CMS Concentrator IC (CIC)

Luigi Caponetto

With contributions from: G. Galbit, B. Nodari, S. Viret, D. Ceresa, S. Scarfi'

Context

CMS upgrade for the High Luminosity LHC

- **increased luminosity**
  - events pileup x5
  - increased cumulative radiation effects
- **CMS Outer Tracker (20cm < R < 120cm)**
  - silicon modules (strip + macro pixels sensors) with increased granularity
  - L1 trigger decision based on tracker data:
    - tracker provides real-time data from each BX
  - increased granularity and data amount
    - higher L1 trigger rate
    - increases in data bandwidth
    - local storage of data
    - increased latency
  - reduced material budget
  - increased radiation tolerance
    - ~100Mrad TID in worst case

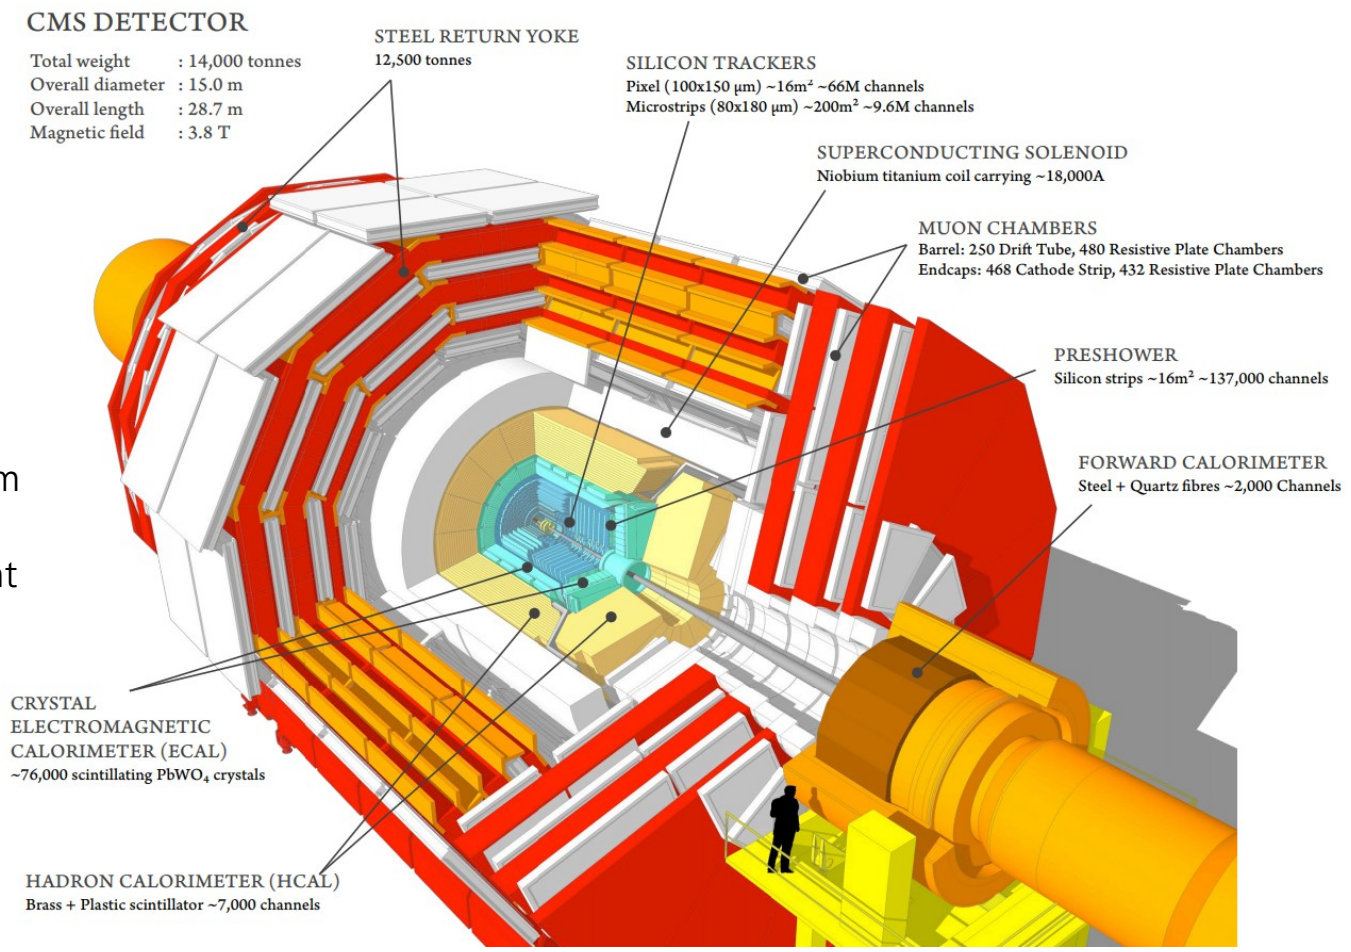

Context

CMS upgrade for the High Luminosity LHC

- **increased luminosity**
  - events pileup x5
  - increased cumulative radiation effects
- **CMS Outer Tracker (20cm < R < 120cm)**
  - silicon modules (strip + macro pixels sensors) with increased granularity
  - L1 trigger decision based on tracker data:
    - tracker provides real-time data from each BX
  - increased granularity and data amount
    - higher L1 trigger rate
    - increases in data bandwidth
    - local storage of data
    - increased latency
  - reduced material budget
  - increased radiation tolerance
    - ~100Mrad TID in worst case
- **reduction of the data rate processed by the L1 trigger**
  - binary readout modules locally select high-p_T (transverse momentum) only particles

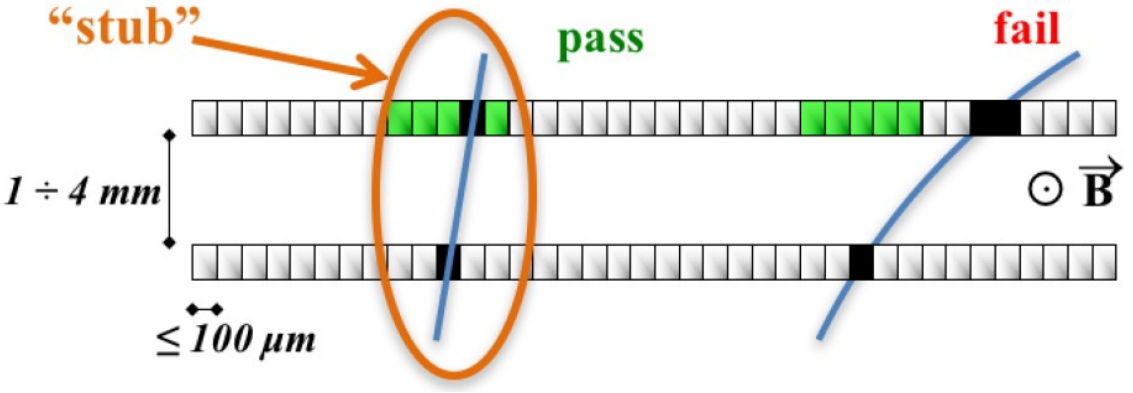

Outer Tracker p_T-modules

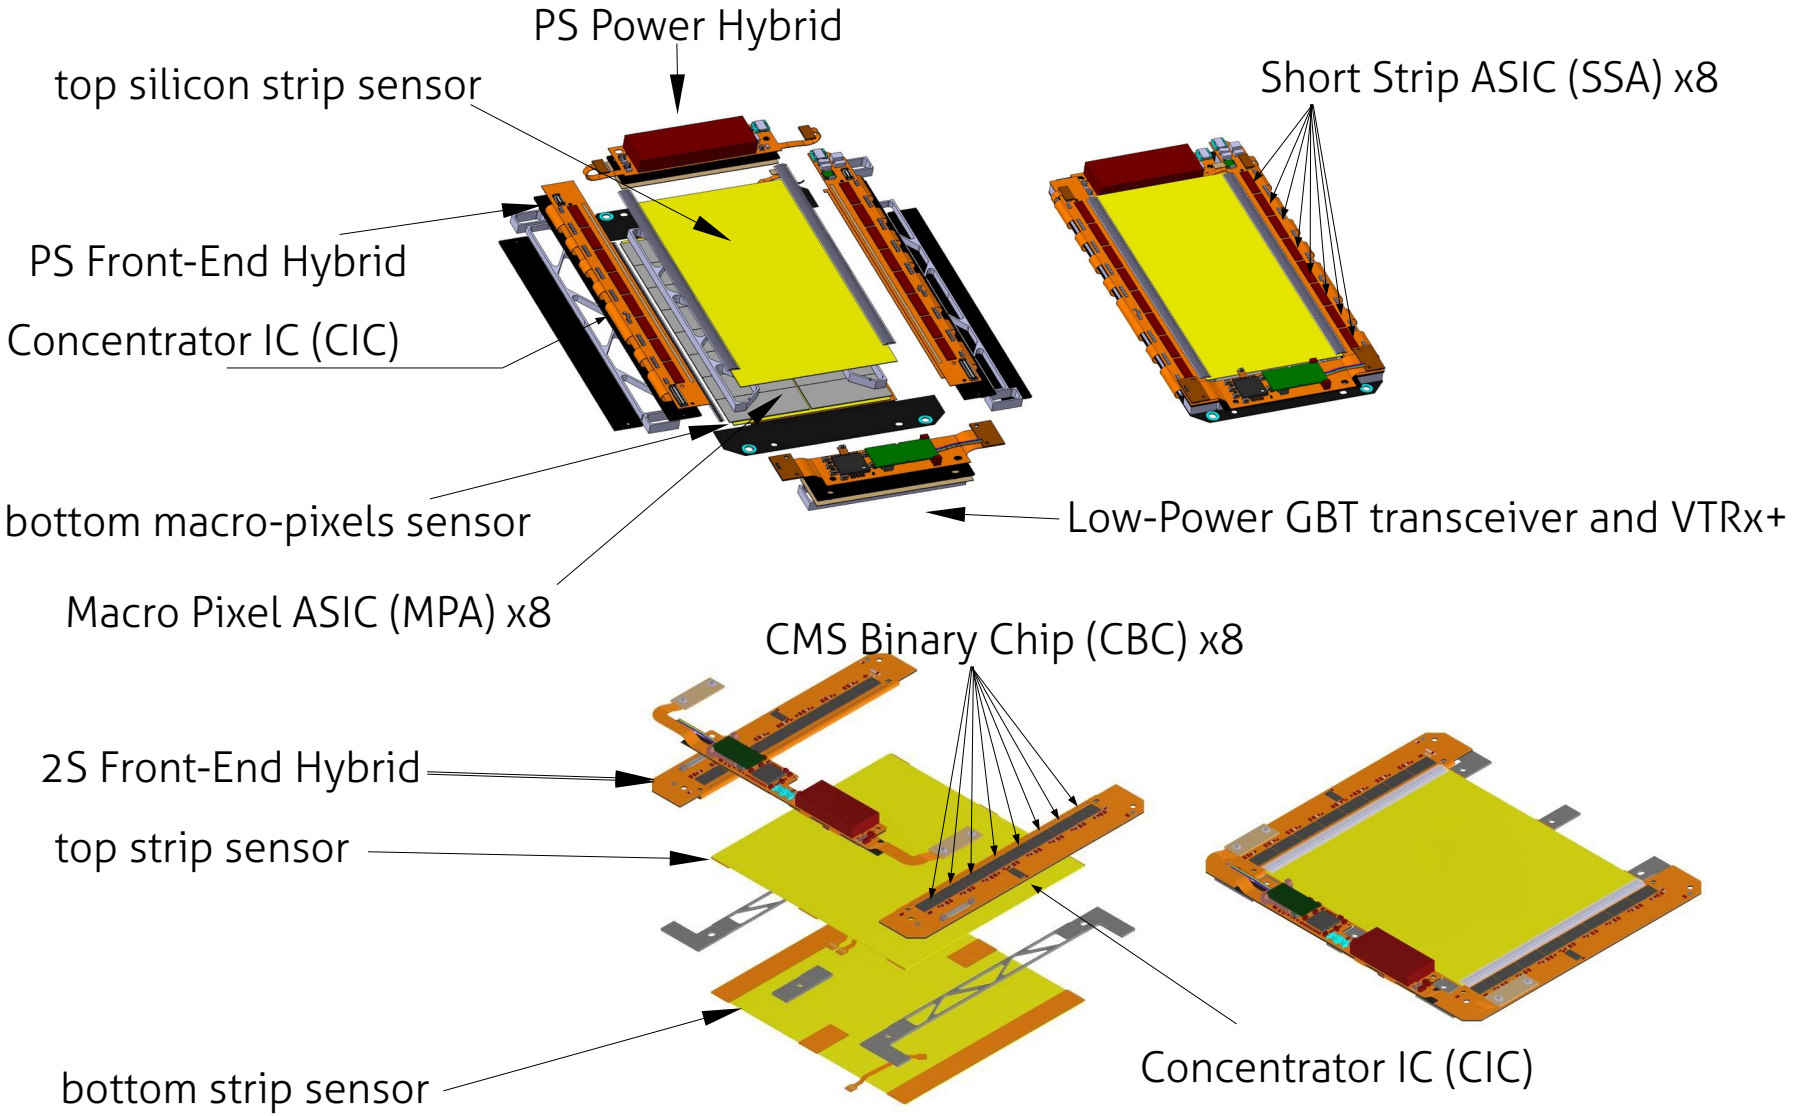

Data digital readout

all readout asics participate to the p_T-module concept:

