safety factor 2
- Most of the data are transmitted within the allocated latency
- Two events exceed very much the latency, they are an artefact of the hits overlay method, they have not been used in the latency plot
- The hits associated with a reconstructed muon are all sent out within the required latency
- This simulation was done with the BMBO logic (FPGA-DCT). 5 BCs to be added for the DCT-BI logic
- For BI-RPC 99.9% of the hits are sent within the required latency

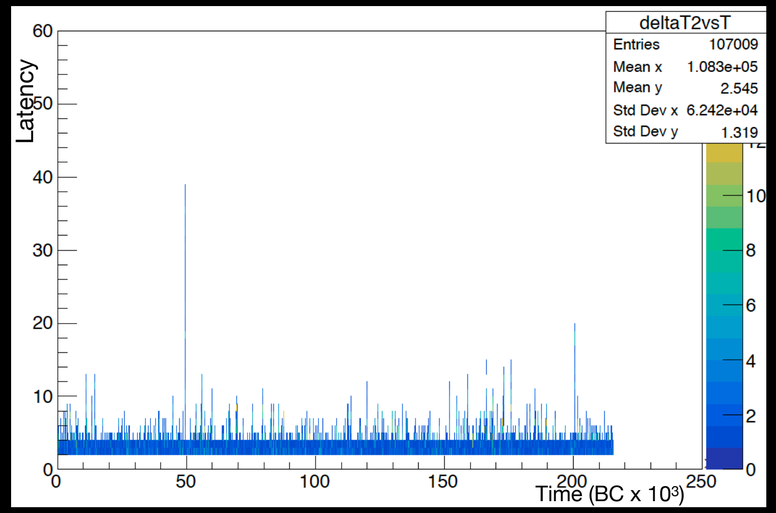

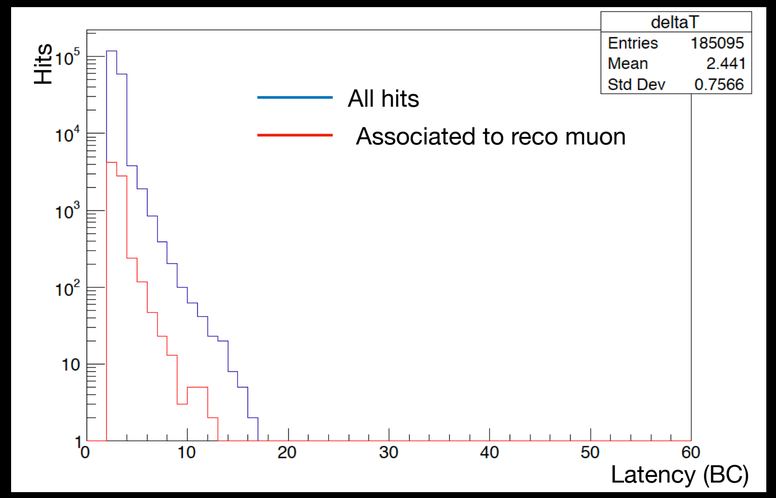

DCT FPGA firmware resources occupancy

BI-DCT FW

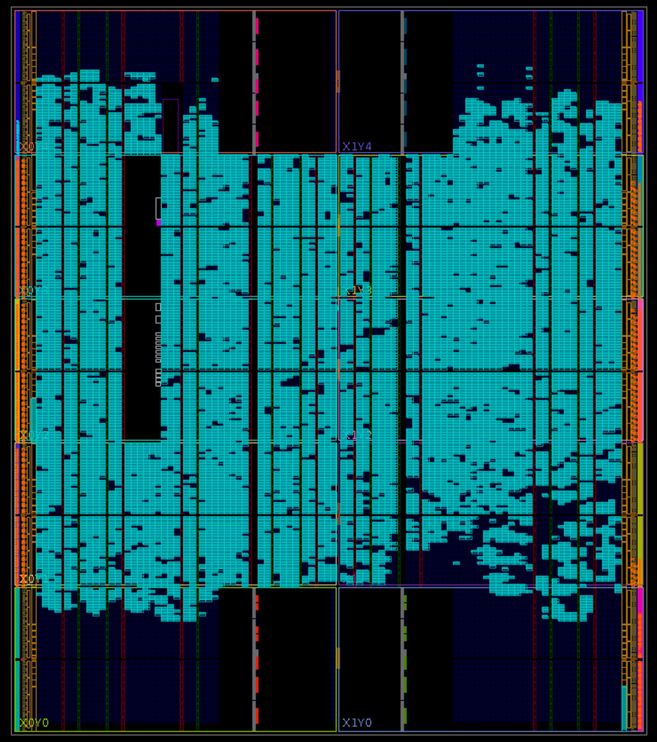

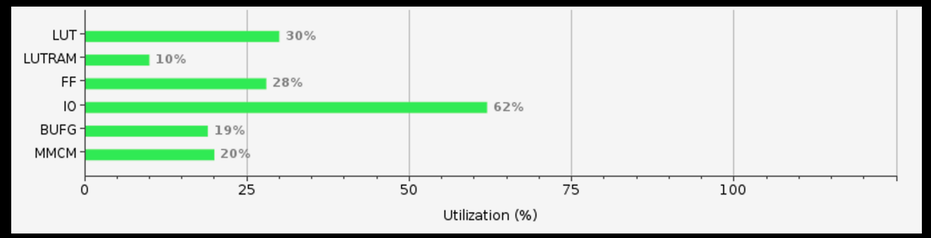

BMBO-DCT FW

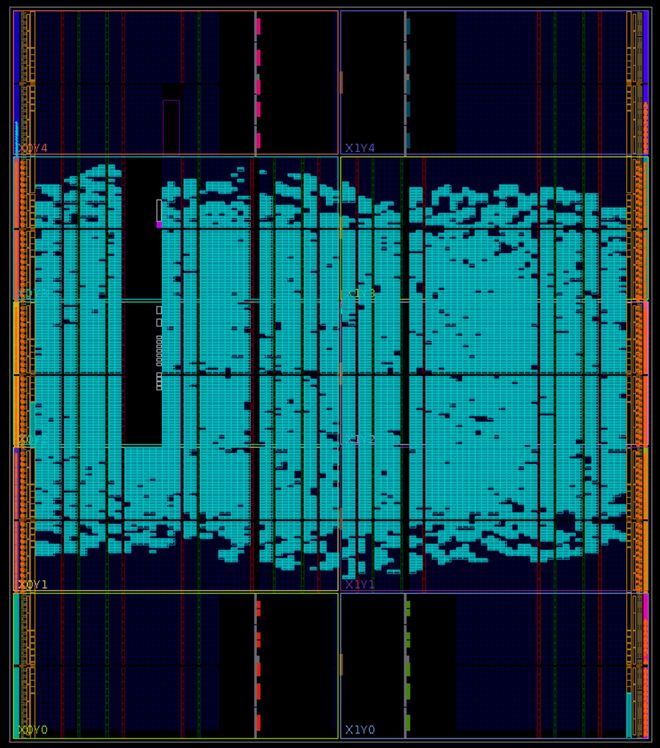

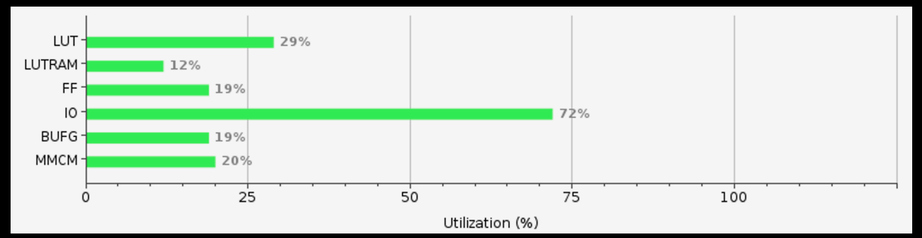

DCT components radiation tests

- lpGBT: certified
- SFP+ commercial optical transceiver: to be tested
- FPGA:
  - XCA200T to be tested with gamma (TID) and protons (SEE)
  - Artix-7 family technology already tested and certified (not the same device)
- Flash memory: choosing a component already tested and certified by other ATLAS groups
- TI DS90LV048 LVDS receiver: to be tested with gamma
- LDO voltage regulator TI TPS7A8500RGRT: successfully tested with neutrons
- BI-DCT additional components to be tested (fan-out chip and CML receivers)

