100 Gbps data acquisition system prototyping

Progress and Challenges

Alberto Dassatti

Jan 23, 2024

Who we are

Hes·so

HEIG VD

HEIG VD | REDS Institute Reconfigurable and Embedded Digital Systems

What we do

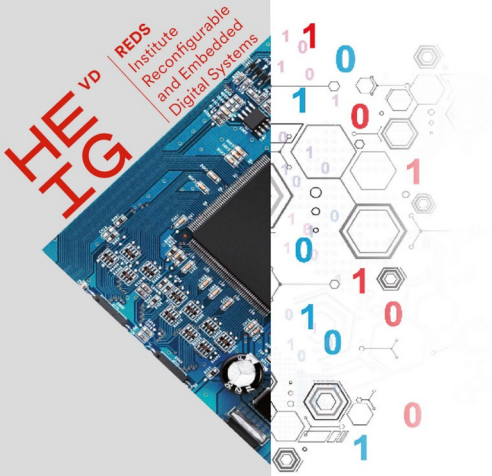
HEIG VD
REDS
Institute
Reconfigurable
and Embedded
Digital Systems

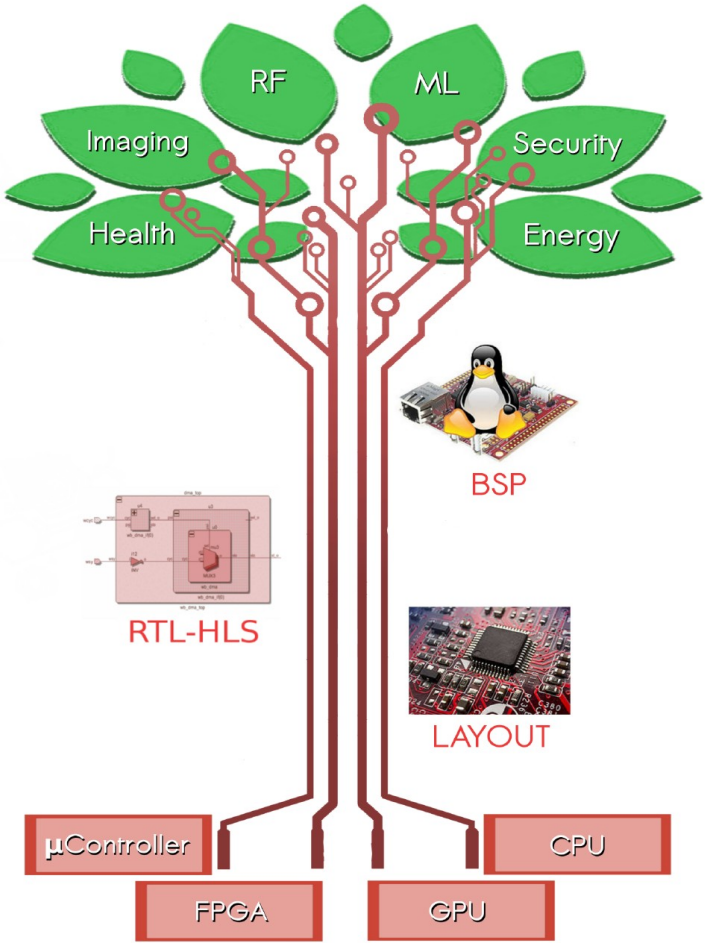
RF
ML
Imaging
Security
Health
Energy
BSP
RTL-HLS
LAYOUT
µController
CPU
FPGA
GPU