"intelligent" discrimination is (digitally) performed on-chip by MPA, SSA, CBC and CIC

- stubs data are created at every BX by the front-end asics
  - this is the real-time information used for L1 trigger decision

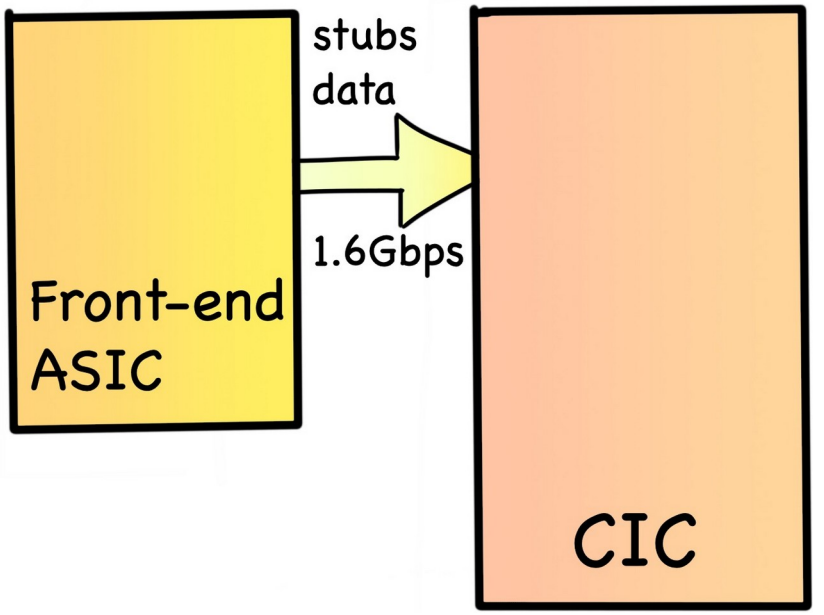

Data digital readout

all readout asics participate to the p_T-module concept:

"intelligent" discrimination is (digitally) performed on-chip by MPA, SSA, CBC and CIC

- stubs data are created at every BX by the front-end asics
  - this is the real-time information used for L1 trigger decision
- stubs data need processing before being sent to the back-end
  - CBC data need zeros-suppression
  - module output bandwidth needs to be efficiently exploited sending out only best event candidates

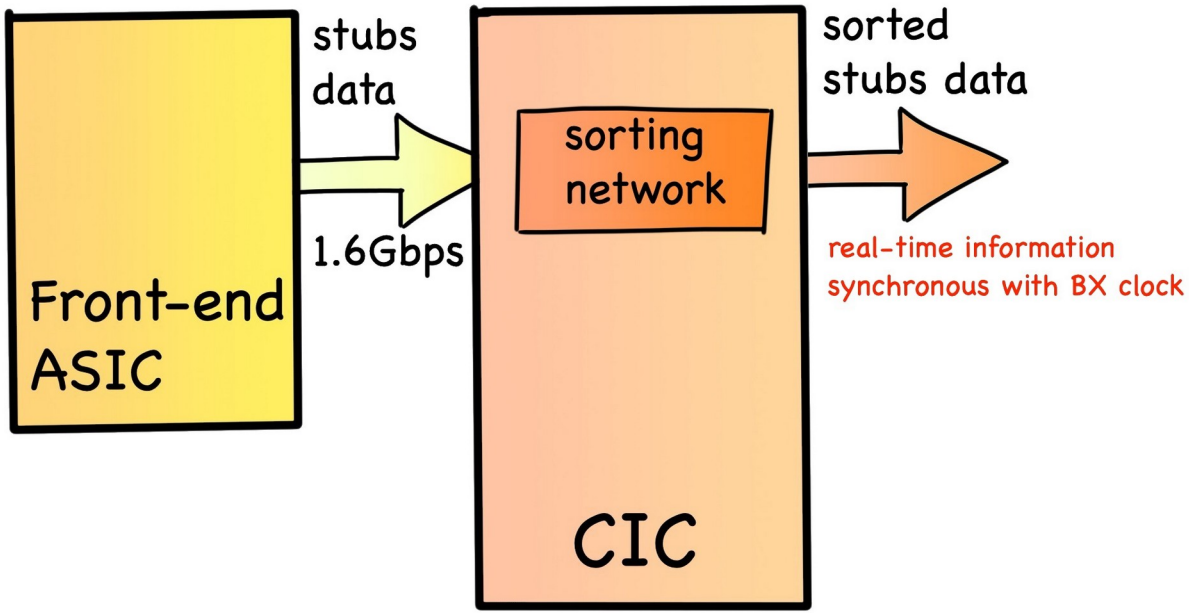

Data digital readout

all readout asics participate to the p_T-module concept:

"intelligent" discrimination is (digitally) performed on-chip by MPA, SSA, CBC and CIC

- stubs data are created at every BX by the front-end asics
  - this is the real-time information used for L1 trigger decision
- stubs data need processing before being sent to the back-end
  - CBC data need zeros-suppression
  - module output bandwidth needs to be efficiently exploited sending out only best event candidates
- upon reception of a **L1 accept (L1A)** signal the full-event image of the detector at a specific BX i sent out by the front-end asics

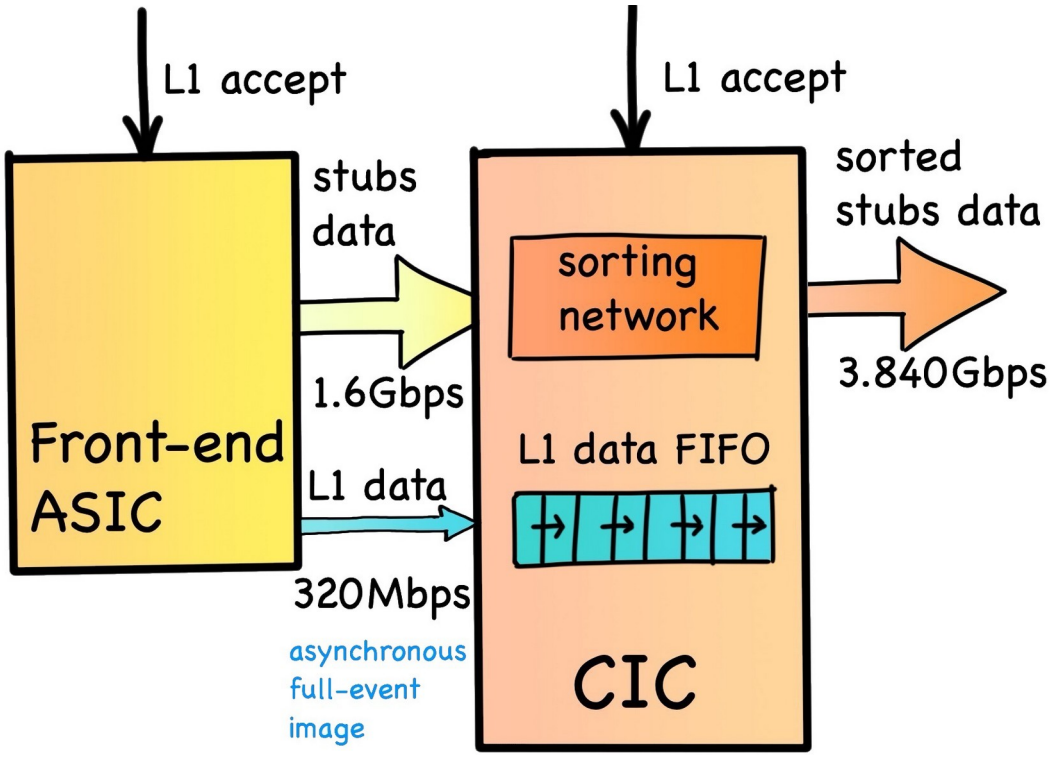

Data digital readout

all readout asics participate to the p_T-module concept:

"intelligent" discrimination is (digitally) performed on-chip by MPA, SSA, CBC and CIC

- stubs data are created at every BX by the front-end asics
  - this is the real-time information used for L1 trigger decision
- stubs data need processing before being sent to the back-end
  - CBC data need zeros-suppression
  - module output bandwidth needs to be efficiently exploited sending out only best event candidates
- upon reception of a **L1 accept (L1A)** signal the full-event image of the detector at a specific BX is sent out by the front-end asics
  - L1A signal rate follows a Poisson distribution centered at 750kHz
  - 16 front-end asics (8 chips per CIC) are required per p_T-module
  - on chip storage (at front-end and CIC level) is required to cope with the limited output bandwidth of the p_T-modules and with the asynchronous nature of the trigger arrival times

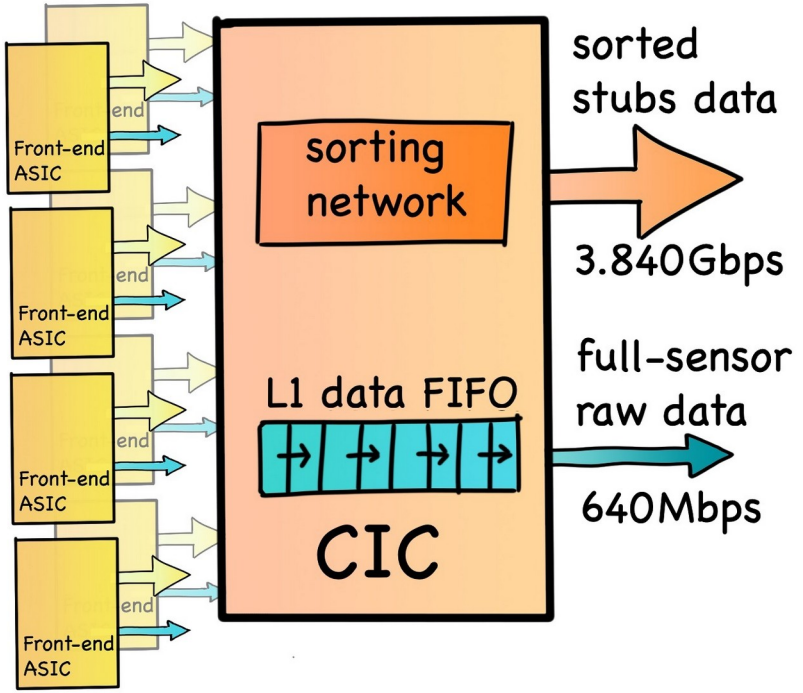

Data digital readout

all readout asics participate to the p_T-module concept:

"intelligent" discrimination is (digitally) performed on-chip by MPA, SSA, CBC and CIC

- stubs data are created at every BX by the front-end asics
  - this is the real-time information used for L1 trigger decision
- stubs data need processing before being sent to the back-end
  - CBC data need zeros-suppression
  - module output bandwidth needs to be efficiently exploited sending out only best event candidates
- upon reception of a **L1 accept (L1A)** signal the full-event image of the detector at a specific BX is sent out by the front-end asics
  - L1A signal rate follows a Poisson distribution centered at 750kHz
  - 16 front-end asics (8 chips per CIC) are required per p_T-module
  - on chip storage (at front-end and CIC level) is required to cope with the limited output bandwidth of the p_T-modules and with the asynchronous nature of the trigger arrival times