	units	dose for 4000 fb^{-1}	sf_sim	sf_ldr		sf_lot		sf_total		requirement		SPR doc
				min	max	min	max	min	max	min	max	
TID	Gy	5,2	1,5	1	5	1	4	1,5	30	7,8	156	420
NIEL	cm^{-2}	2,80E+11	2	1	1	1	4	2	8	5,6E+11	2,24E+12	5,00E+12
Hadrons with E>20 MeV (SEE)	cm^{-2}	7,00E+10	2	1	1	1	4	2	8	1,4E+11	5,6E+11	1,00E+12

BMBO-DCT FPGA firmware

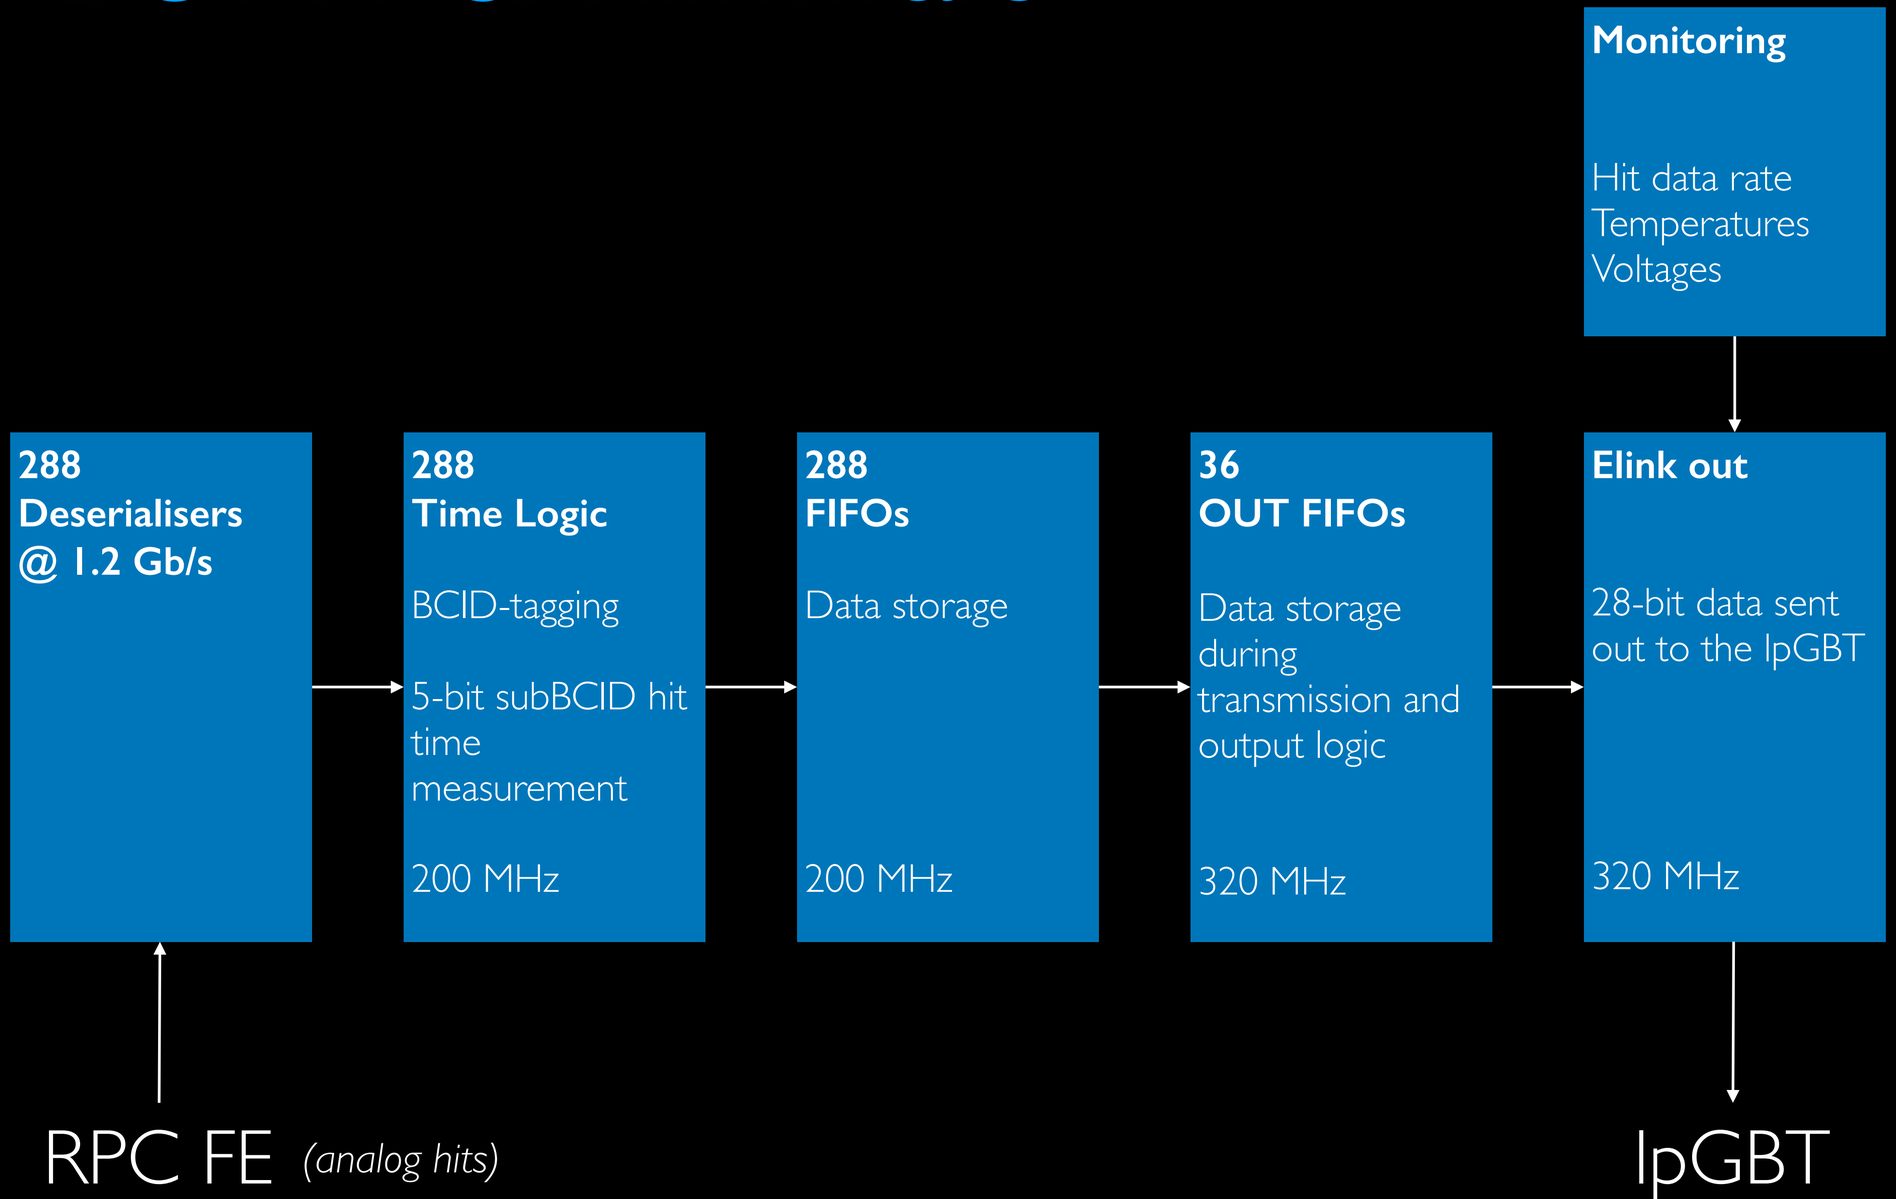