System View

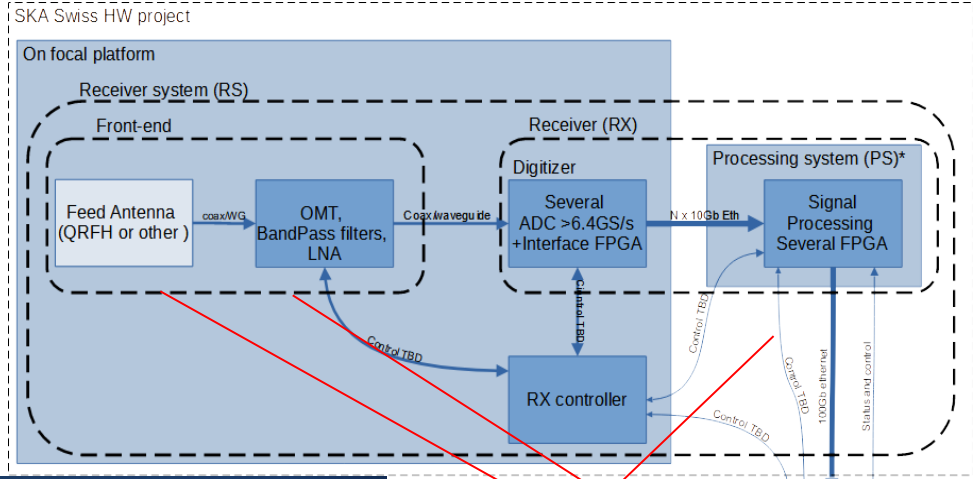

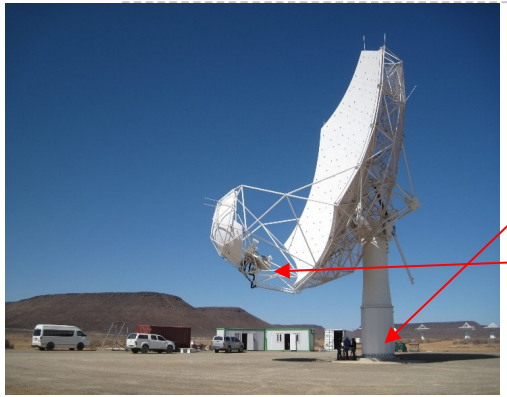

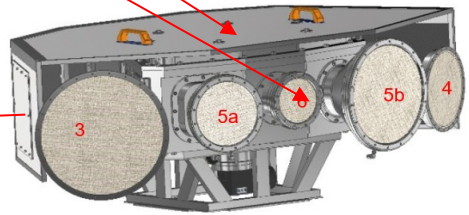

Detail of indexer

System View

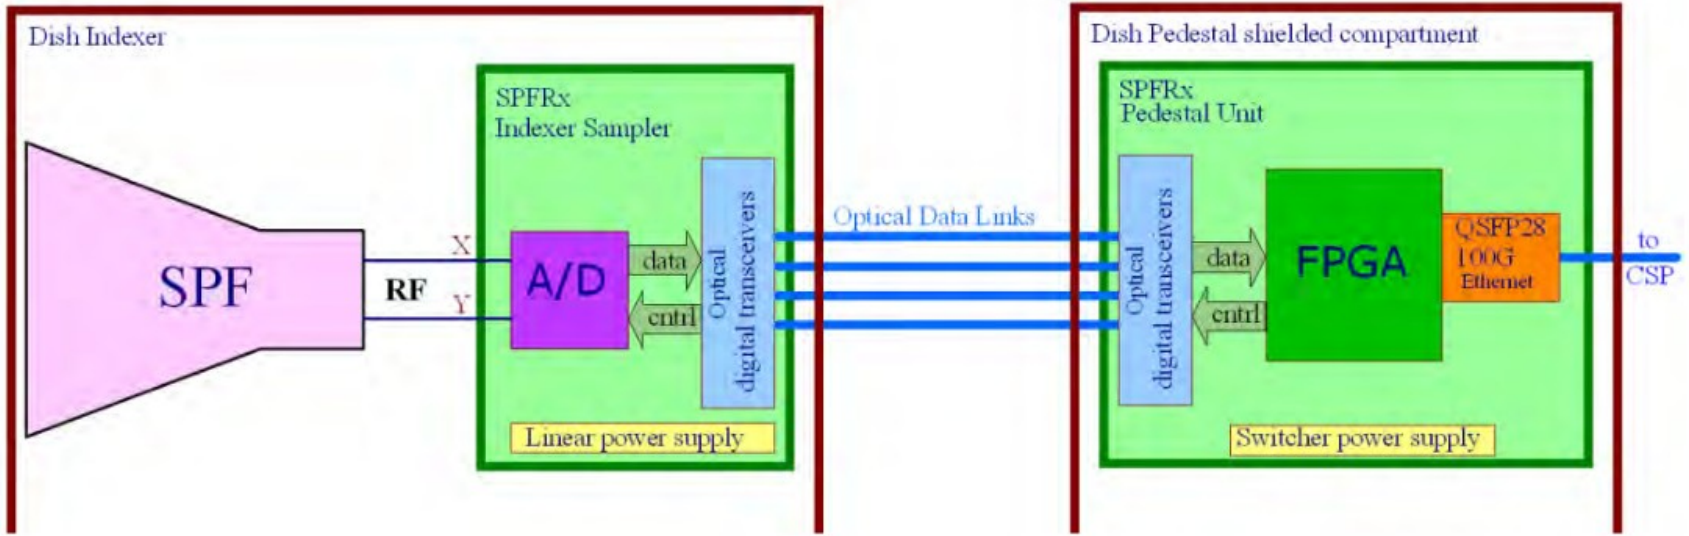

Figure 4 - SPFRx ODL architectural scheme

RXPU

Image from SKA-TEL-DSH-0000021-06_SPFRxDesignReportSigned.pdf