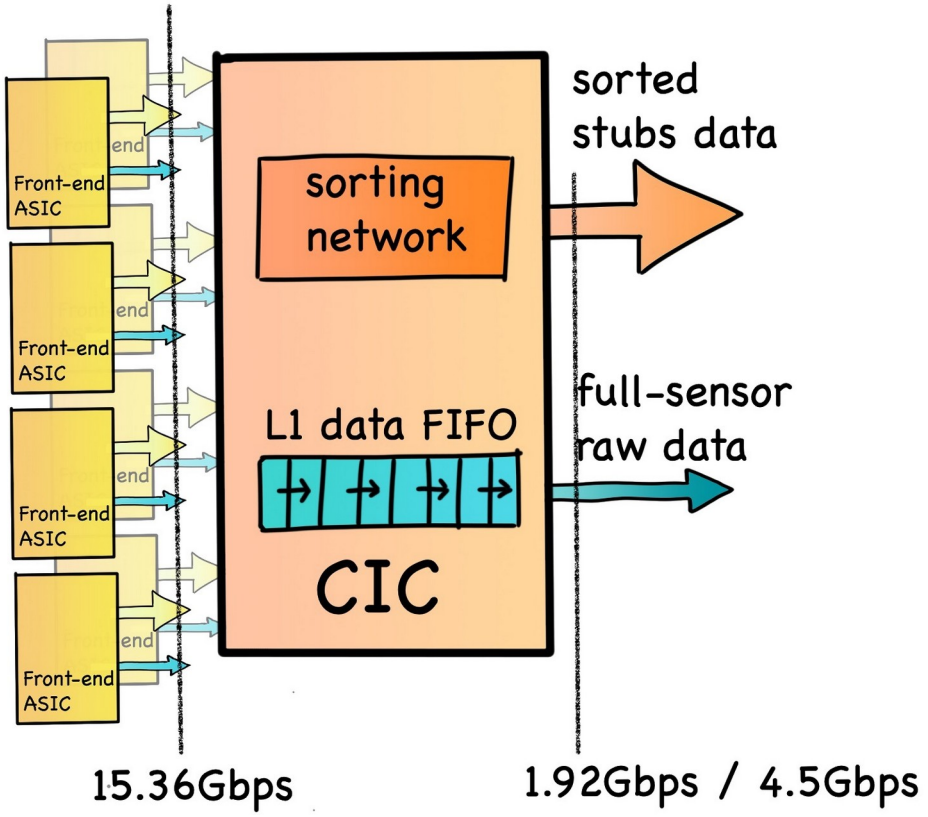

- CIC allows a data compression factor between 3 and 8 for the data produced by the front-end asics

CIC block diagram (stubs path)

- front-end asics (MPA and CBC) have different stub rates per BX and different data formats

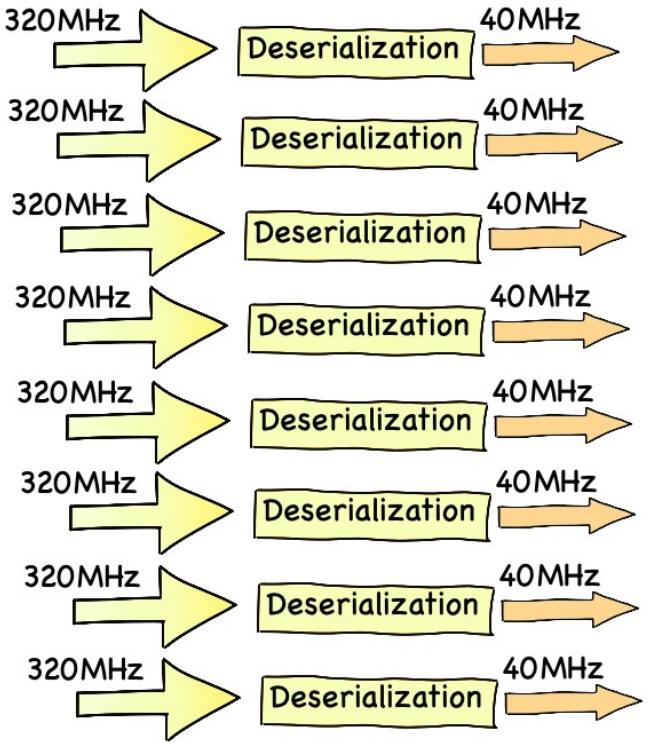

CIC block diagram (stubs path)

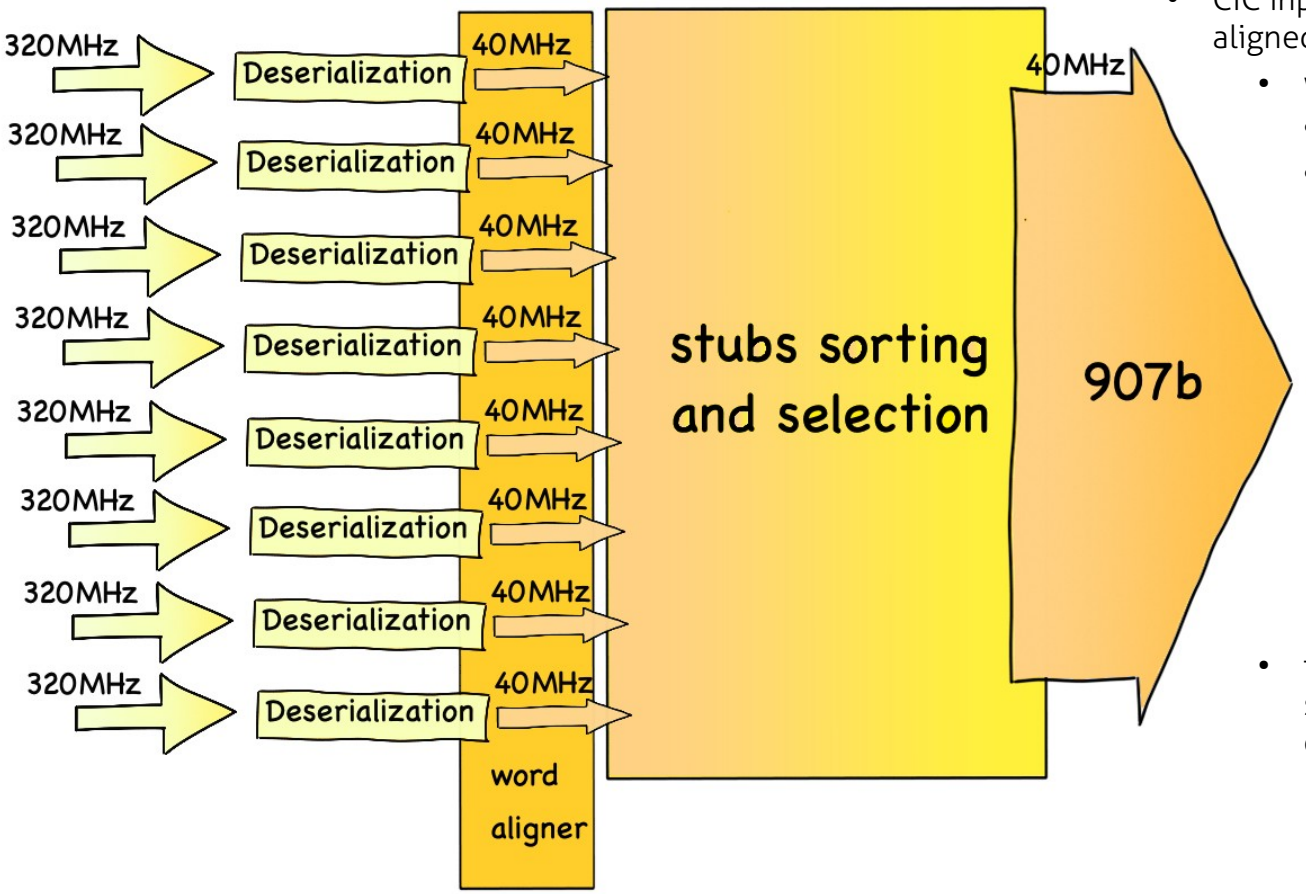

- front-end asics (MPA and CBC) have different stub rates per BX and different data formats
- CIC input deserialized data need to be word aligned per front-end channel
  - within a period of 8 BXs, stubs are sorted according to their bending: high-p_T values are prioritized
- the sorting network implements a bitonic sort algorithm and it's implementation is critical for the timing closure

CIC block diagram (stubs path)

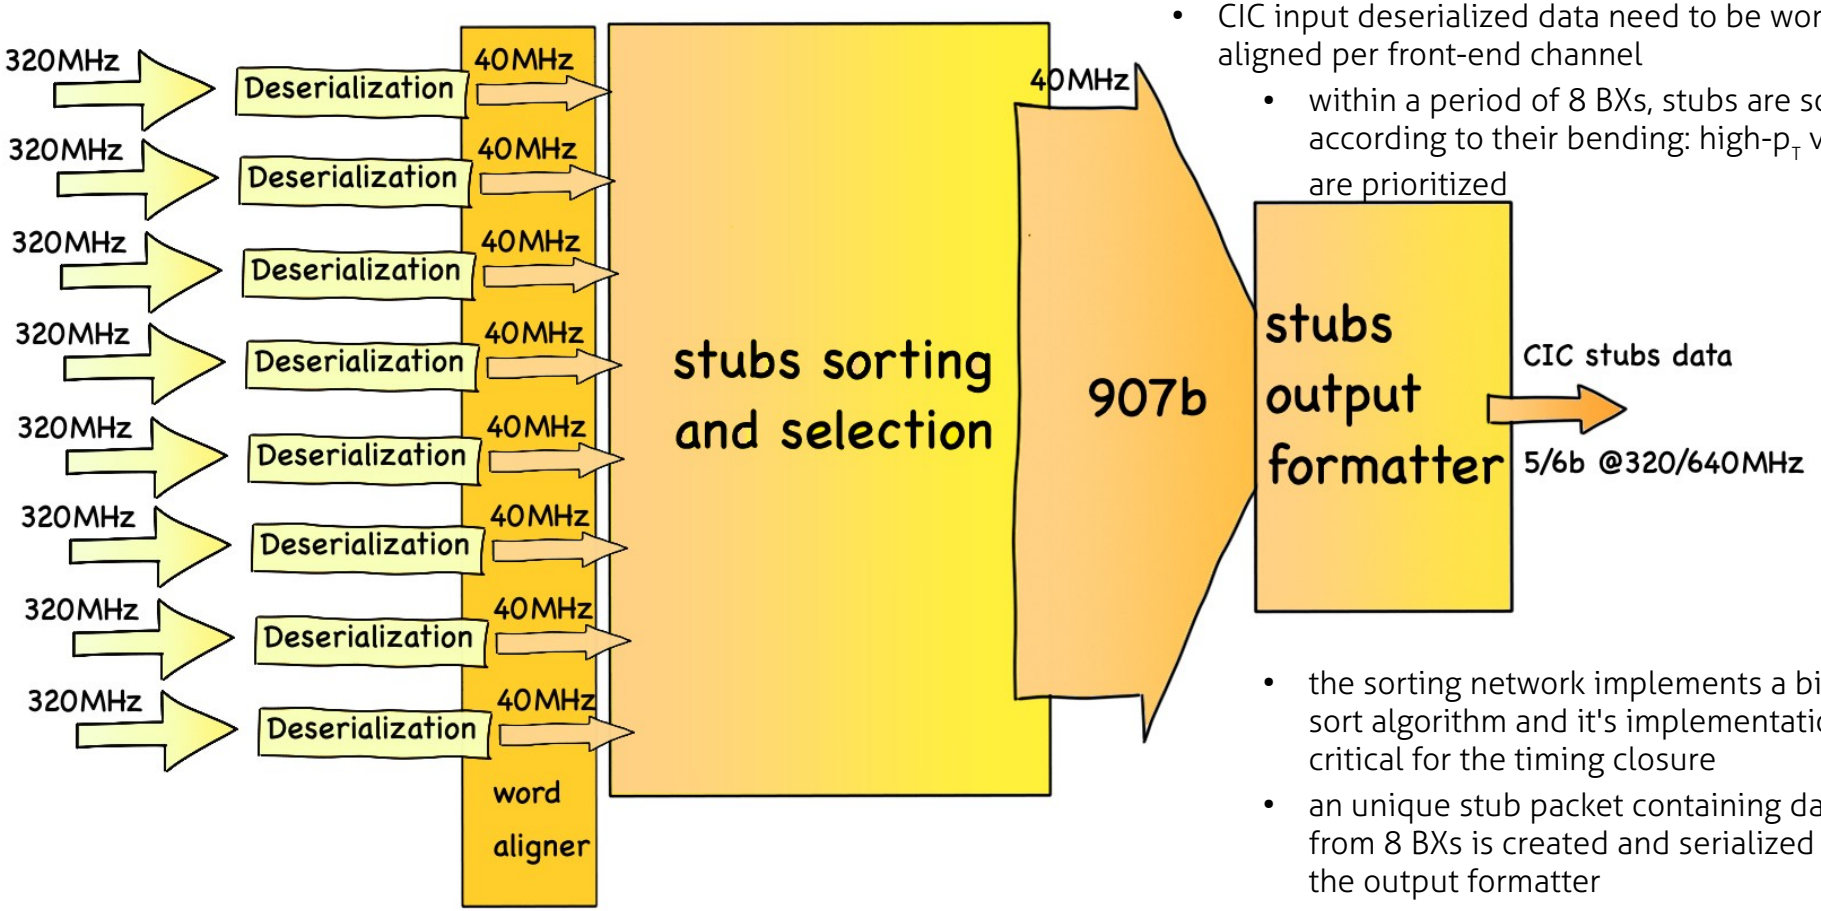

- front-end asics (MPA and CBC) have different stub rates per BX and different data formats
- CIC input deserialized data need to be word aligned per front-end channel
  - within a period of 8 BXs, stubs are sorted according to their bending: high-p_T values are prioritized
- the sorting network implements a bitonic sort algorithm and it's implementation is critical for the timing closure
- an unique stub packet containing data from 8 BXs is created and serialized by the output formatter