BI-DCT FPGA firmware

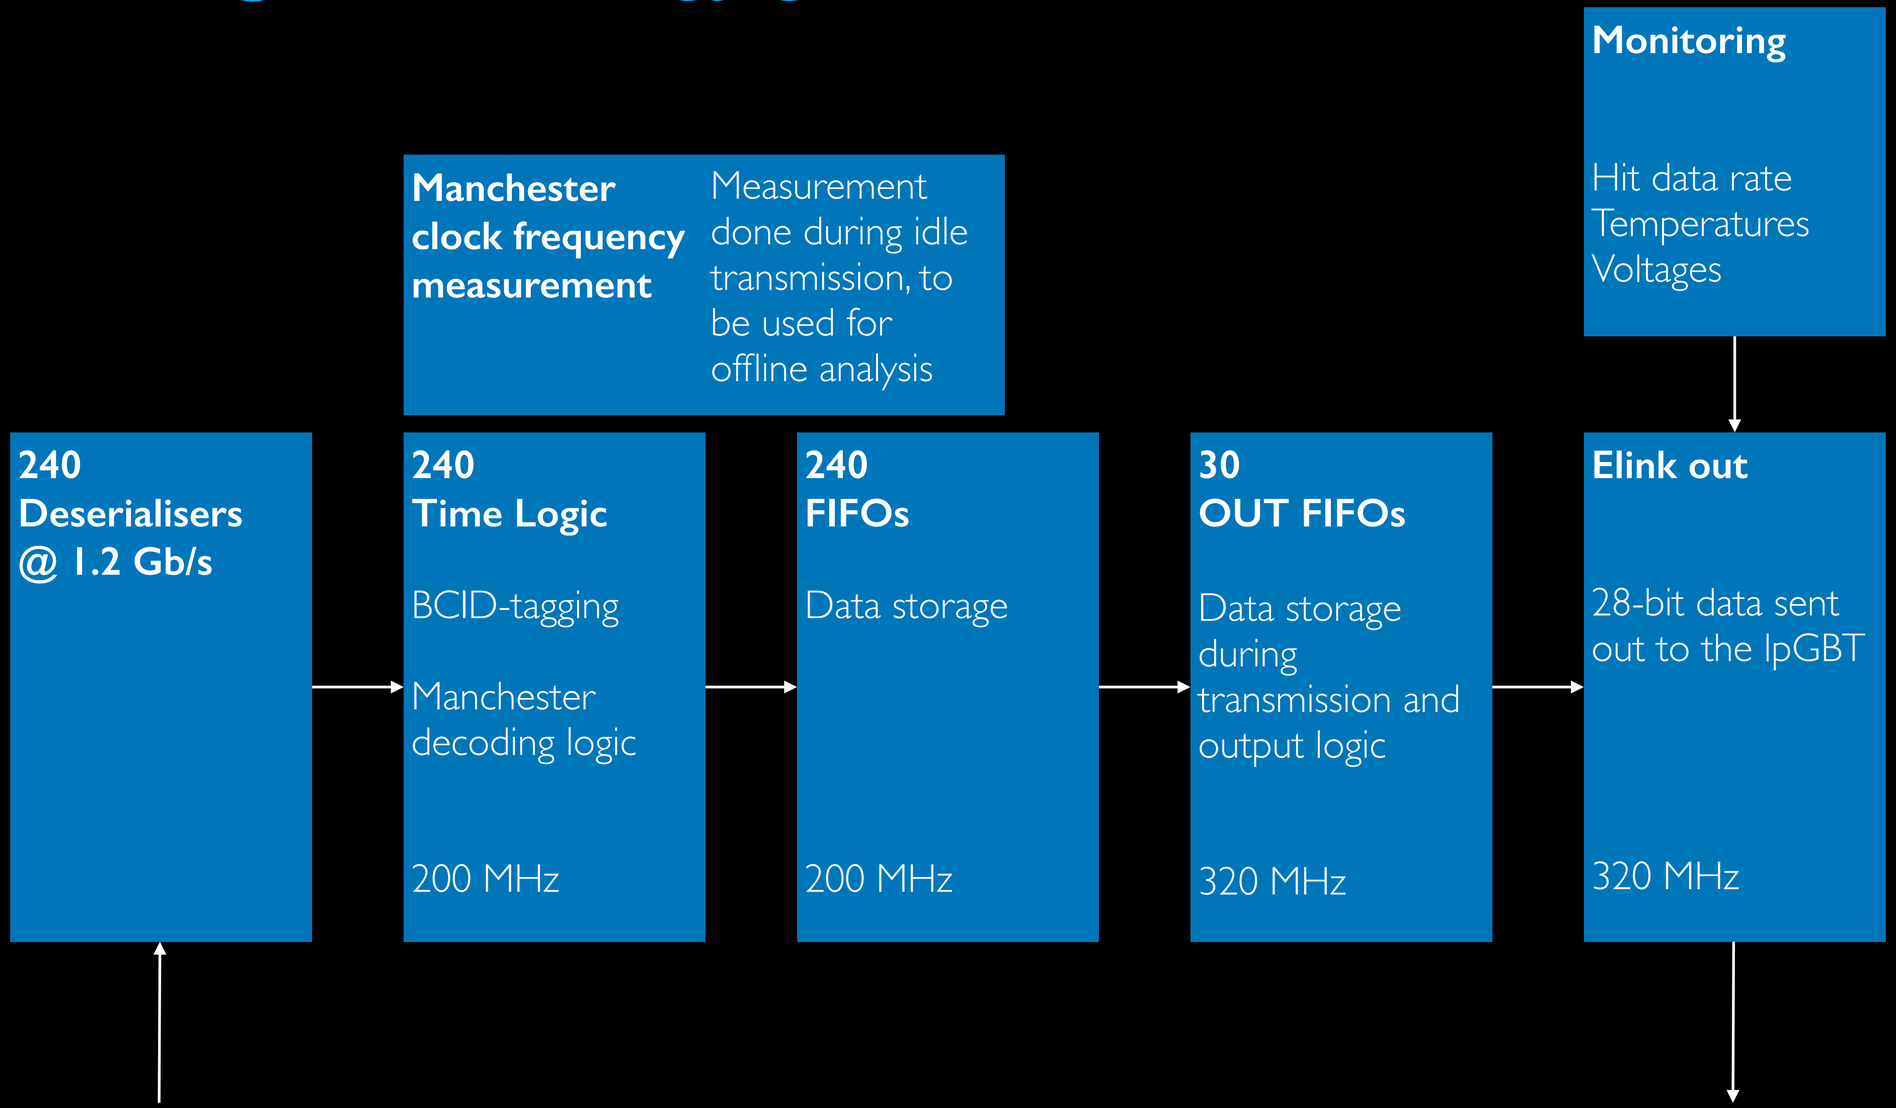

SL prototype

- ATCA board
- Xilinx Virtex Ultrascale+ FPGA (XCVU13P-FLGA2577-1-E)
- 20 FireFly optical transceivers, each one connecting 12 TX or 12 RX to the FPGA (total 120 TX + 120 RX)
- 10 x 24-fiber bundles on front panel + ethernet and SFP+ optical
- SoC mezzanine module with Xilinx Zynq Ultrascape+ FPGA (Mercury XU5-5EV)
