The new readout board for LHCb : PCIe400

P. Bibron, CPPM
J.-P. Cachemiche, CPPM

Contents

LHCb experiment

Constraints

FPGA overview

Architecture

Foreseen development

Core Power Estimation

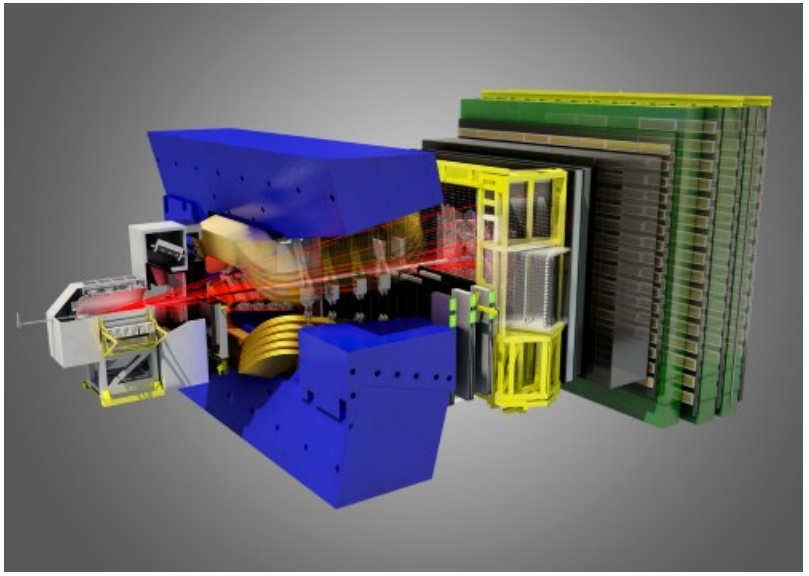

LHCb experiment

LHCb goal is to study Matter/Antimatter asymetries

Curent system is based on a Trigger-less architecture

- All data are sent to computer farm
- Real time event builder without pre-filtering
- More accurate algorithms can target the most interesting events

Interface between Front End and servers is key

- Current board PCIe40 has 100Gbit/s server interface
- 2030 upgrade is aiming 10 times higher
- First step is 4 times faster for 2026
    - 400 Gbits/s link with CPU
    - Input bandwidth minimum 480 Gbits/s, more if possible

This new board is named PCIe400

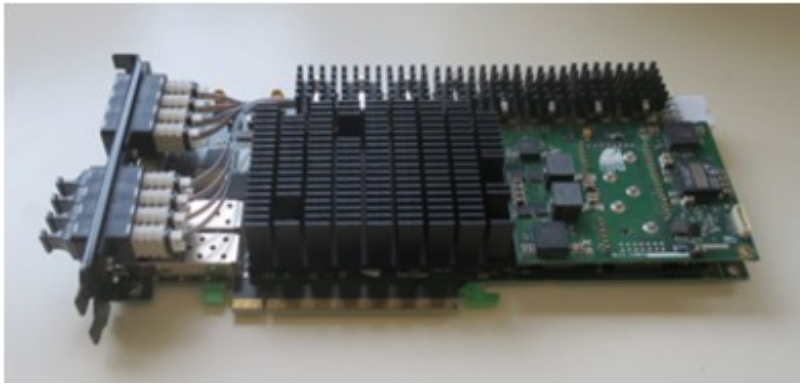

FPGA Chosen : Intel AgileX I-Series

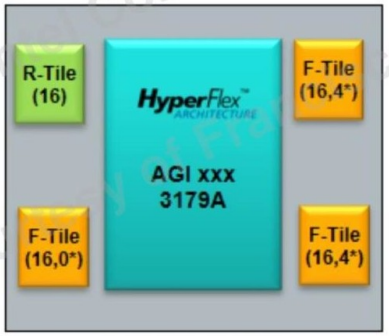

3x F-tile, 1x R-tile
(72 XCVRs)

Notion of Tiles for Transceivers

- Usually, FPGAs have a very high speed links on few links (~ 28, 56, 112 Gbits/s), whereas we need many links at a rather « low » speed (~10 Gbits/s)

- F-Tiles - for Front End : **48 NRZ links or (exclusive) 36 PAM4**
- R-Tiles - for CPU link : **PCIe Gen5 x 16**

Number of Logic Elements

- Hasn't increased as much as in the past (2.7MLEs Max Currently)

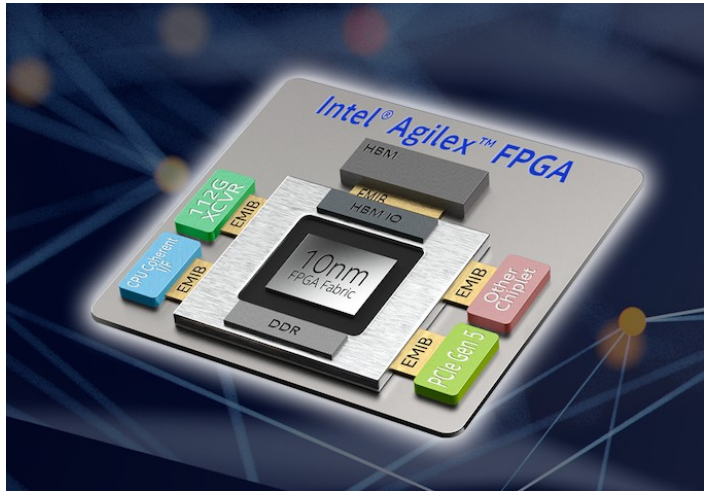

Architecture :

Same principle as PCIe40
PCIe board Gen5

- 48 x 10 = 480 Gb/s - Can be increased up to 2016 Gbits/s if PAM4 used
- 400 Gb/s **connection with CPU by PCIe Gen5** (acquisition + slow control)
- Slow control **through PCIe Gen5**
- Clock distribution over copper or through optics if needed