CIC block diagram (L1 path)

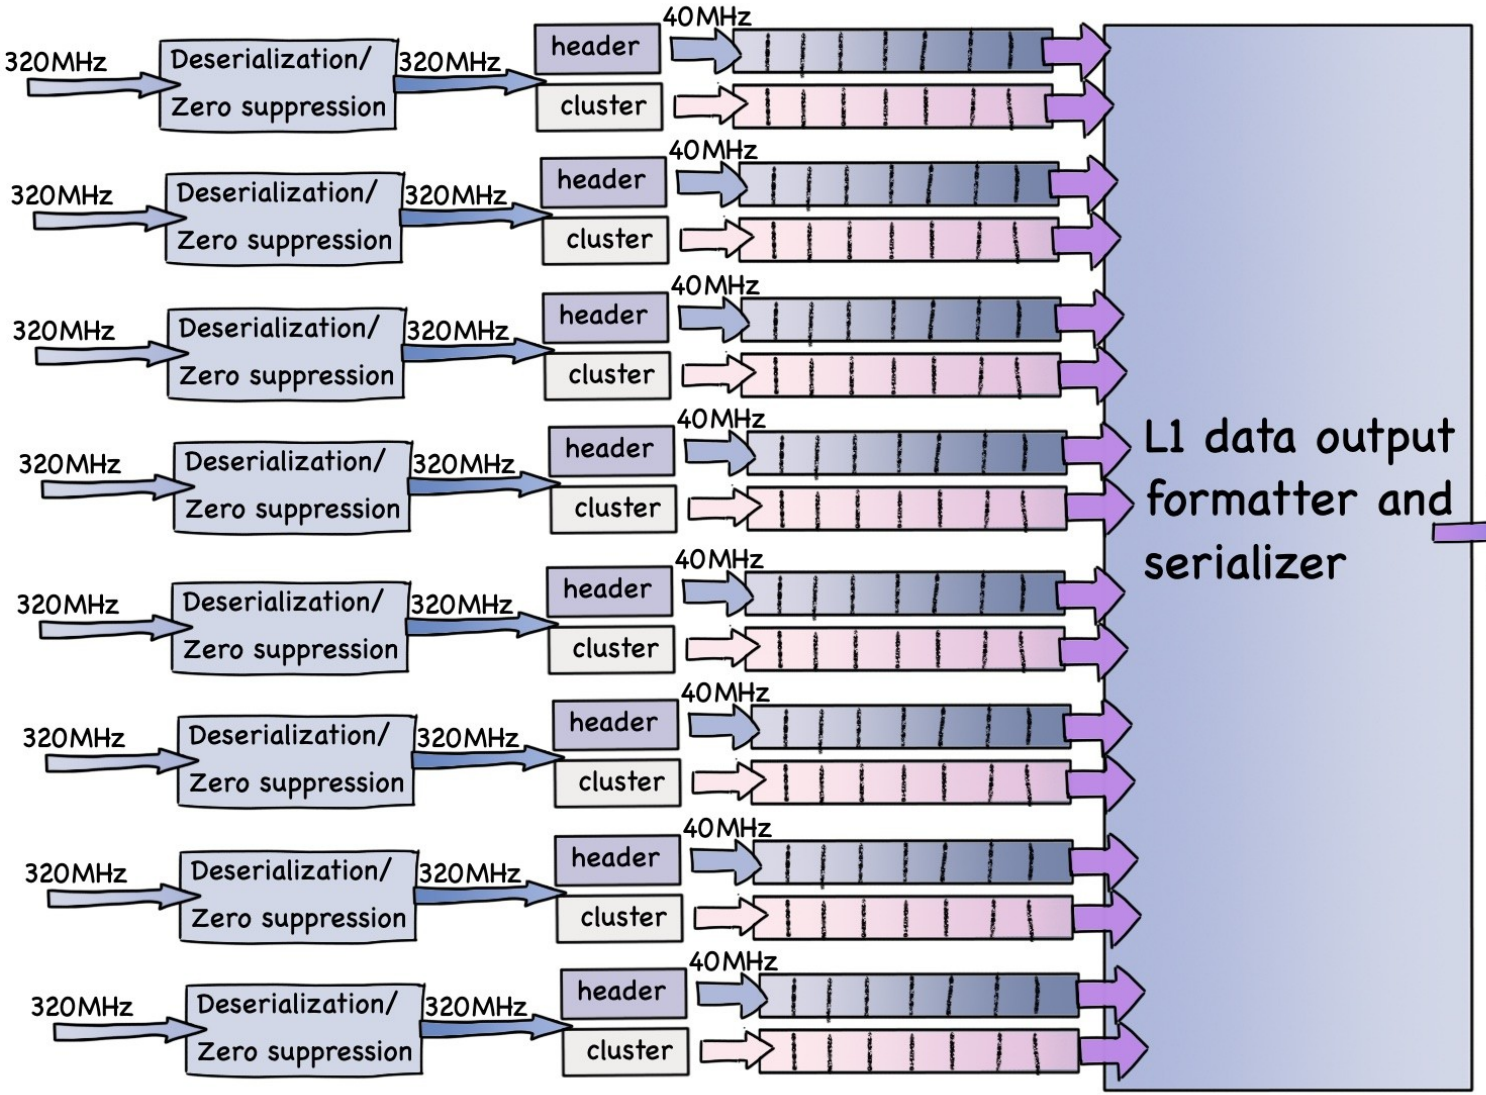

- 1b data lines at 320Mbps
- zero suppression for CBC data (it requires 320MHz clock)
- FIFO storage due to the limited output bandwidth

- FIFO dimensioning: studied and optimized for the full readout chain with detector simulated data and within the verification framework
- FIFOs are triplicated (only for the synchronization-critical data)

CIC block diagram (L1 path)

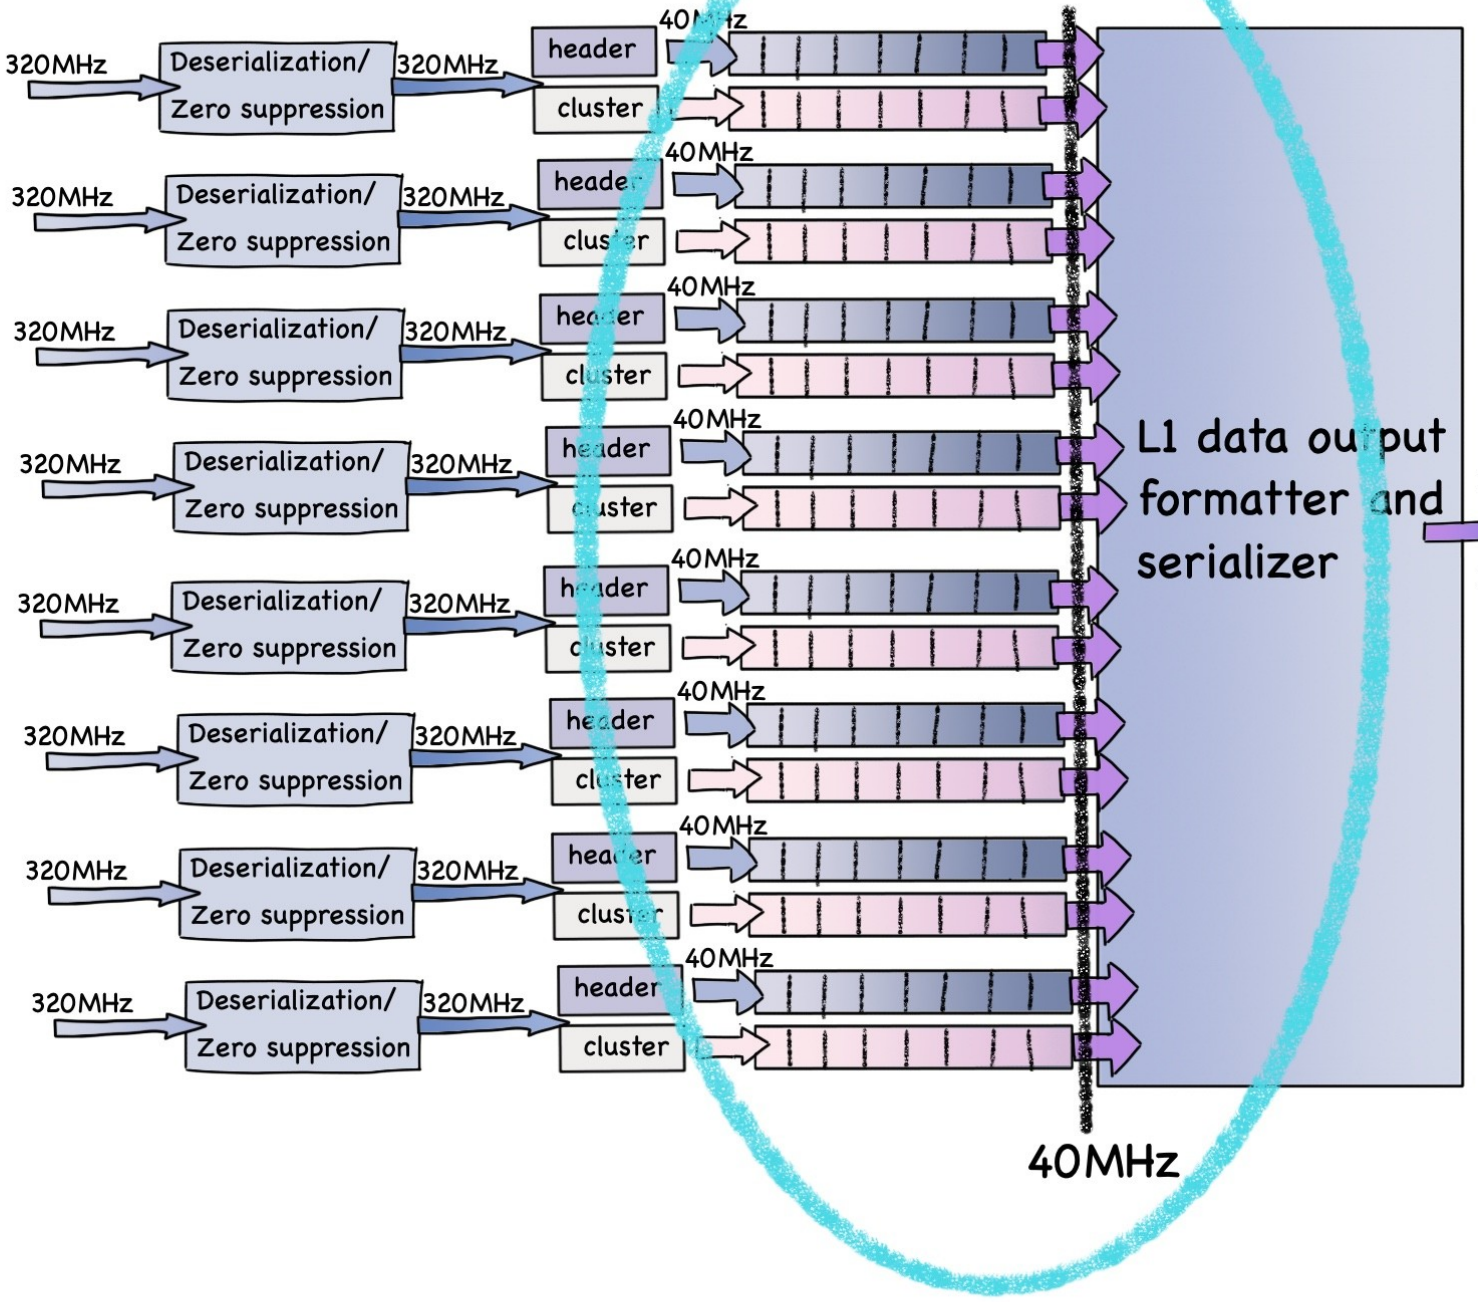

- 1b data lines at 320Mbps
- zero suppression for CBC data (it requires 320MHz clock)
- FIFO storage due to the limited output bandwidth