- IPMC module for ATCA control

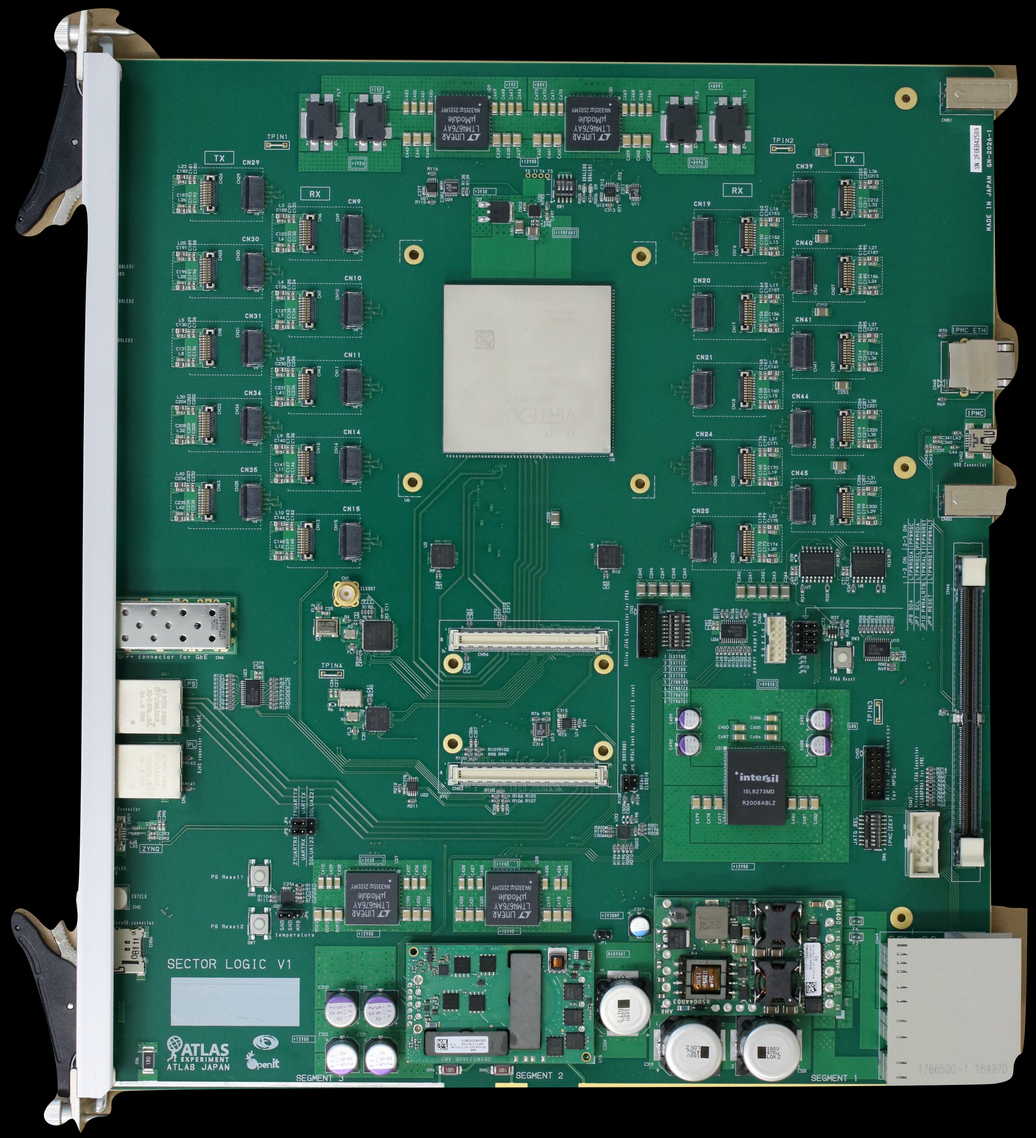

SL firmware block schema

- XCVU13P-FLGA2577-1-E device:
  - 3.8 million logic cells
  - 3.5 million flip flops
  - 1.7 million LUTs
  - 48.3 Mb distributed RAM
  - 94.5 Mb block RAM
  - 360 Mb UltraRAM
  - 12.3 thousands DSP slices
  - 128 GTY transceivers (up to 32.75 Gb/s)