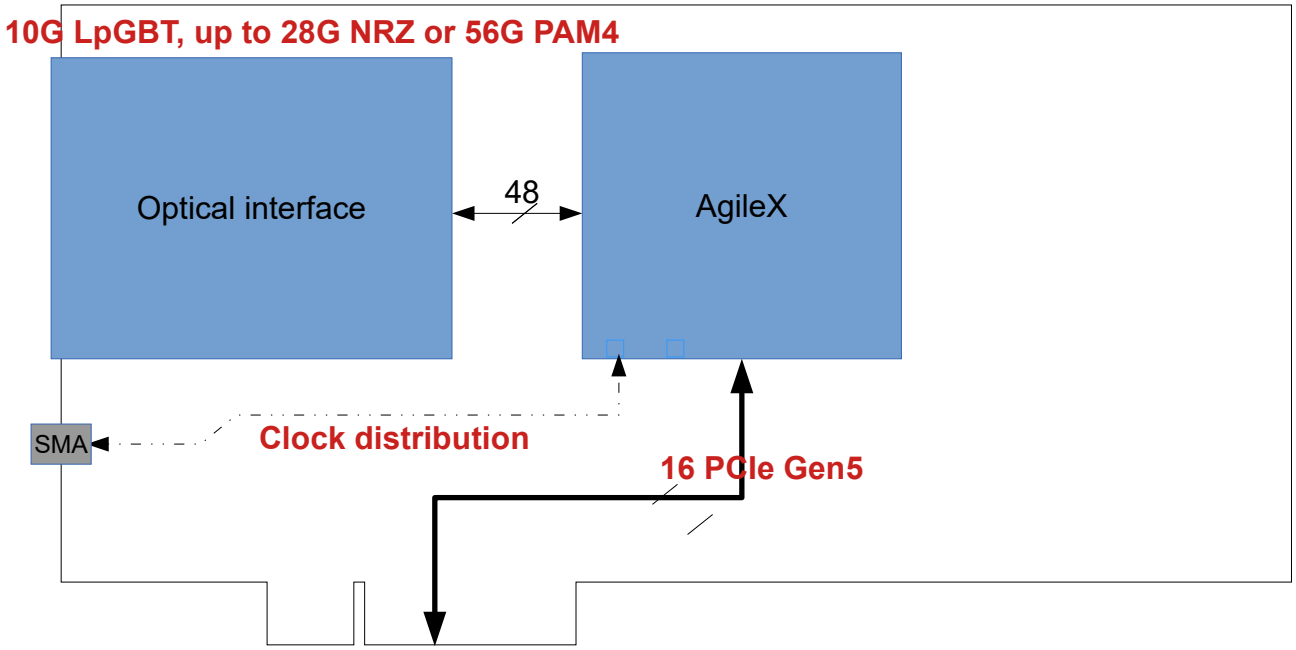

Foreseen development

CPPM has started general studies

- Power supplies
  - Core Power estimations
  - Voltage rail sequencing
  - Power delivery network
  - Power integrity simulations
- Accurate time distribution
  - PLL comparisons and choice
- 28/56 Gbits/s transmissions
  - Simulations
  - Implementation on Intel dev kit

Aim for a prototype in 2023

- Compatible with release of I-Series Agilex FPGA in Late 2021

Agilex core Power estimations

Agilex F-Series Development kit measurements

- Agilex F Series 1.4MLEs (AGFB014) – Not exactly our target (I-Series 2.7MLEs)
- Enpirion Core Power supply can go to 100/120A
- External PLLs can vary from 0 to 700MHz

- Monitoring (Current, Voltage & Temp) done externally by Max 10 FPGA
- Intel HMI provided for
  - PLL configuration
  - Voltage/Current/power monitoring
  - Temperature monitoring

How to make the FPGA consume power ?

- Depends on Occupancy/Frequency/Toggle rate
- Use of random Pattern generators (50 % Toggle rate) to fill up the FPGA