- FIFO dimensioning: studied and optimized for the full readout chain with detector simulated data and within the verification framework
- FIFOs are triplicated (only for the synchronization-critical data)
- careful optimization of FIFOs readout to limit power consumption

CIC block diagram

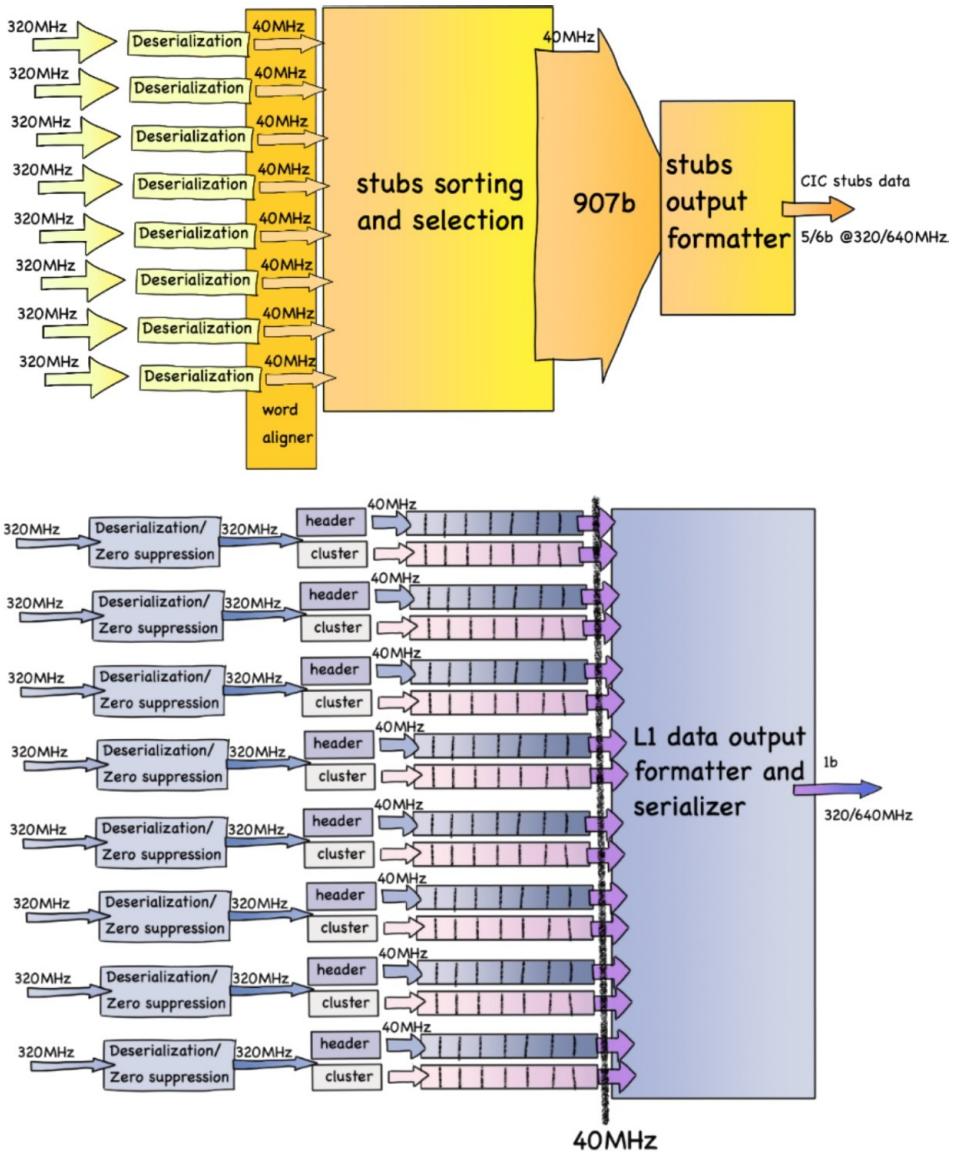

- independent L1 and stubs data paths
- phase-aligning with respect to the internal clock required for input data bits

CIC block diagram

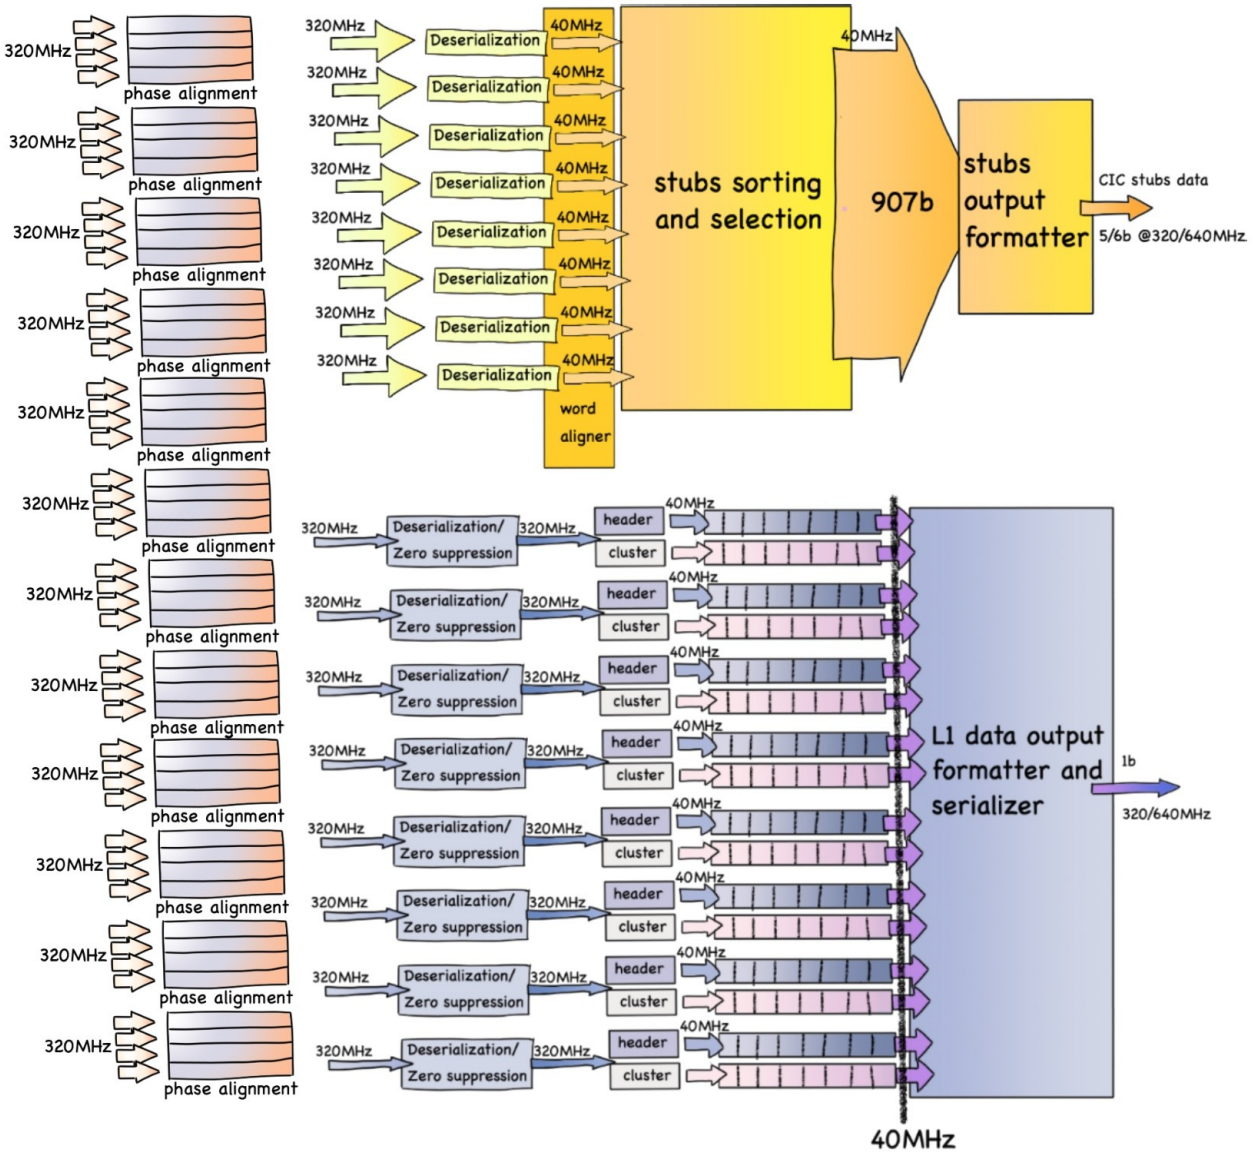

- independent L1 and stubs data paths
- phase-aligning with respect to the internal clock required for input data bits:

  analog macrocell design by SMU University for the lpGBT project

CIC block diagram

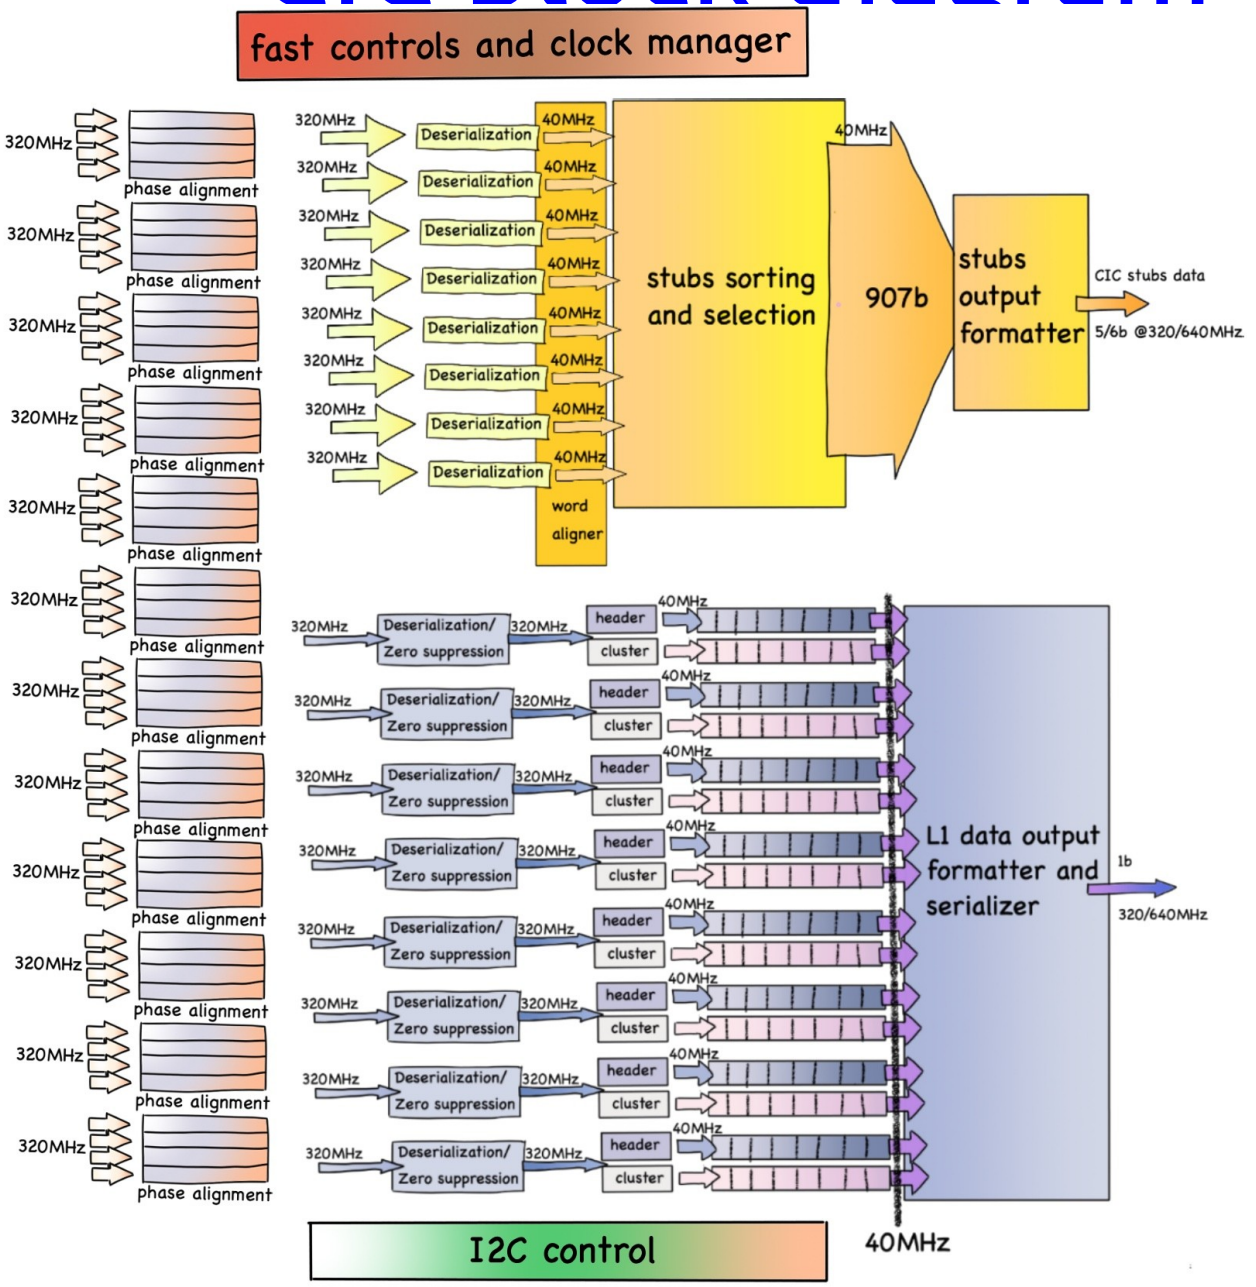

- independent L1 and stubs data paths
- phase-aligning with respect to the internal clock required for input data bits:
  analog macrocell design by SMU University for the lpGBT project
- switchable (320MHz and 640MHz) system clock
- configuration through I2C interface: the internal registers are clock-gated and triplicated