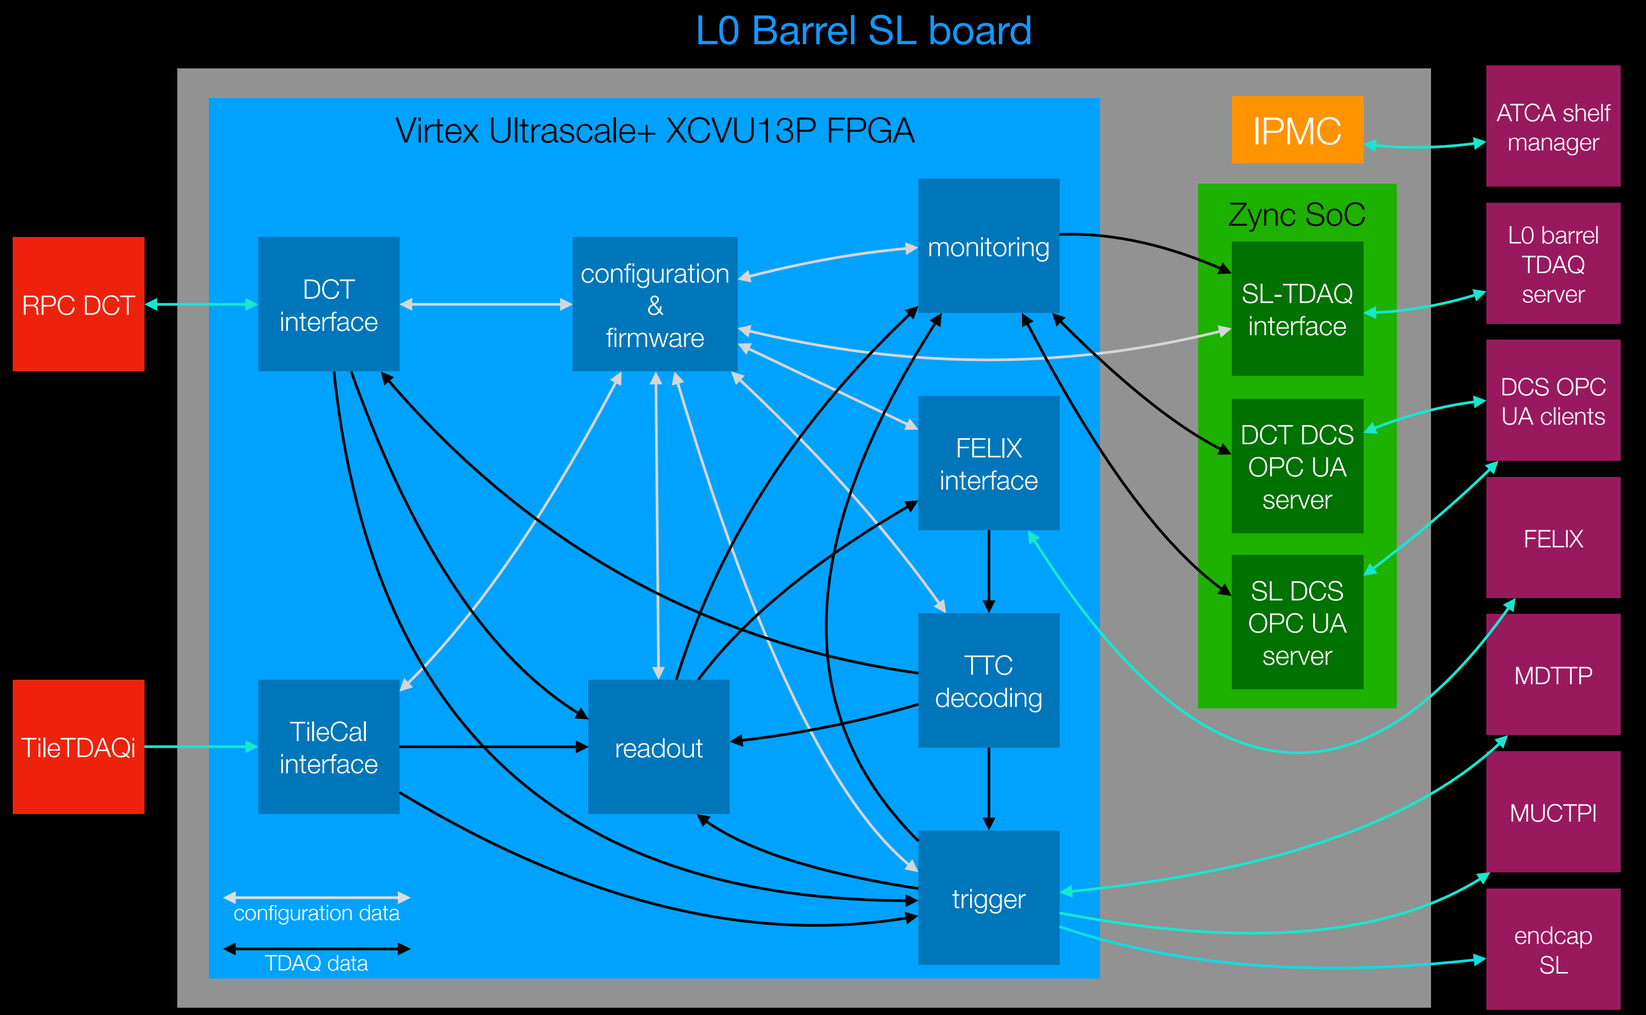

SL interface logic with DCT

- Uplink logic (lpGBT to SL):
  - Line rate: 10.24 Gbps (user data rate: 8.96 Gbps)
  - FEC5 coding (can correct up to 5 consecutive wrong bits)
  - User data: RPC data and monitoring data
- Downlink logic (SL to lpGBT):
  - Line rate: 2.56 Gbps (user data rate: 1.28 Gbps)
  - FEC12 coding (can correct up to 12 consecutive wrong bits)
  - User data: TTC, configuration commands and firmware
- lpGBT-FPGA core firmware provided by CERN, we applied some custom modifications
- Downlink components: scrambler, encoder, interleaver, TXgearbox
- Uplink components: frame-aligner, RXgearbox, deinterleaver, decoder, descrambler