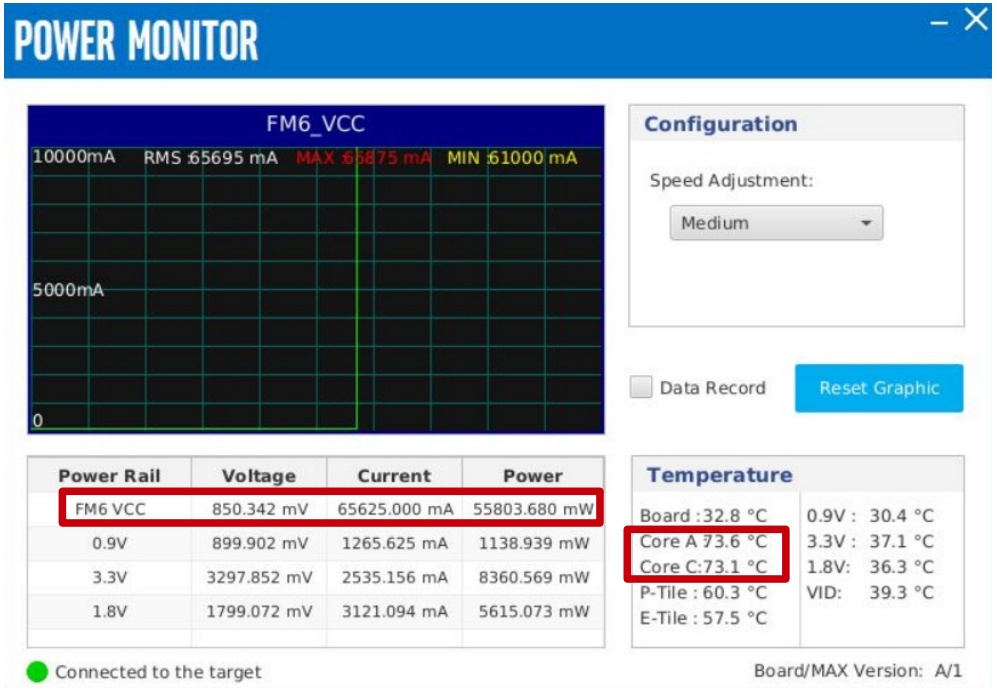

Agilex core Power estimations

Agilex F-Series Development kit measurements

- Core Current vs frequency and occupancy

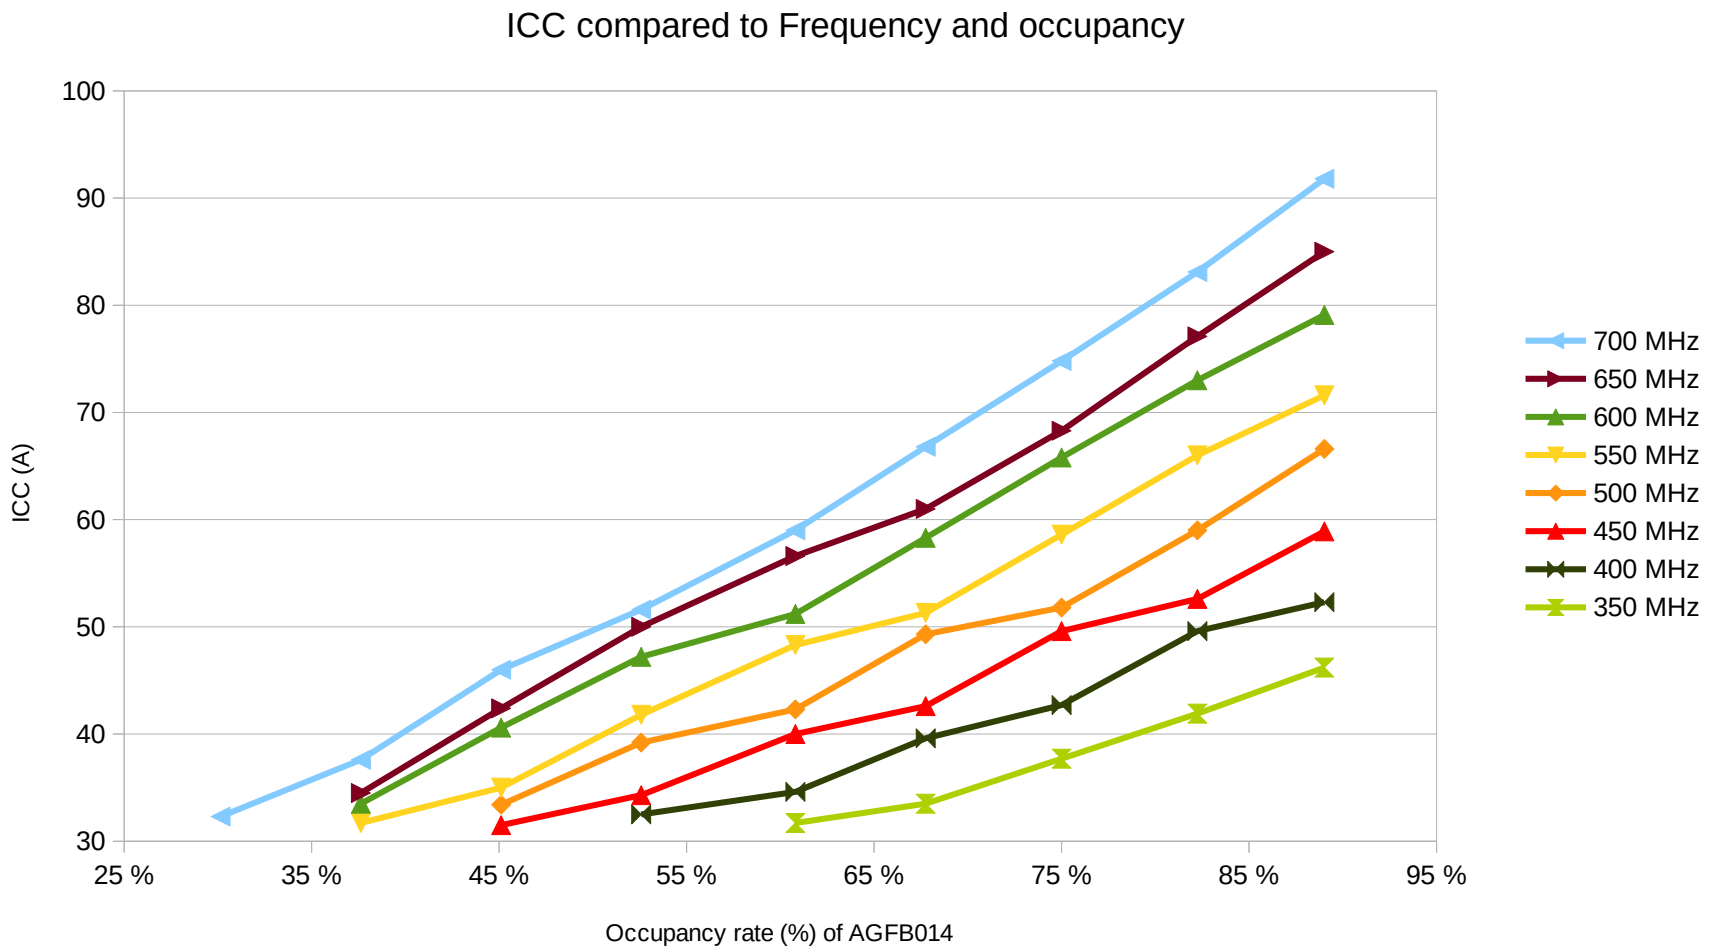

Agilex core Power estimations

Agilex F-Series Development kit measurements

- Extrapolation to target FPGA (from 1.4MLEs to 2.7MLEs)