Talon Dx Board

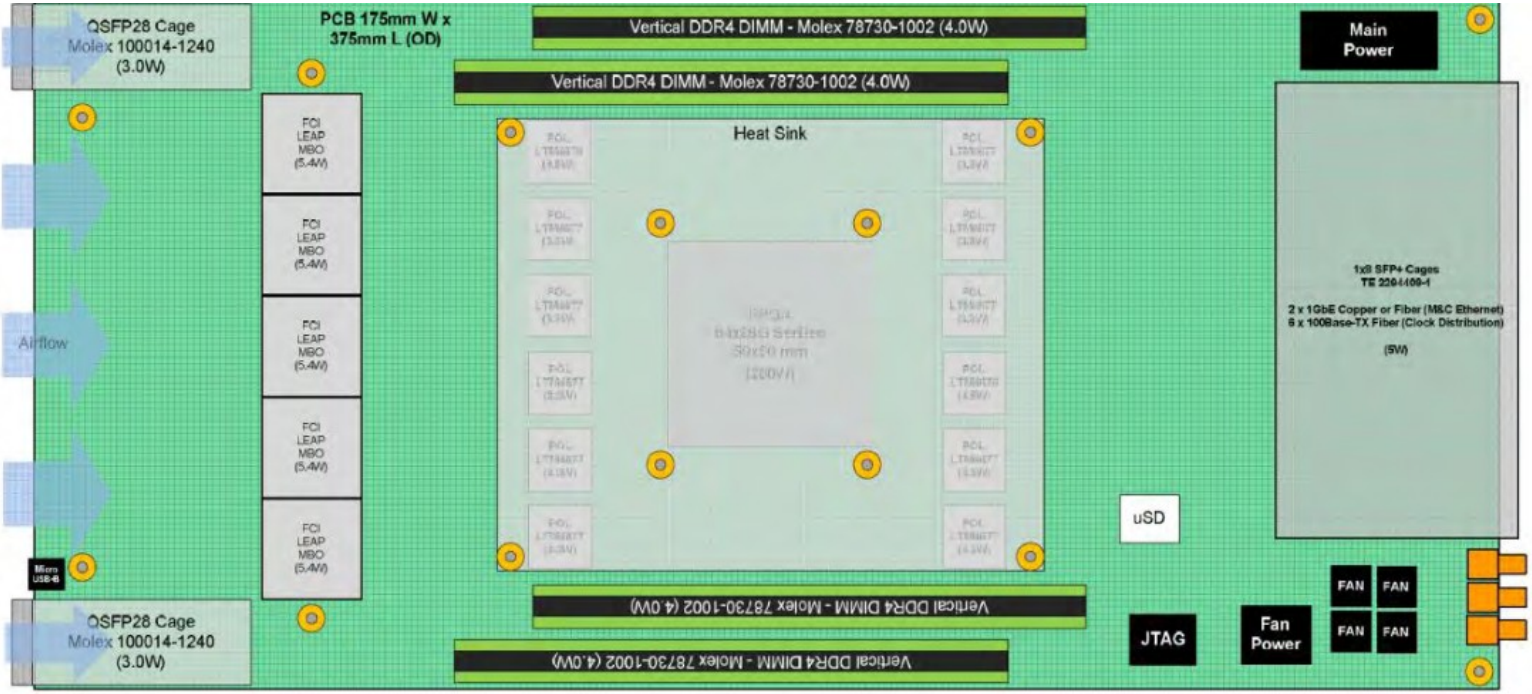

Figure 88 - TALON-DX Board by NRC –components placement on top surface

Image from SKA-TEL-DSH-0000021-06_SPFRxDesignReportSigned.pdf

EBB Specs and Interfaces

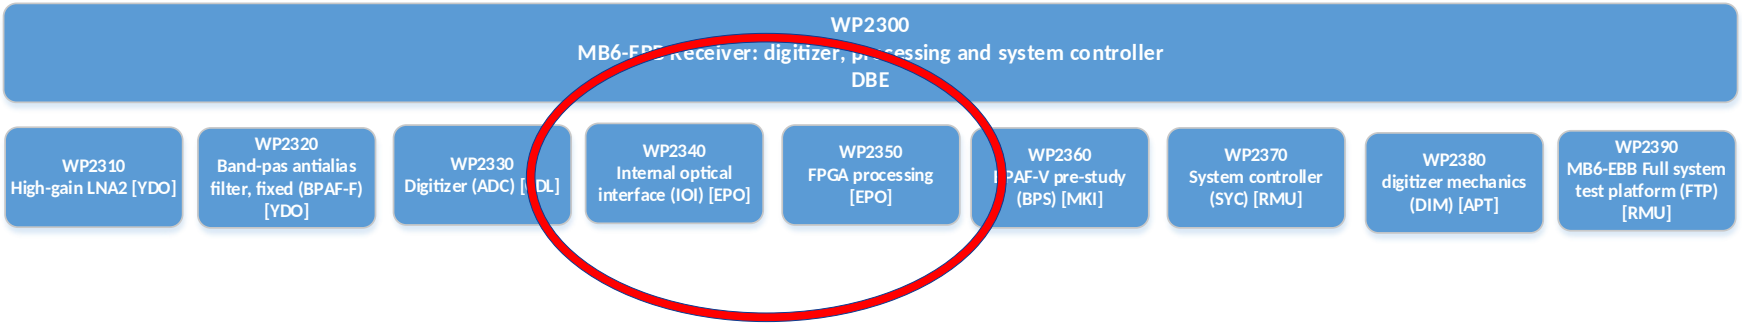

From ADCs to 100Gb/s Ethernet:

- 2 ADCs – 12.5 Gs/s – Data format: 10 b/s - FMC connector