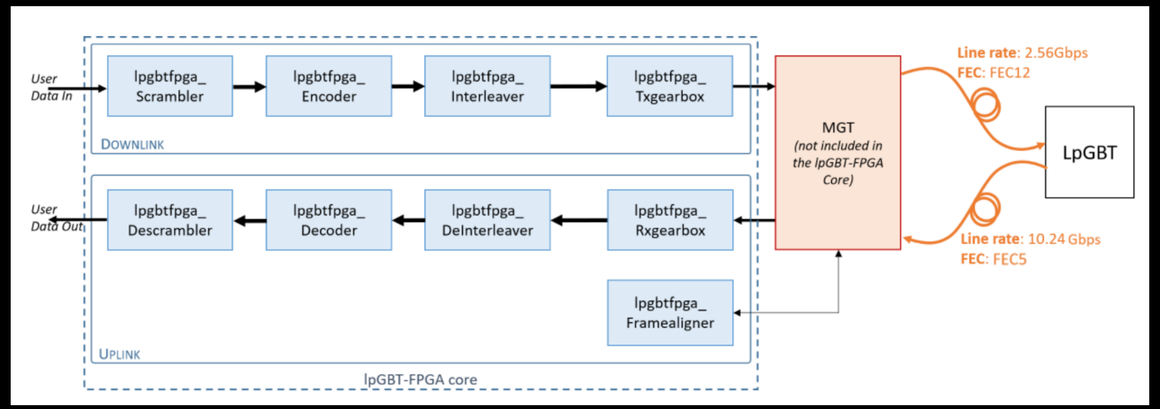

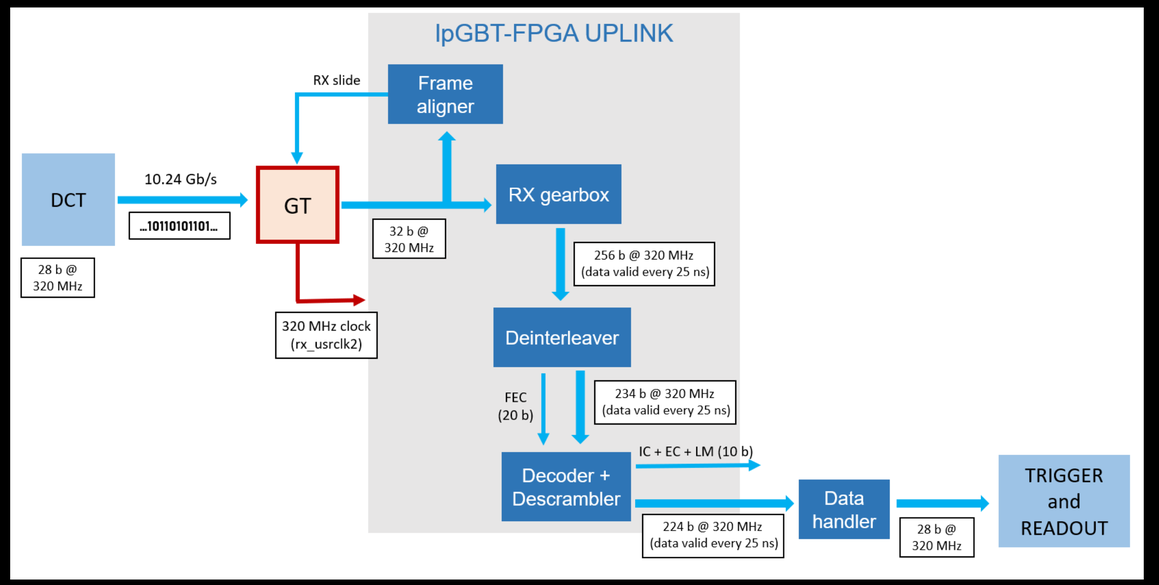

SL trigger logic

- Simulations used for the studies are based on MC single muons events from 1 to 100 GeV and 100% RPC collection efficiency
- Expected hit rates from one sector @ 7.5 x 10^{34} is in the range 100-180 MHz depending from the region
- The algorithm looks for hits in a coincidence windows around a pivot candidate, the width of the window depending on the p_T
- BI-BM-BO (3/4) or BM-BO (3/3) or BI-BO (2/2) majorities are required
- Up to four candidates are selected and sent to MDT-TP every BC
- NN-based alternative algorithm is also under study

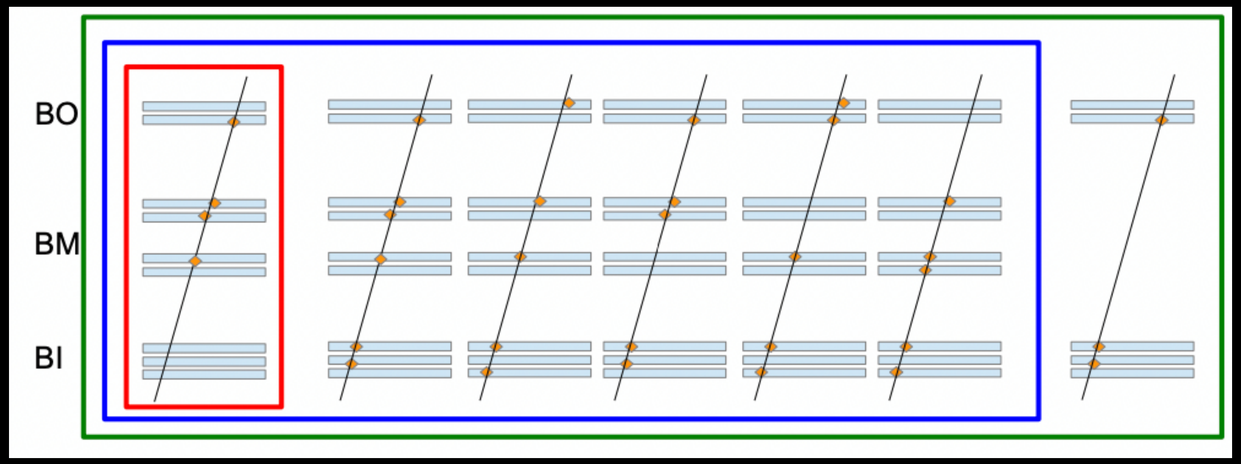

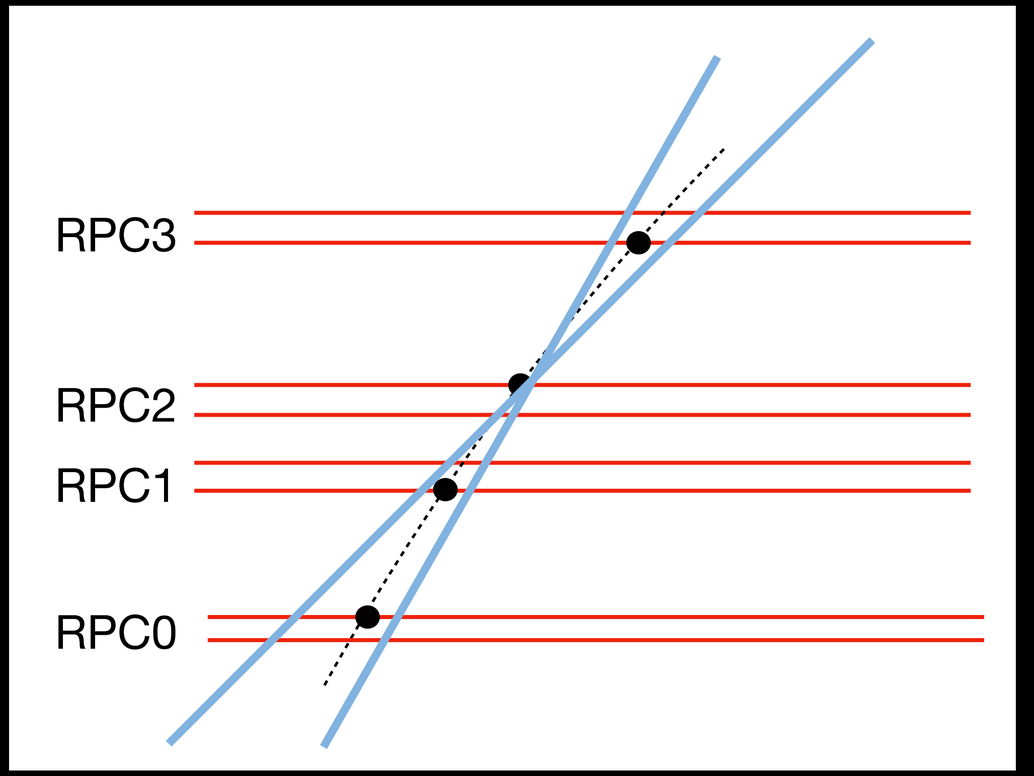

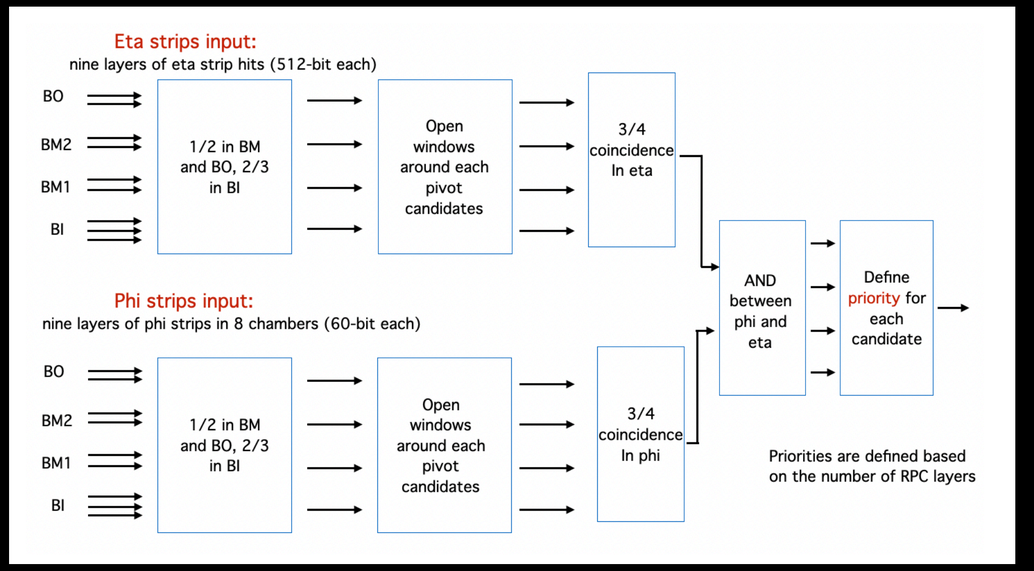

SL trigger latency requirements

- Last bit received from DCT: 1110 ns from BC
- Last bit received from Tile Calorimeter:1425 ns from BC
- SL trigger algorithm: 390 ns
- Time to transmit the candidates to the MDT-TP: 175 ns
- Last bit in from MDT-TP: 3810 ns from BC
- Final candidate decision: 25 ns
- Time to send the candidates to the MUCTPI:175 ns