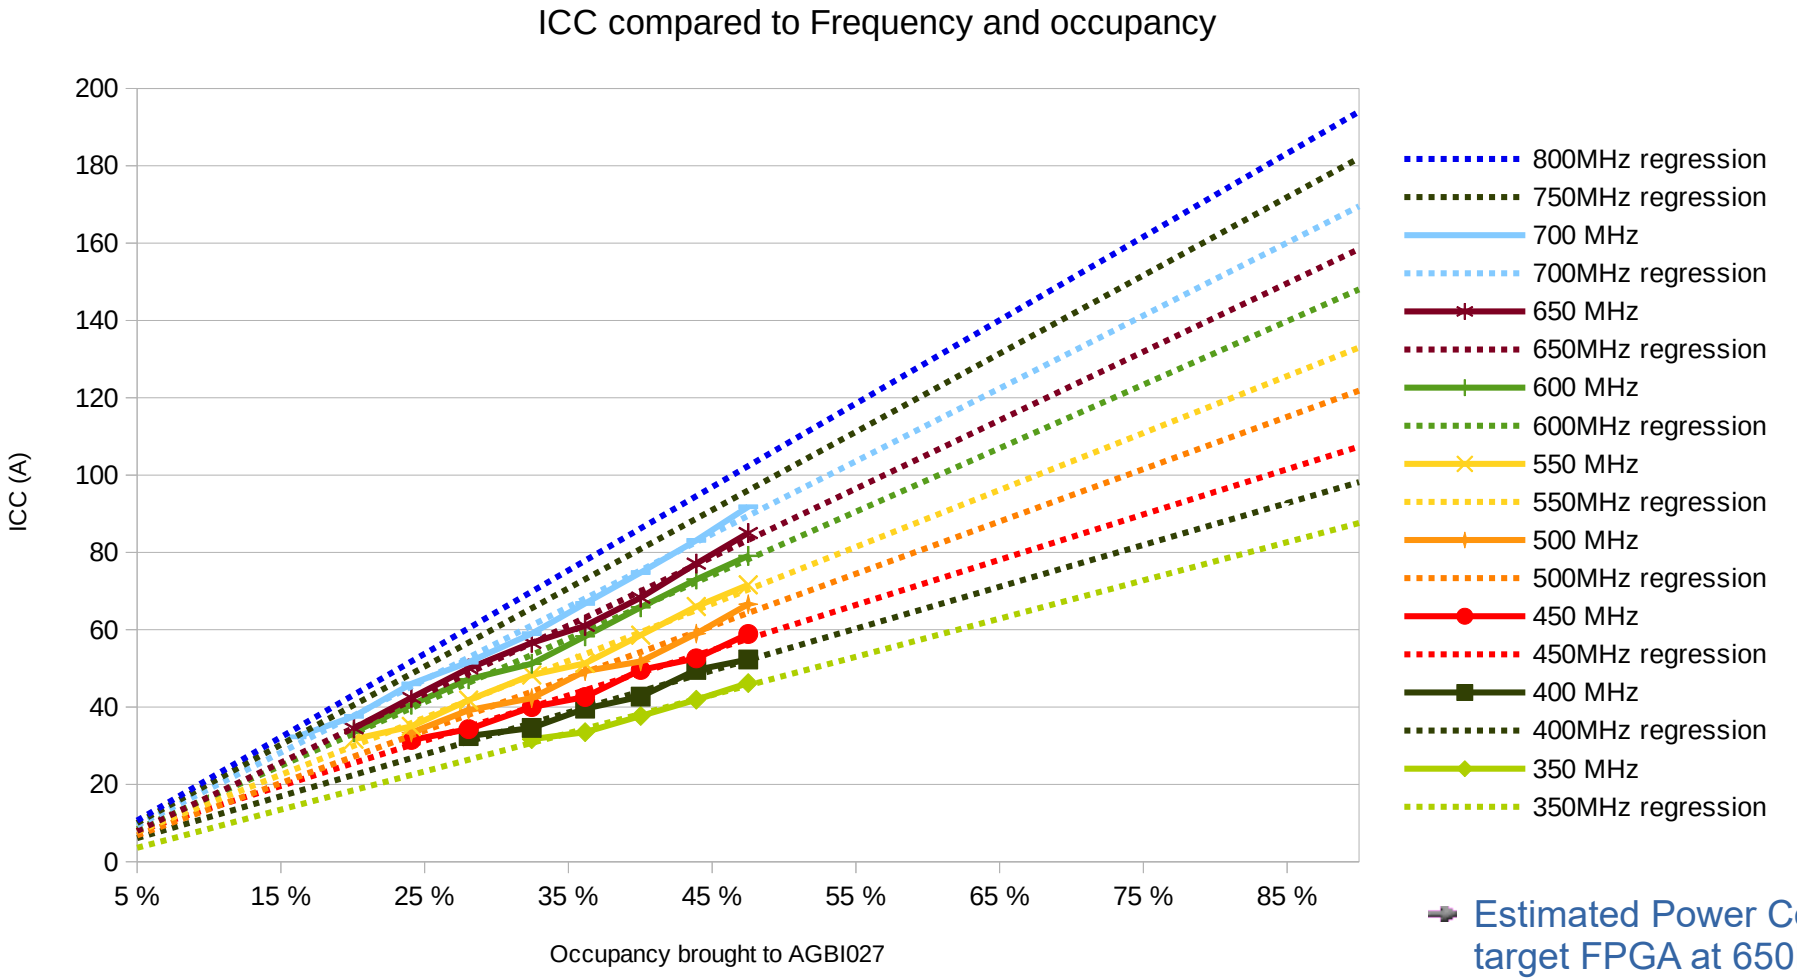

- Estimated Power Consumption for target FPGA at 650MHz is ~160A

Agilex core Power estimations

Simulation Method :

- Final firmware not available: use of random Pattern generators to simulate power consumption
- Compilation on Intel Quartus pro using Power Analyzer Tool
- Use of post fit simulation to simulate toggle rate data in QuestaSim
- Export of vcd file to Quartus Power Analyzer
- Estimation of power based on real toggle rate and temperature