CIC physical implementation

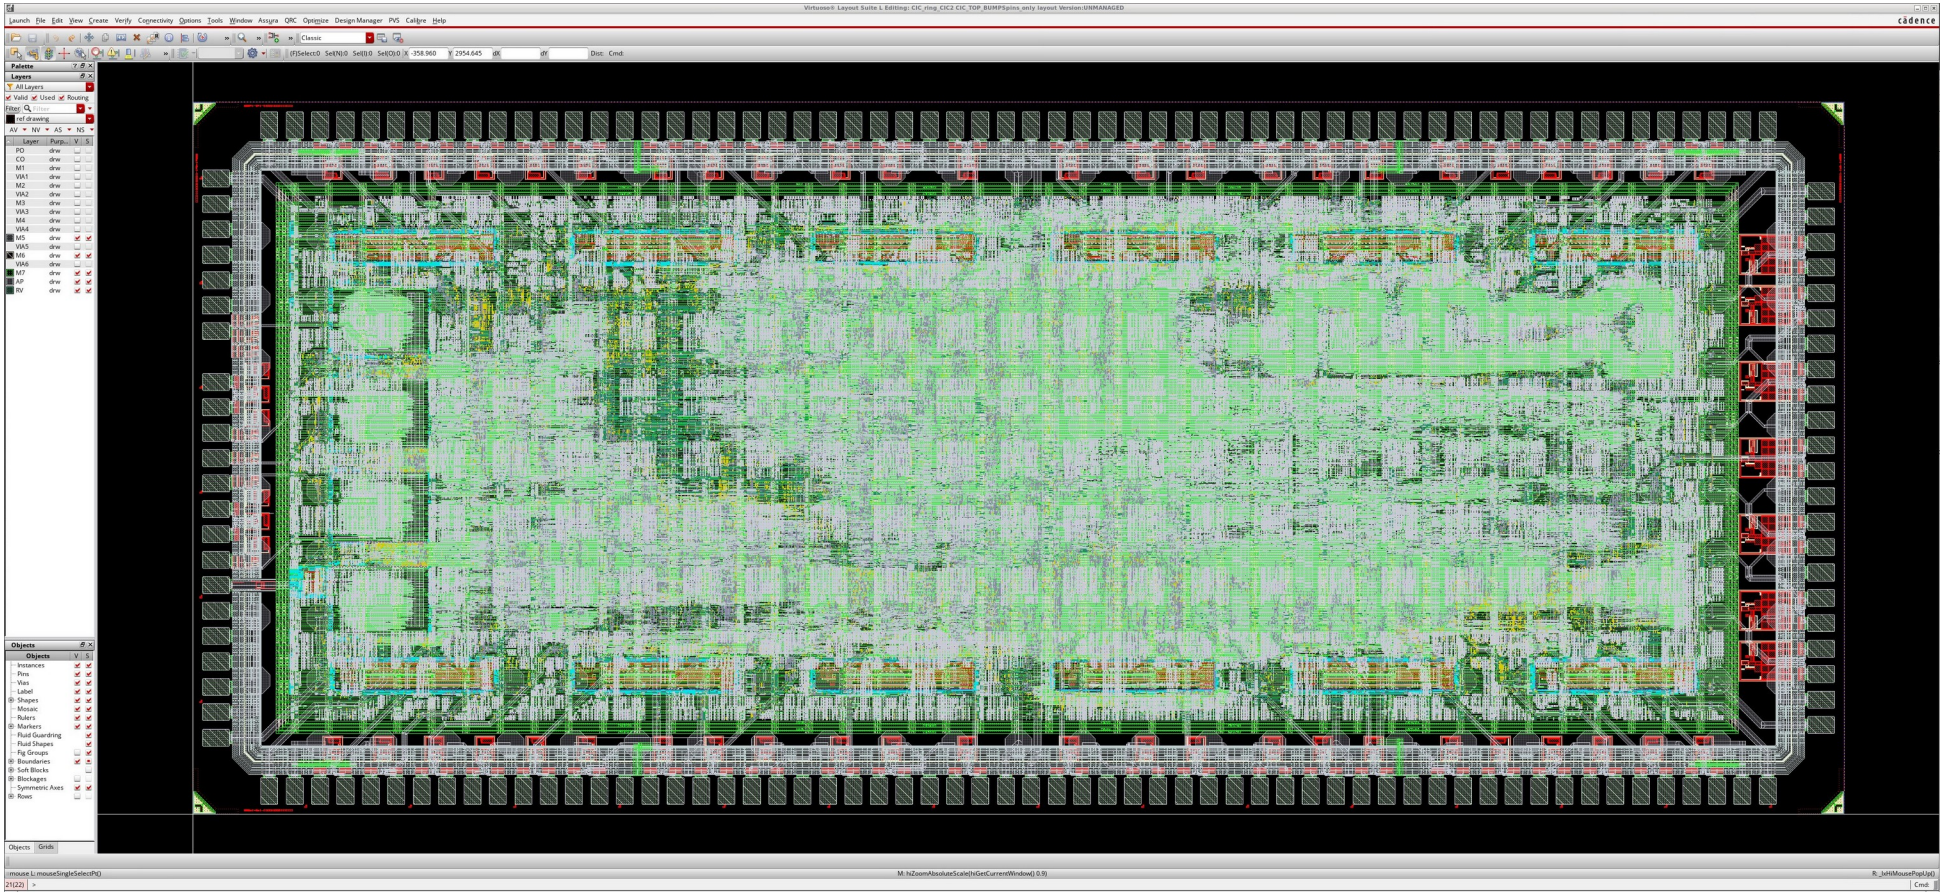

Digital on Top (innovus) flipchip implementation in tsmc65 node (1p7m option): 6500um x 2800um

- different CORE supply voltages for different 2S and PS use (1.2V and 1.0V)
- configurable input system clock (320MHz / 640MHz) and output data rate
  - phaseAligner (analog macro designed @SMU) + ePortRxGroup (digital IP designed @CERN and modified to our specs)
  - sLVS receiver + transmitter cells (designed @INFN BG-PV): used in periphery ring
  - radiation tolerant ESD protections (designed by SOFICS): used in periphery ring
  - eFuse macro (TSMC, triplicated and modified by the CERN team)

CIC physical implementation

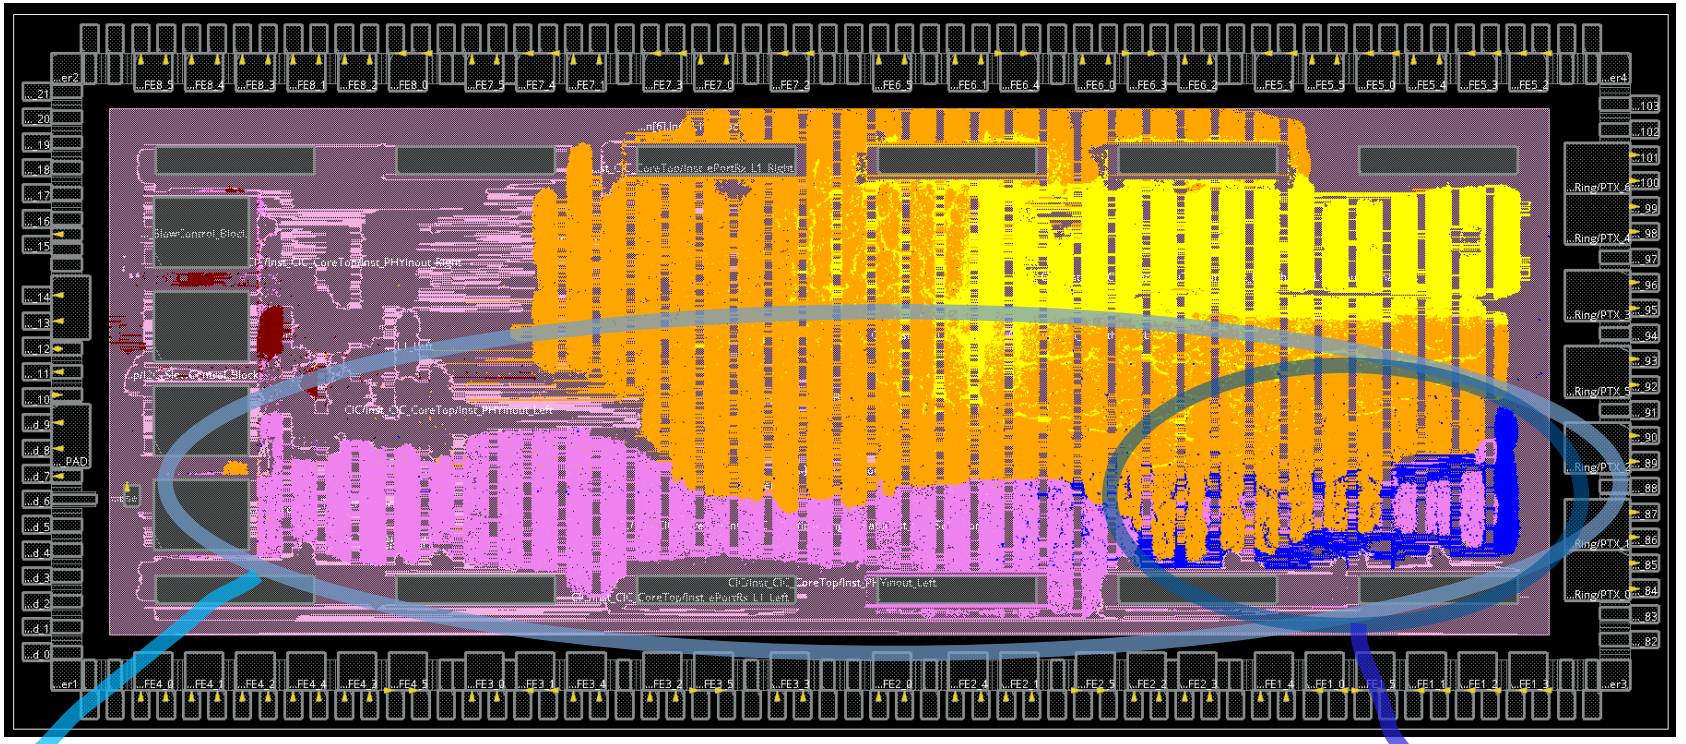

stubs data path

stubs output formatter

- it contains the stub selection and sorting unit (most critical for timing)
- almost only combinatorial logic working @ 40MHz
- Cells ~ 79K
  (StubSelection: Cells ~ 44K)