SL readout logic

- 50 RAM based L0 latency circular buffers, one per each DCT
- Each memory buffer stores the DCT incoming data, ordered by their BCID
- Data are read out and sent to Felix in case a L0 signal is issued by ATLAS CTP
- Old data are discarded

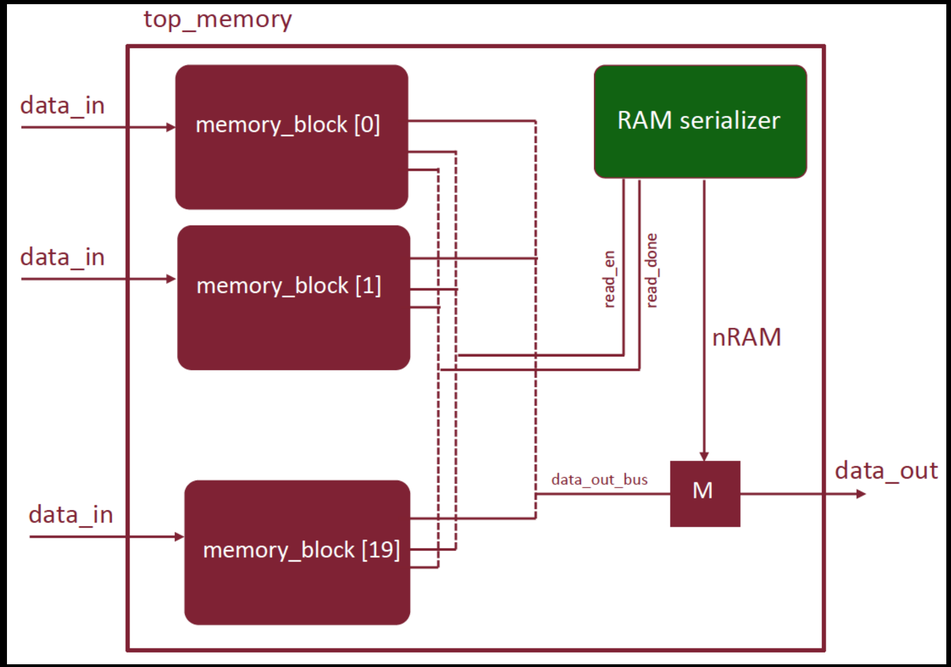

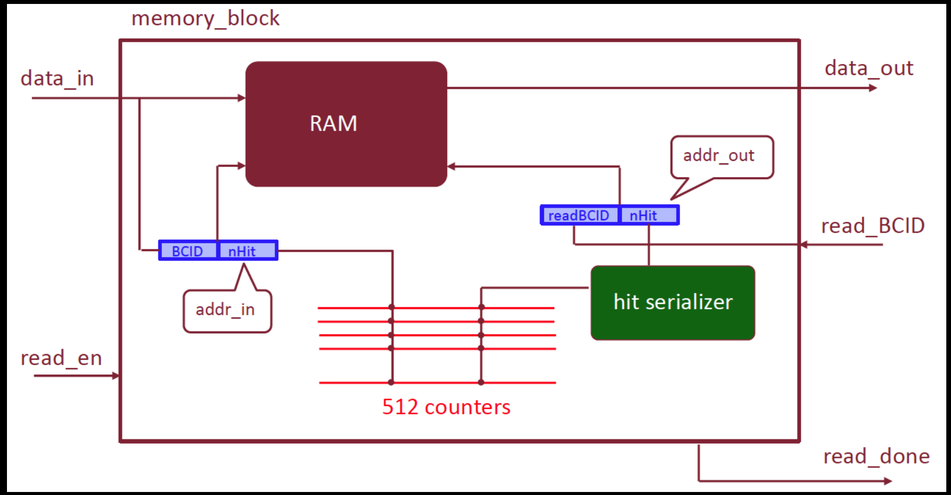

SL firmware floor planning

- XCVU13P FPGA has 128 GTY transceivers
- 120 GTY are connected to the FireFly optical transceivers modules on the SL board
- FPGA is made by four Super Logic Regions (SLR)
- Trigger and readout logic partitioned in the four SLRs, each one for a different part of the RPC sector
- SLR0 and SLR2 receive BMBO-DCT data from half-phi BMBO sector, SLR1 receives BI-DCT data
- High FPGA resource occupancy, 320 MHz internal clock and high number of signals crossing SLRs result in very difficult timing closure