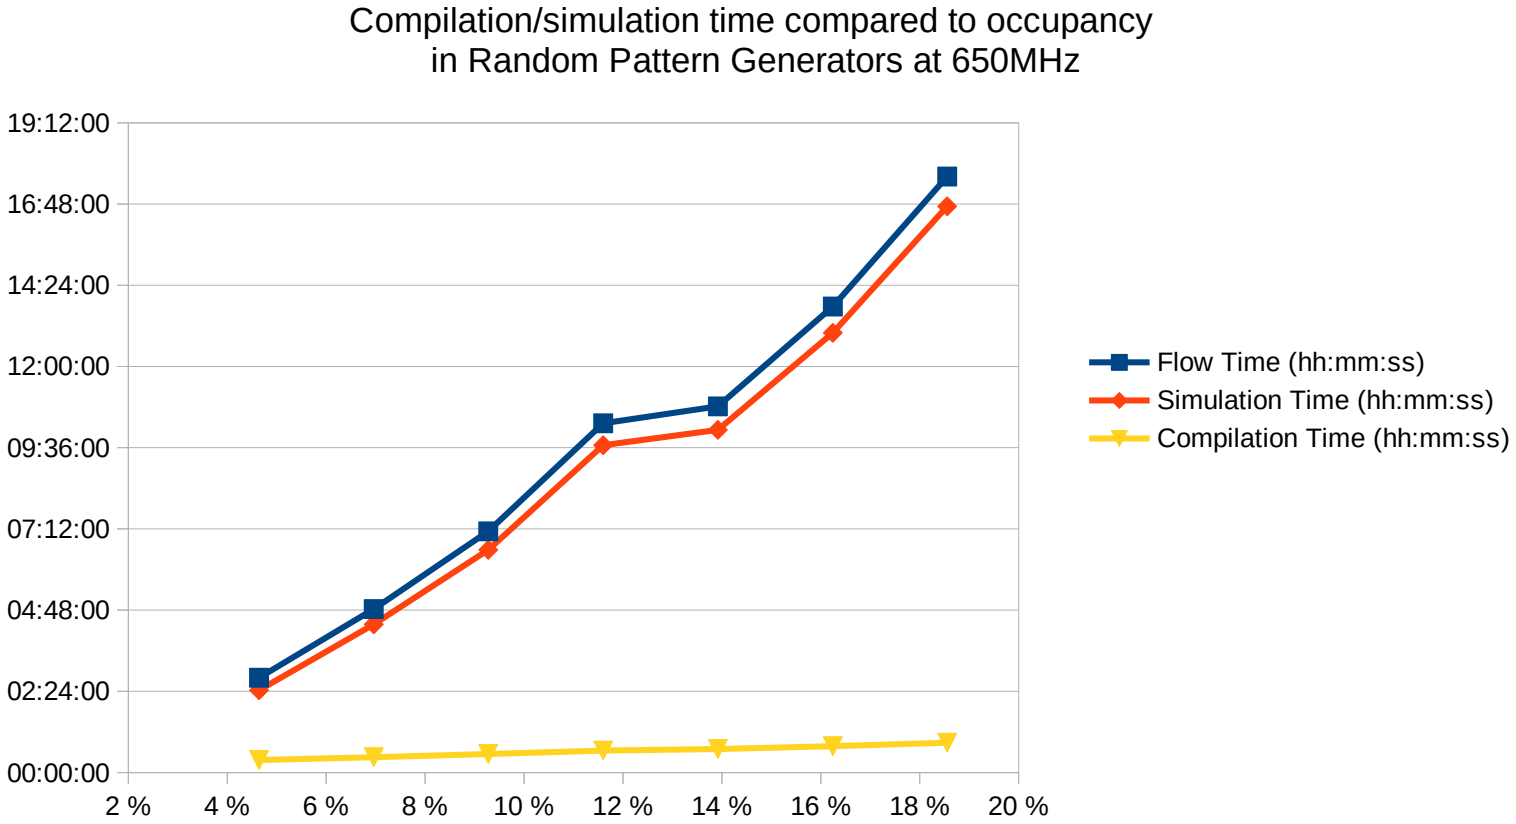

Agilex core Power estimations

Simulation Method :

- Incremental scripts, Up to 16 % occupancy at 3 frequencies
- Temperature measured on devkit is an additionnal input to improve power analysis

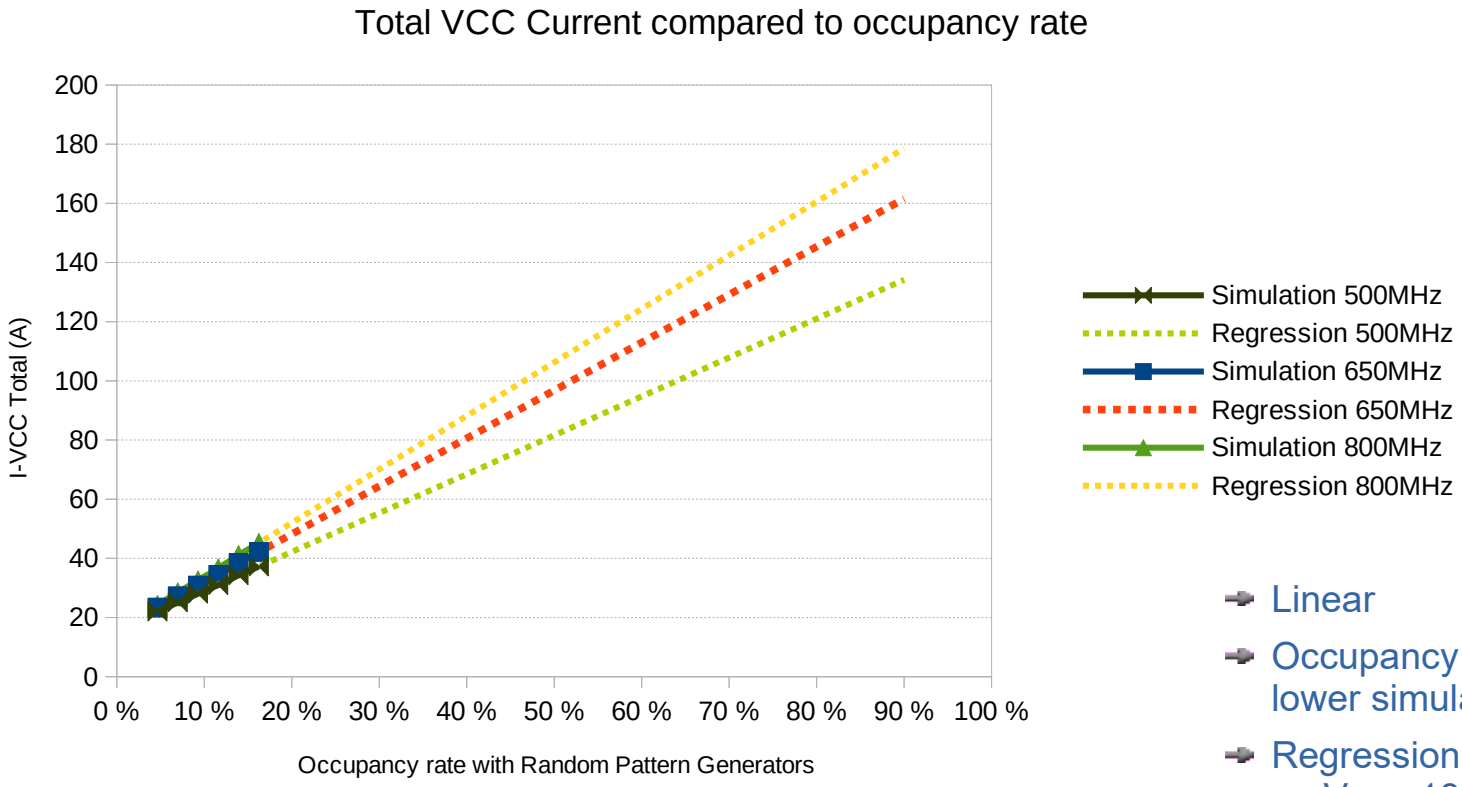

- Linear
- Occupancy can be reduced to get lower simulation time
- Regression : for 90 %, Max Current on Vcc ~161A at 650MHz

Agilex core Power estimations

Simulation vs Measurement extrapolations

- Simulation are not accurate if FPGA junction tempereture isn't an input
- As Power increases temperature increases