CIC physical implementation

L1 data path: largest contribution (~80%) to the overall power consumption

L1 data output formatter and serializer

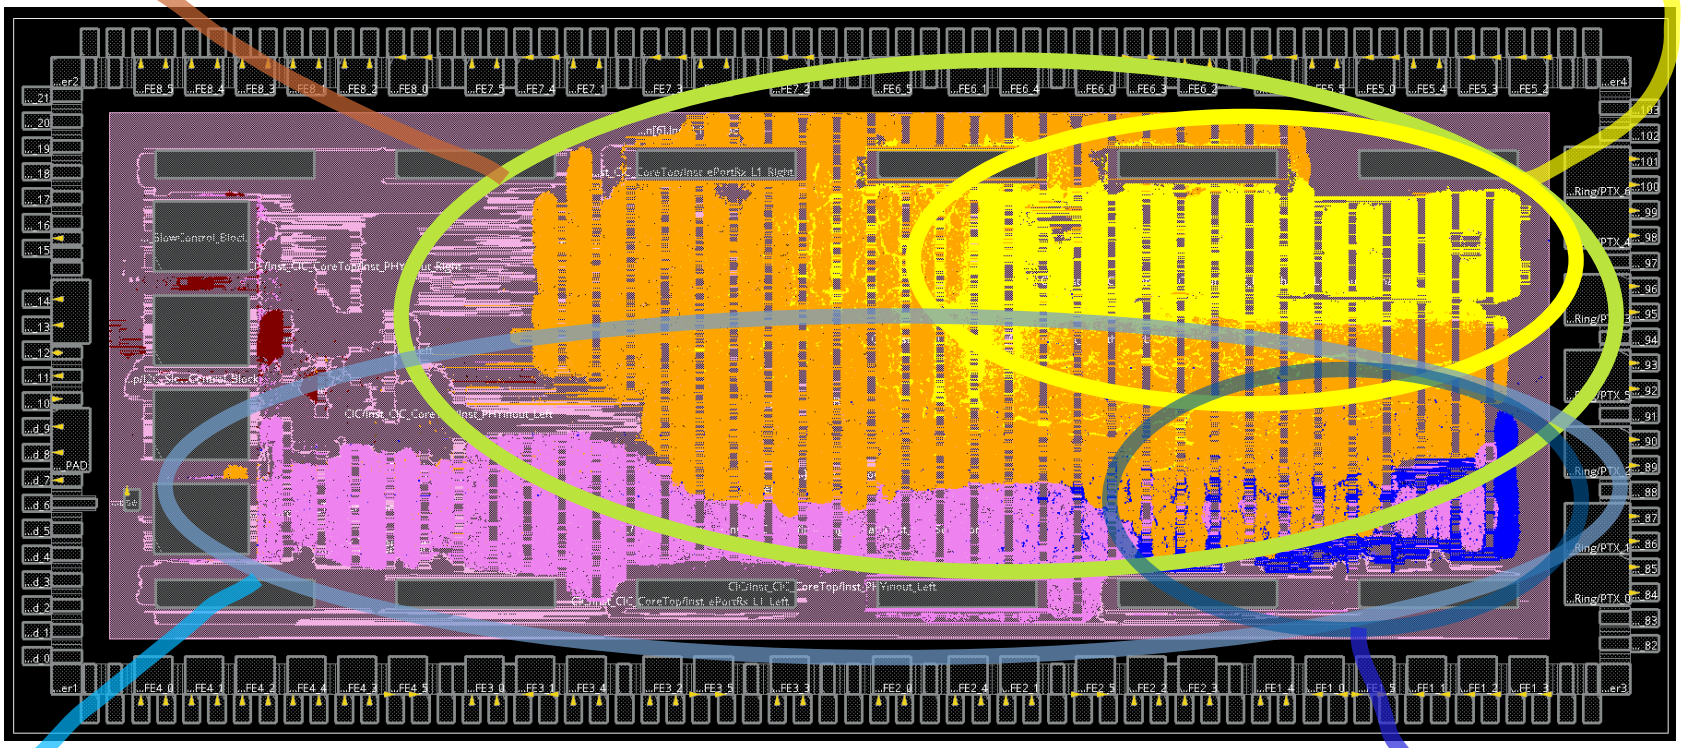

stubs data path

stubs output formatter

- it contains the stub selection and sorting unit (most critical for timing)
- almost only combinatorial logic working @ 40MHz
- Cells ~ 79K
  (StubSelection: Cells ~ 44K)

CIC prototype tests and characterization

functional test:

stand-alone testbench based on a commercial FPGA card + asic mezzanine

- front-end data are software generated and stored in the RAM of the FPGA card: data are then sent to the CIC through an interface board and the outputs can be checked against the expected values through millions of BX cycles
- functional test (based on the verfication verilog testbench)
- power and temperature characterization

Module	$V_{core}(V)$	$P_{tot}(mW)$
PS-module @640MHz	0.9	123
	1.0	144
	1.1	169
2S-module @320MHz	1.1	156
	1.2	181
	1.3	210

200 pileup occupancy input data frames
750 kHz L1A trigger rate

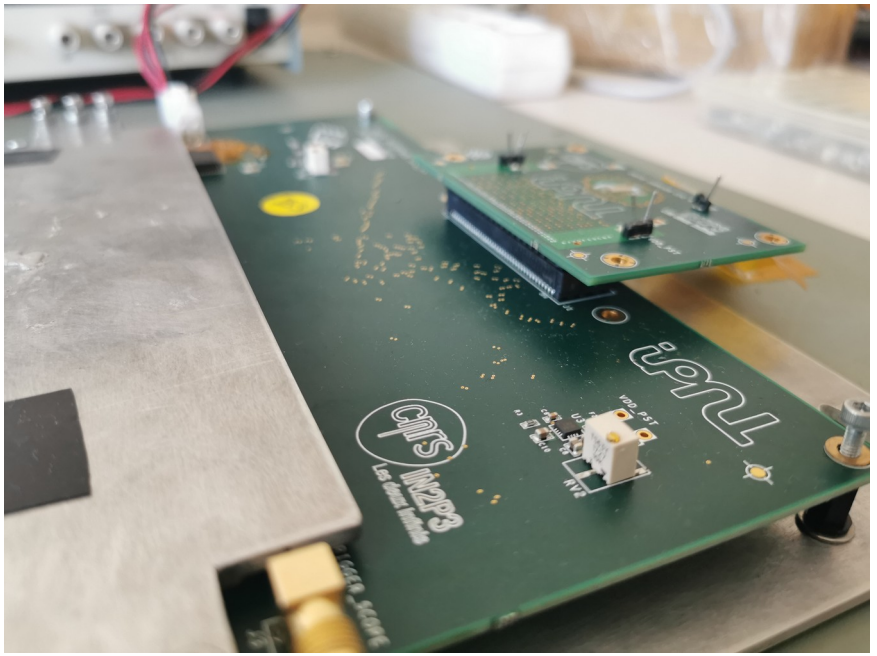

CIC test mezzanine and interface board
(stand-alone testbench configuration)

CIC prototype tests and characterization

functional test:

stand-alone testbench based on a commercial FPGA card + asic mezzanine

- front-end data are software generated and stored in the RAM of the FPGA card: data are then sent to the CIC through an interface board and the outputs can be checked against the expected values through millions of BX cycles
- functional test (based on the verfication verilog testbench)
- power and temperature characterization

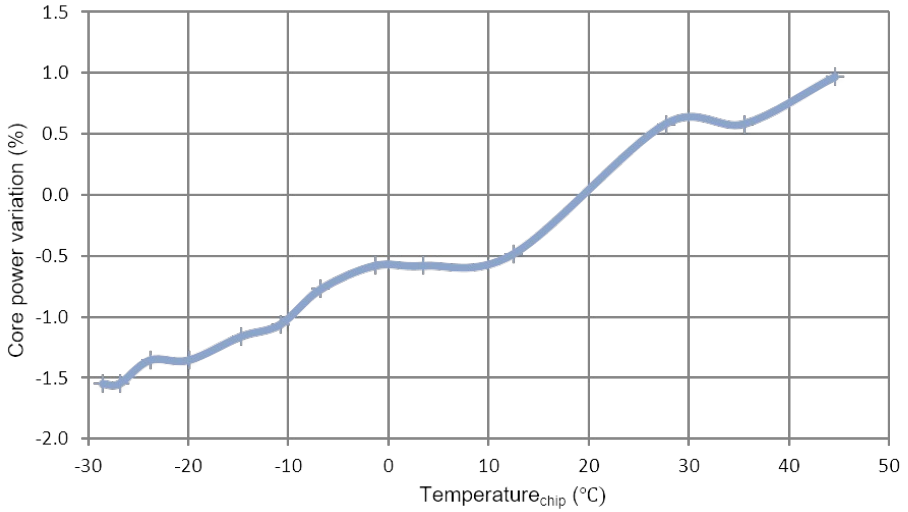

core power variation at different temperatures

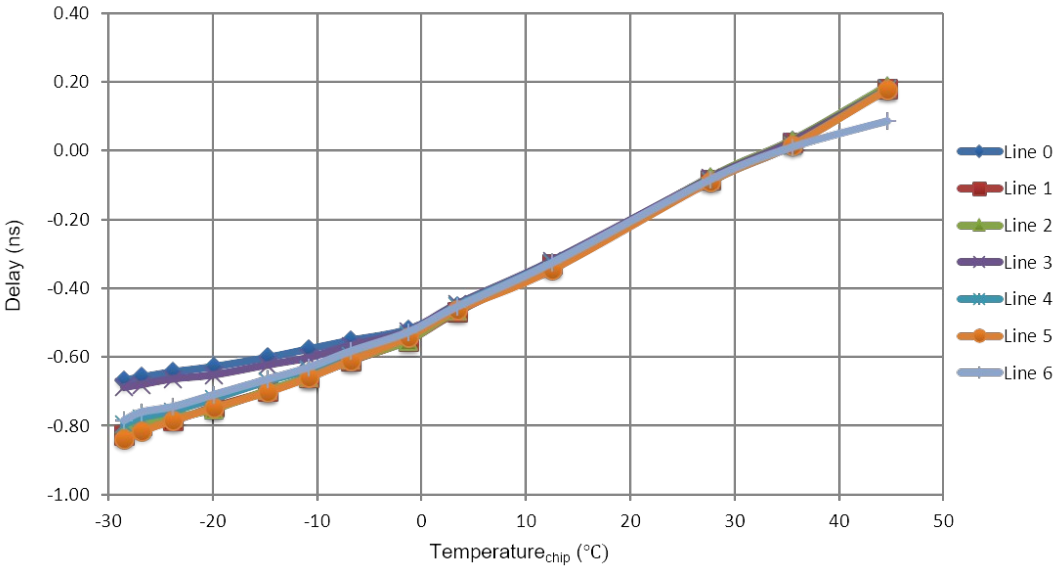

output delay variation at different temperatures

CIC prototype tests and characterization

functional test:

stand-alone testbench based on a commercial FPGA card + asic mezzanine

- front-end data are software generated and stored in the RAM of the FPGA card: data are then sent to the CIC through an interface board and the outputs can be checked against the expected values through millions of BX cycles
- functional test (based on the verfication verilog testbench)
- power and temperature characterization
- TID sensitivity:
  - increase of leakage current and cells propagation delay
  - irradiation tests with X-ray beam up to 200 Mrad (7.6 Mrad/h)