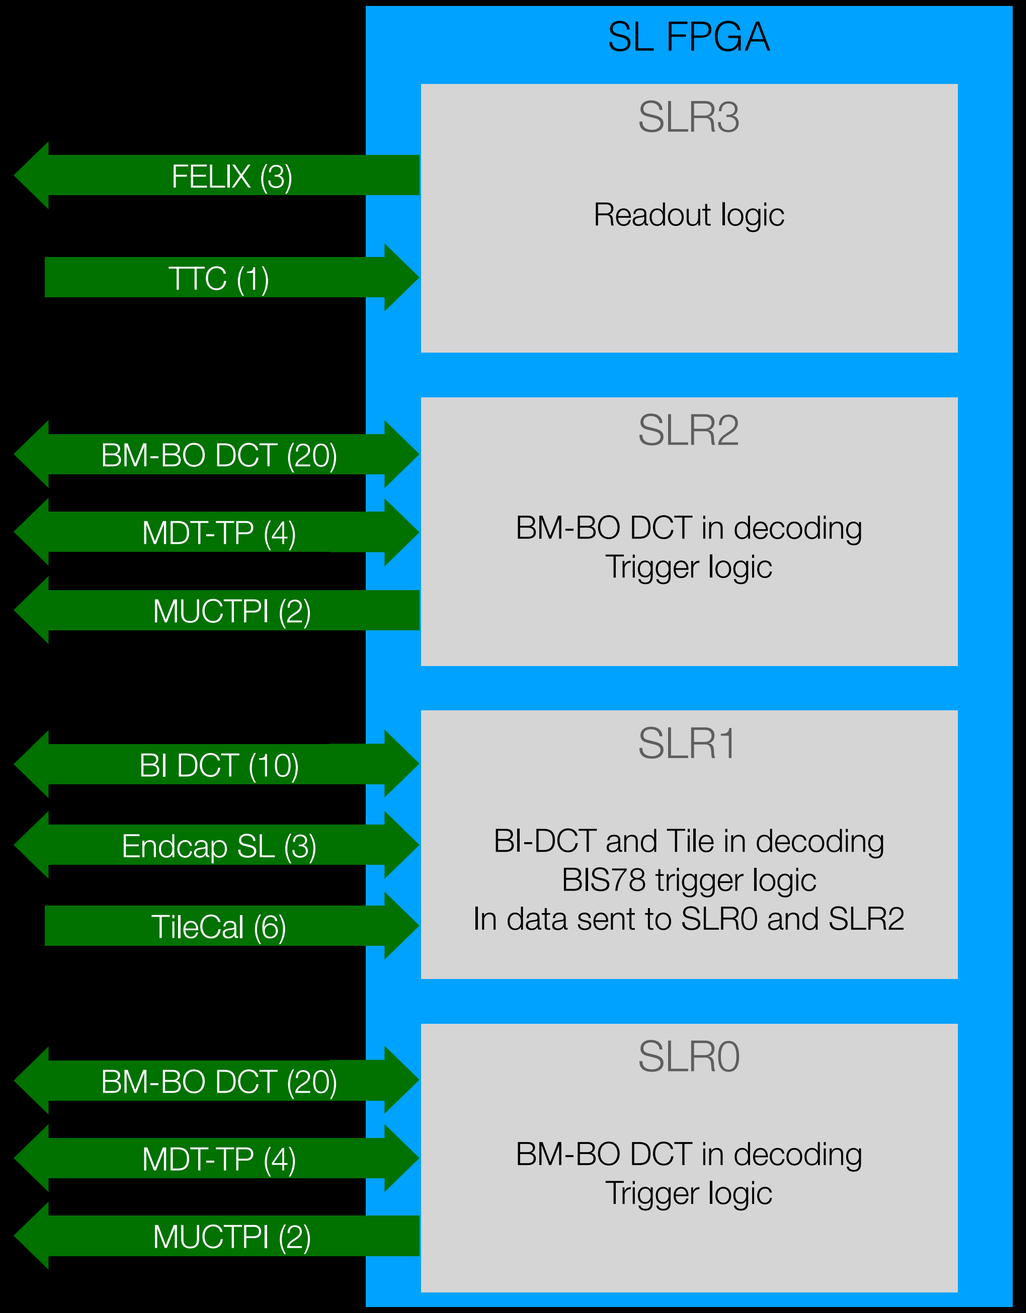

SL SoC

- The Zynq SoC provides the interface with the barrel TDAQ server for SL/DCT control, configuration, monitoring and firmware upload
- The Zynq SoC also runs the SL/DCT server to provide DCS data to ATLAS DCS clients
- AXI Chip2Chip with Aurora protocol is used for the Virtex-Zynq communication
- The Zynq runs Petalinux OS
- Many tests done on FPGA commercial boards, next step is to repeat the tests on the SL prototype

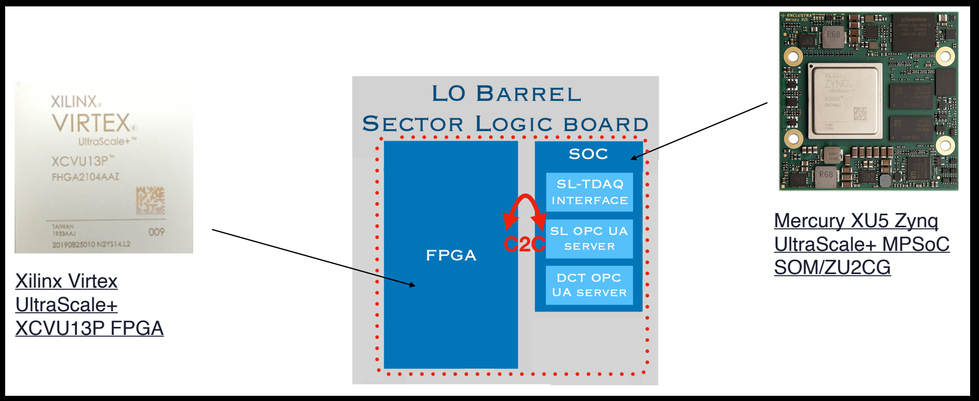

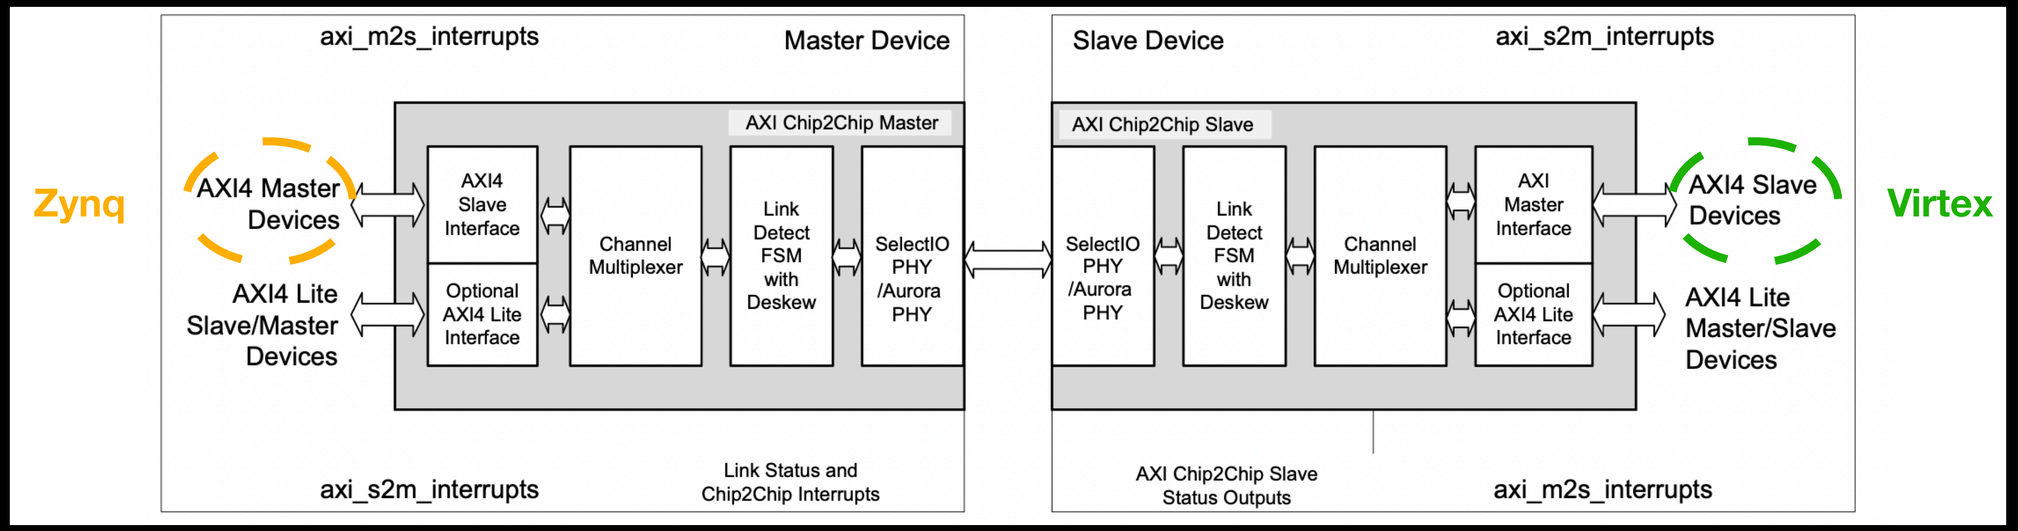

Conclusions

- BMBO-DCT and SL prototypes available, BI-DCT expected by early 2023
- Tests are ongoing, next step is integration tests at CERN, where we are currently assembling the L0 Muon testbed
- Huge effort taken from firmware development
- Hardware and firmware testing will continue for the whole 2023
- Moving towards final prototypes in 2023 (warning: global chip shortage might cause delays)
- Pre-production foreseen in 2024, production in 2026