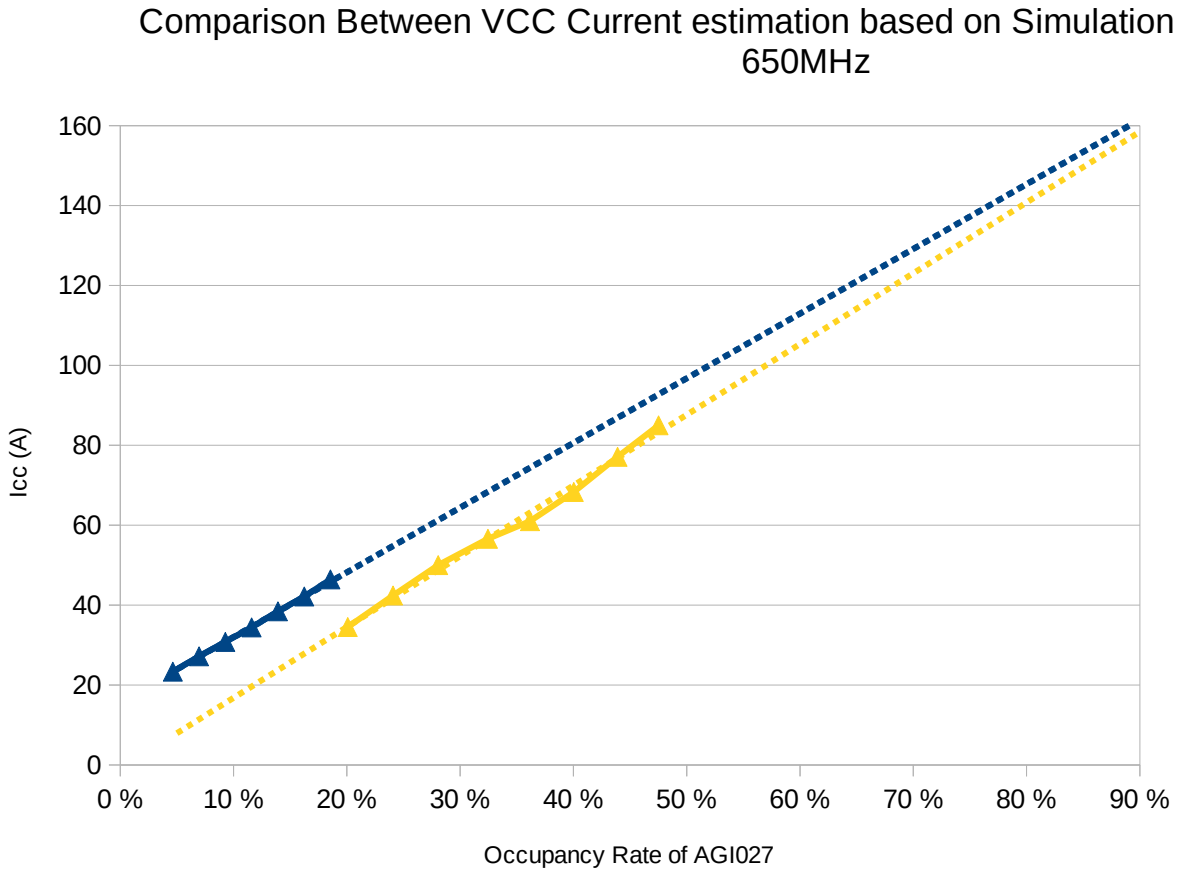

- Both methods confirm 160A max
- Simulations are pessimistic
- The difference is under 10 % for highest occupancy

Agilex Power Tree

Agilex FPGA requires :

- 11 Regulators / 24 rails
- Power On Sequencer (3 groups)
- Low ripple voltage / High accuracy (0,5%)

High Power regulators characteristics :

- Digital Power Management (PMBus, Voltage/ Current/Temperature Telemetry)
- Analog Devices (ex Linear) µmodules chosen
  - Low Surface
  - Easy to simulate & implement
  - Expensive but easy to find