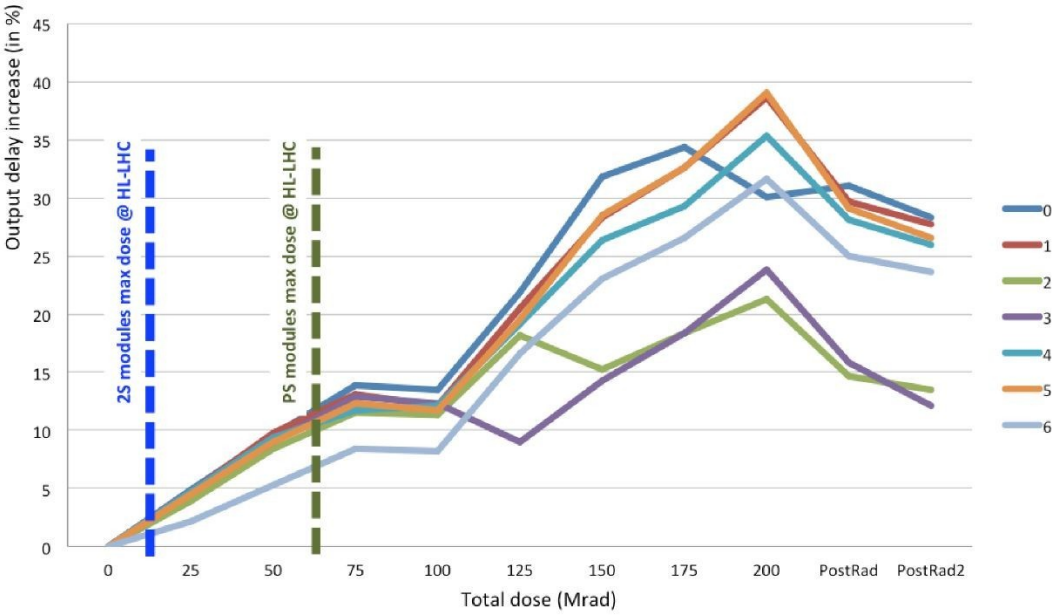

output delay variation wrt total irradiated dose

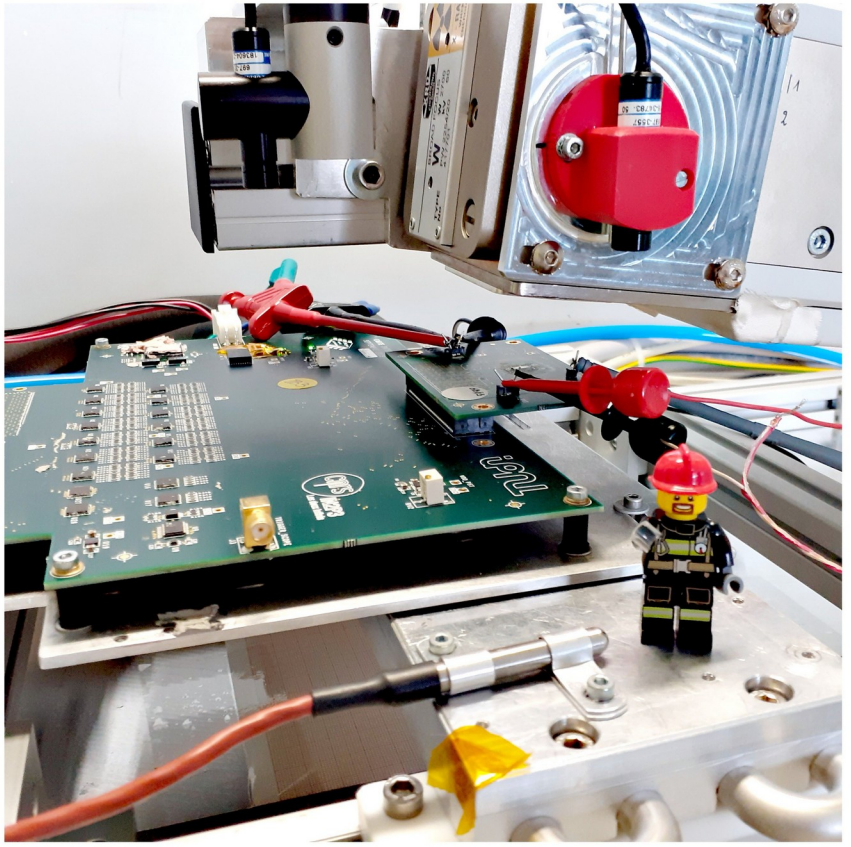

CIC bonded on test mezzanine and interface board in TID (X-ray beam) configuration at CERN

CIC prototype tests and characterization

functional test:

stand-alone testbench based on a commercial FPGA card + asic mezzanine

- front-end data are software generated and stored in the RAM of the FPGA card: data are then sent to the CIC through an interface board and the outputs can be checked against the expected values through millions of BX cycles
- functional test (based on the verfication verilog testbench)
- power and temperature characterization
- TID sensitivity
- SEU tests:
  - control paths and configuration registers are triplicated
  - SEU errors should be internally corrected by the logic: if correction fails, no blockages of the state machines should appear
  - configuration registers should never be upset
  - a small number of errors - not compromising functionality - is expected
  - maximum error rates due to SEU observed during the SEU test campaign were lower than the expected rates for the p_T-modules

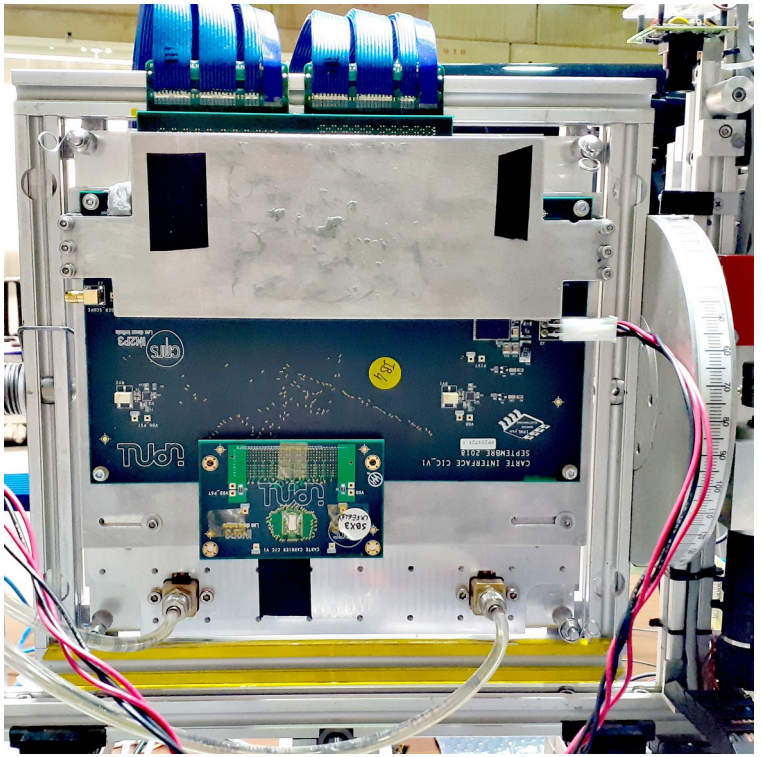

CIC bonded on test mezzanine and interface board with cooling plates in test beam configuration at Louvain

CIC project timeline

2012/2013 Project start (software simulations, input/output data formats definition, chip specs definition)

2015 (end) PDK 65nm available

2018 CIC1 1st submission

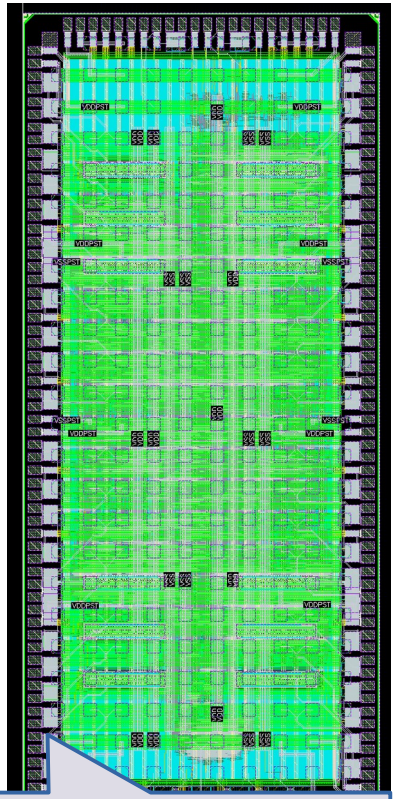

CIC1:
- 320MHz clock only
- untriplicated design
- october 2018 (wirebonded)
- may 2019 (bumped)

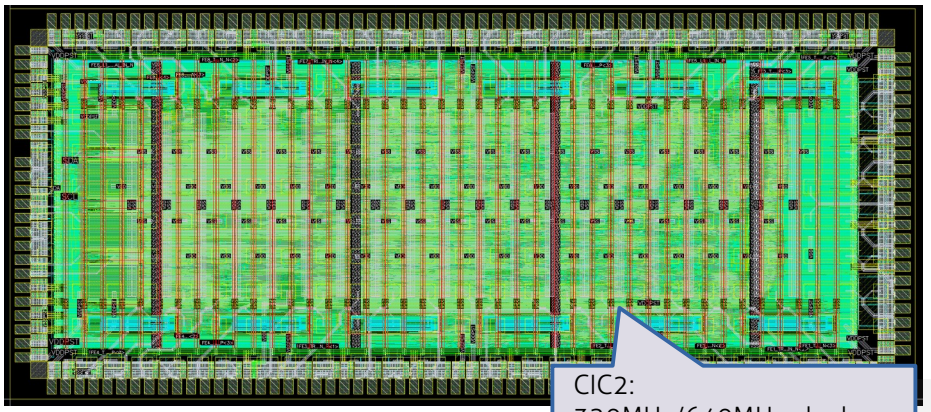

CIC2:
320MHz/640MHz clock
triplicated design
SEU+TID characterised
august 2019 (wirebonded)
february 2020 (bumped)

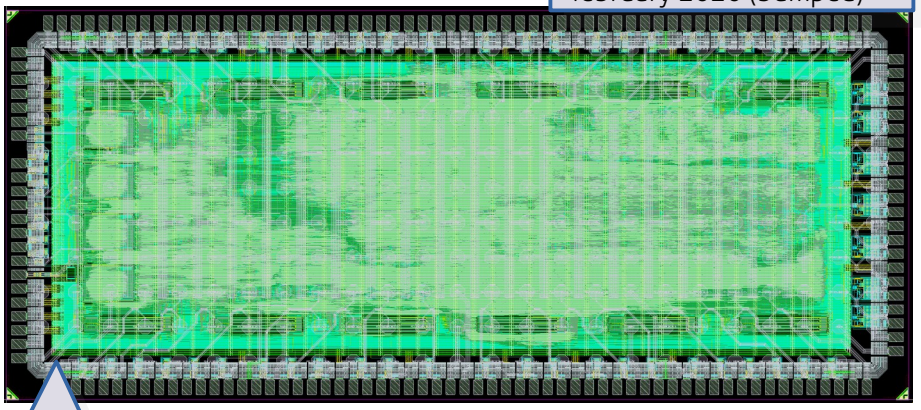

CIC2.1

2019 CIC2 submission

2020 CIC2 characterization campaign (TID, SEU)

CIC2 integration in PS module

2021 implementation respin (CIC2.1)

2021 (sept.) Launch of production

CIC1 implements all functionalities without the 640MHz clock mode and without SEU-resistent solutions

in the CIC2 and CIC2.1 the control structures and the reset lines are triplicated (not the data paths)

Summary

- the Concentrator IC (CIC) is part of the digital readout chipset developed for the p_T-modules of the new phase-II Outer Tracker of the CMS experiment
- the chip main purpose is to utilise more efficiently the bandwidth of the links needed to readout the OT: it achieves this goal by compressing and merging data coming from a group of 8 front-end asics populating the same hybrid
- two CICs are used per p_T-module for a total production quantity of about 30K parts
- the chip is a complex digital design implemented in a 65n technological node by exploitation of the Digital On Top methodology
- the second version of the design, the CIC2, has been successfully tested and characterized in 2020: a respin of the design with added features before the production has been nevertheless launched
- this CIC2.1 version will be tested and characterized in January 2022
- the mass production of the two chip's versions CIC2 and CIC2.1 has been launched in September

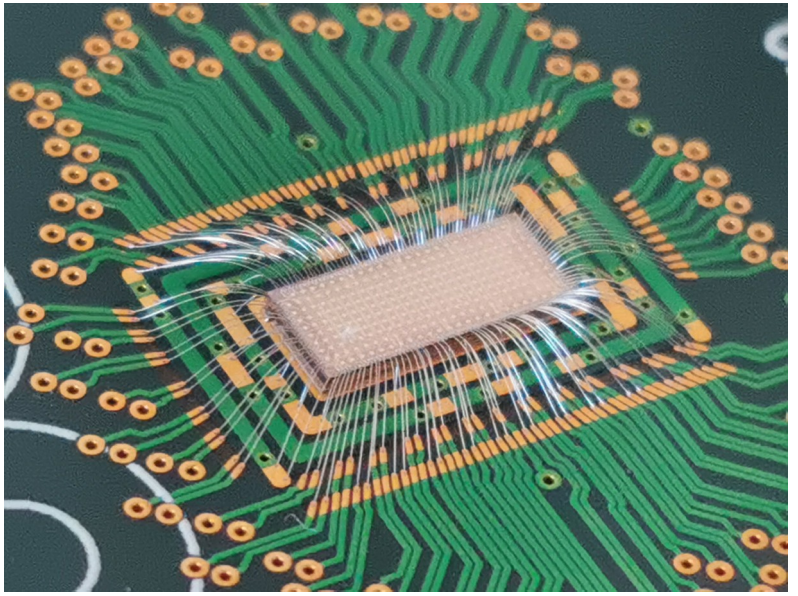

Thanks to F. Agnese, L.Gross and C. Wabnitz – IPHC Strasbourg for the bonding of the CIC2 prototypes

An important part of the project consisted in the definition of the data formats used by the four asics in the chipset and took contributions from several people all along the lifespan of the project.

There is strong interdependence between all the readout asics of the chipset: the CIC design benefited of the complex simulation and verification UVM framework developed by the CERN team for the identification of the best architecture.

The design team includes members of the IP2I as well as from CERN with various responsibilities and different implications in the project.

All of them equally contributed to the project and I'd like to express to them my gratitude here

D. Contardo, G. Galbit, C. Guerin, B. Nodari, S. Jain, W. Tromeur, S. Viret, Y. Zoccarato - IPN Lyon
G. Bergamin, A. Caratelli, D. Ceresa, K. Kloukinas, S. Scarfi' - CERN