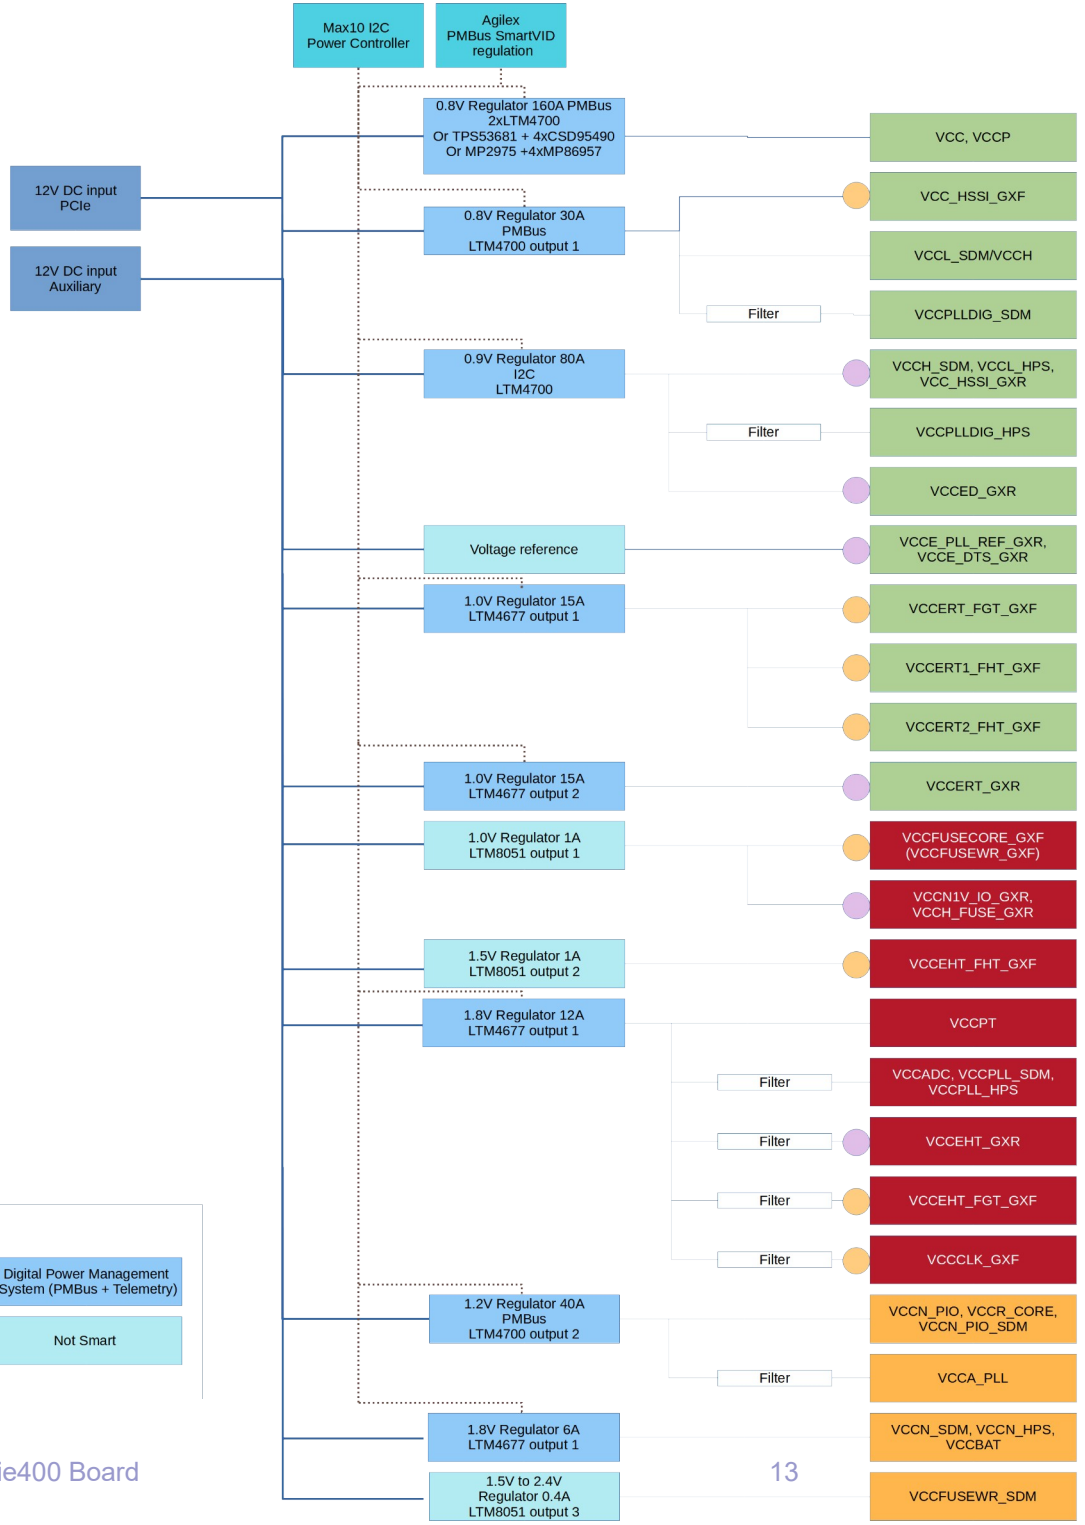

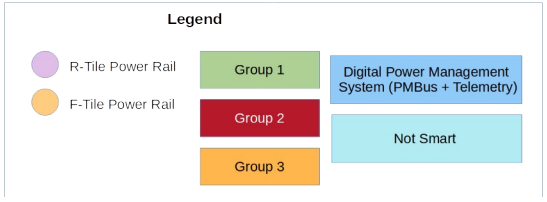

Conclusion

A lot of uncertainty for component procurement

- Many critical components are already ordered
- Producing prototypes in 2023 will be challenging

Challenges to come

- High component density
- Power Integrity
- Thermal studies
- Signal integrity (56Gbits/s signals)
- Will become a multi laboratory project

Backup

Increasing the individual link speed

Above 25 Gbits/s FPGAs use PAM4 instead of NRZ

- Multi-level signaling

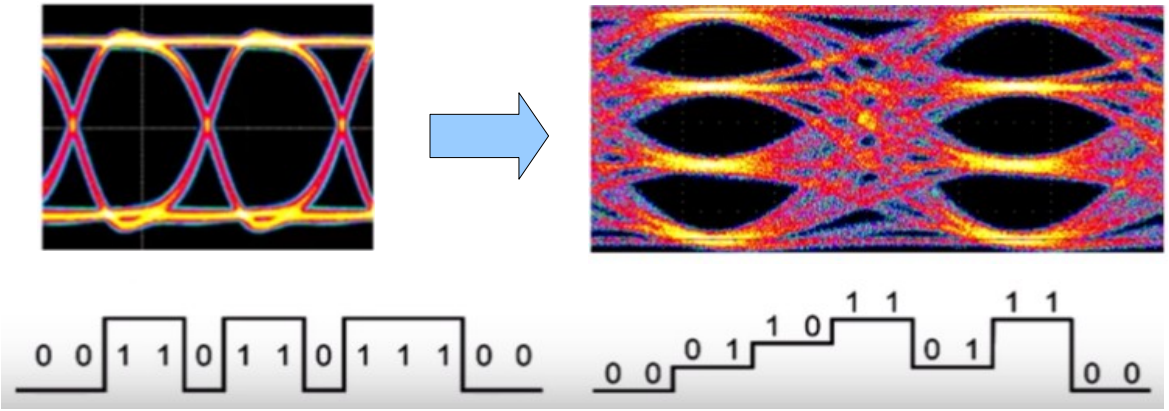

- More speed with constant bandwidth
- Data rate commonly found: 28 or 56 Gbits/s

Feasibility of serializer in HEP

- Research papers show that 40 Gbits is theorically possible with 65nM technologies
- Future developments could be in 28 nM

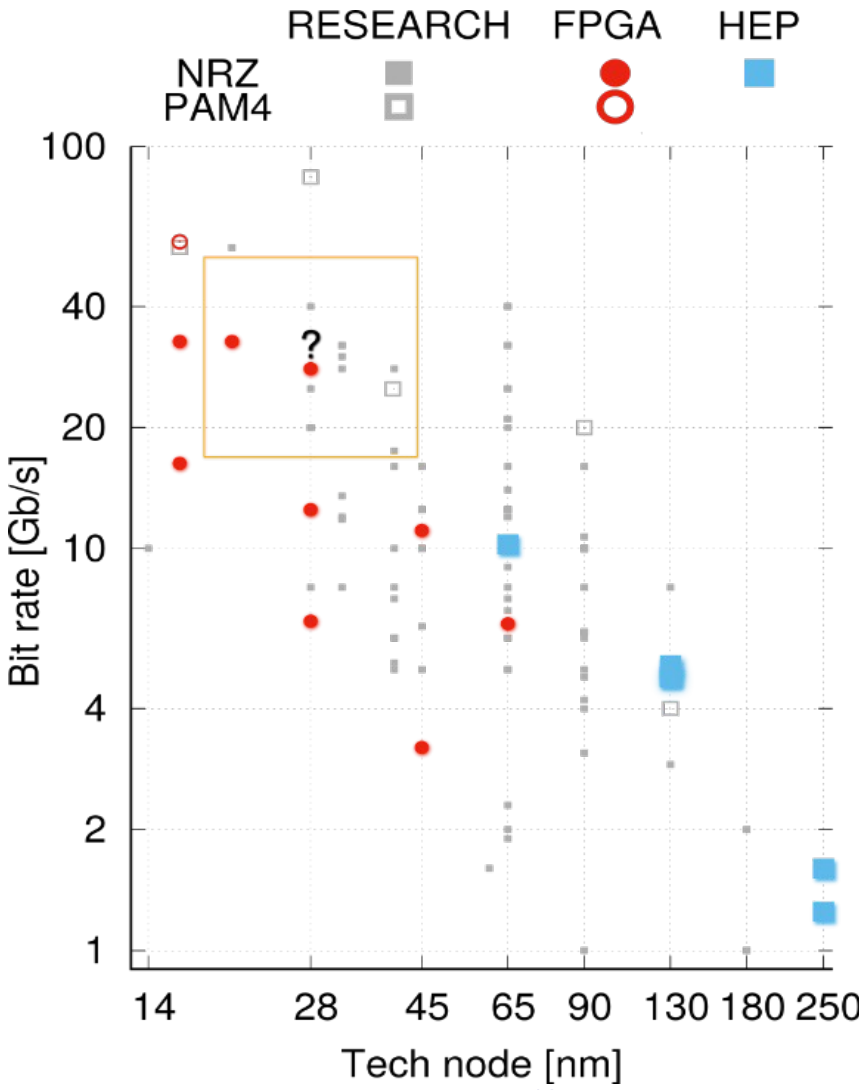

Evolution of high speed links at CERN (1/2)

R&D on experimental technologies

- WG6 : high speed links
  https://espace.cern.ch/ep-rdet-wg6-links/Shared%20Documents/WP6_Report_v1.03_19102018.pdf

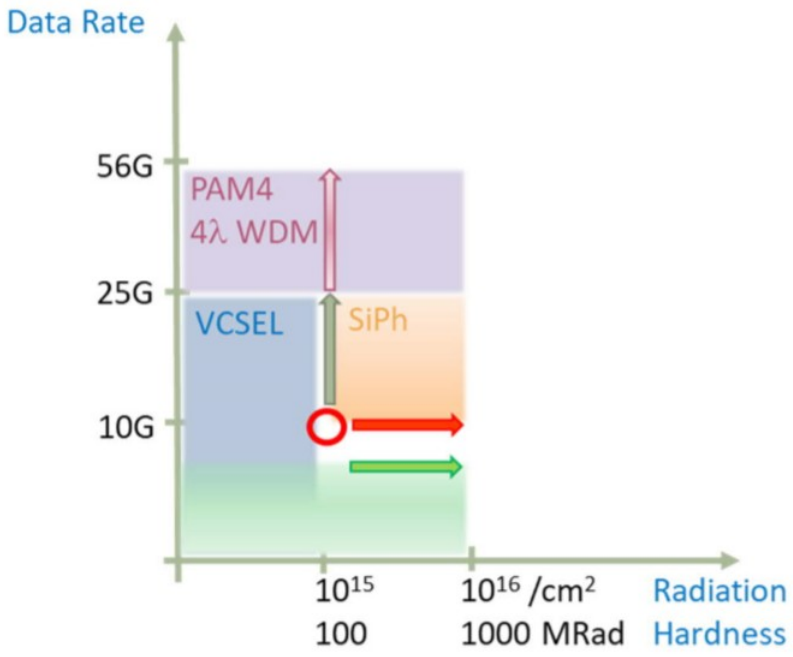

Evolution of high speed links at CERN (2/2)

Data aggregation

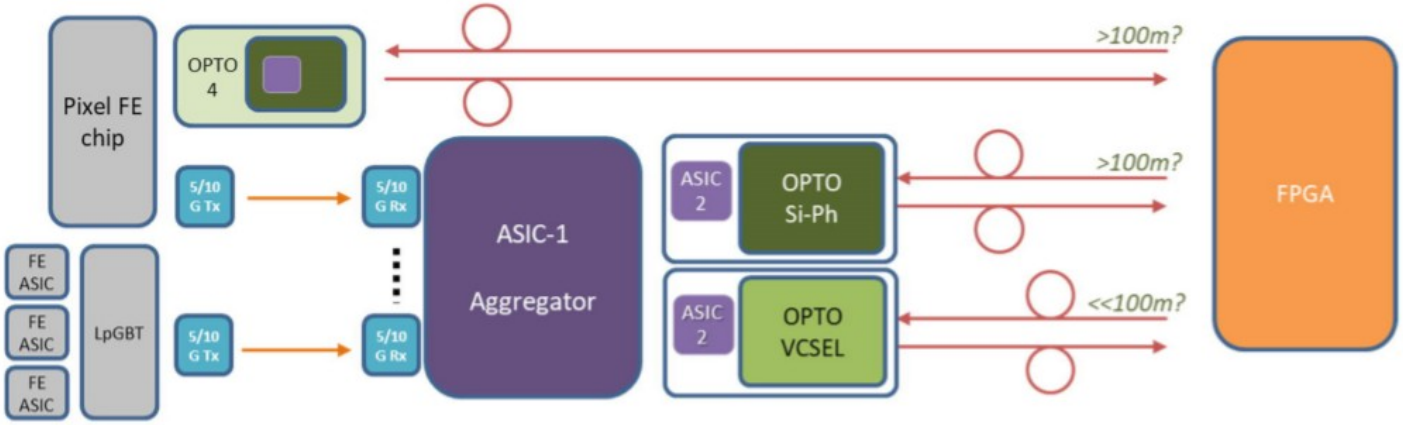

Activity	Task	Description
ASICs	ASIC-1	Very high data rate aggregator/transmitter
	ASIC-2	Optoelectronics drivers
	ASIC-3	Low-mass electrical cable transmission (active cable)
FPGA	FPGA-1	FPGA-based system testing and emulation
OPTO	OPTO-1 & 2	Silicon Photonics System & Chip Design
		Silicon Photonics Radiation Hardness
	OPTO-3	Next-generation VCSEL-based optical link
	OPTO-4	Silicon Photonics packaging