- Optical fiber isolation between ADCs and FPGA processing
- Talon compatible (future integration?) as much as possible

Prototyping Set-Up

2 Custom made Boards to interface the ADC board with the FPGA board via optical fibers.

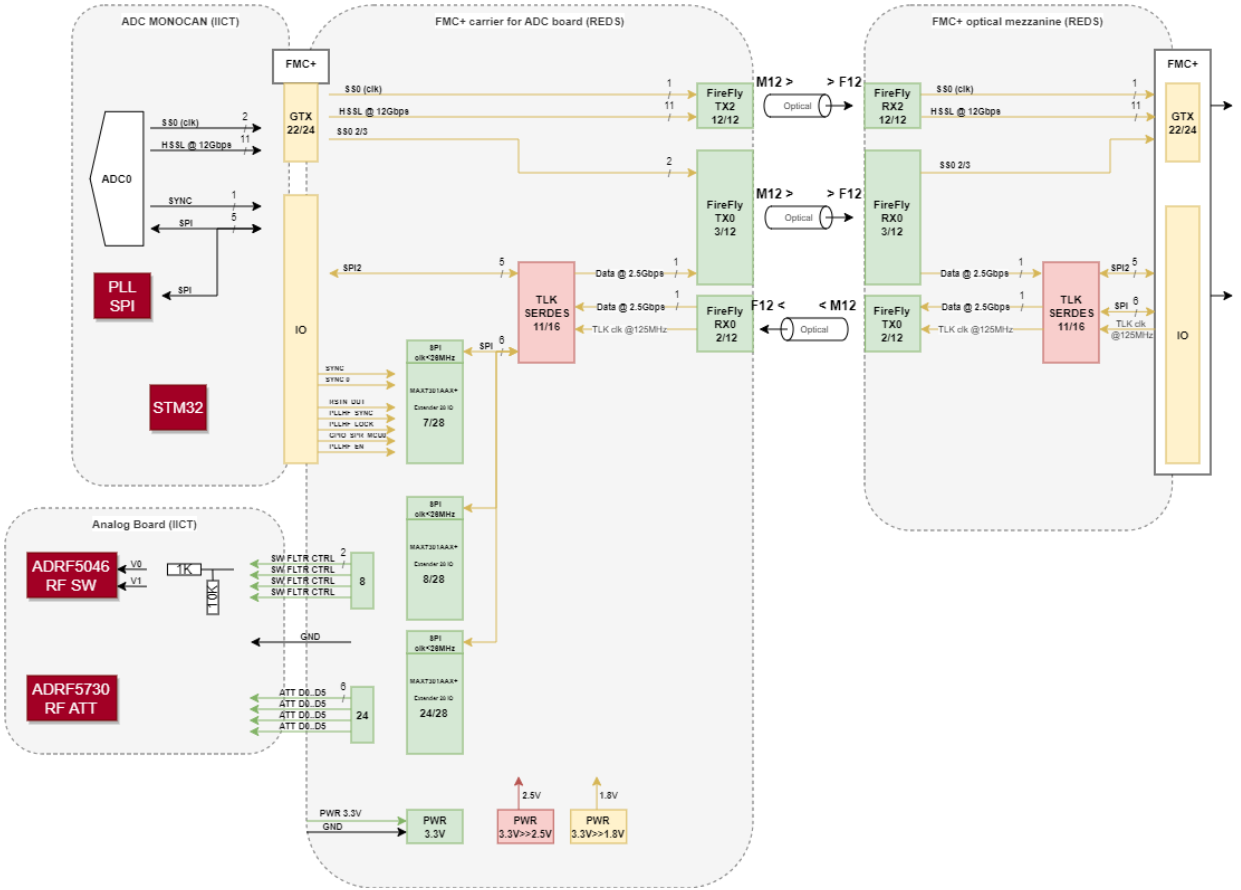

Prototyping Set-Up

SKACH

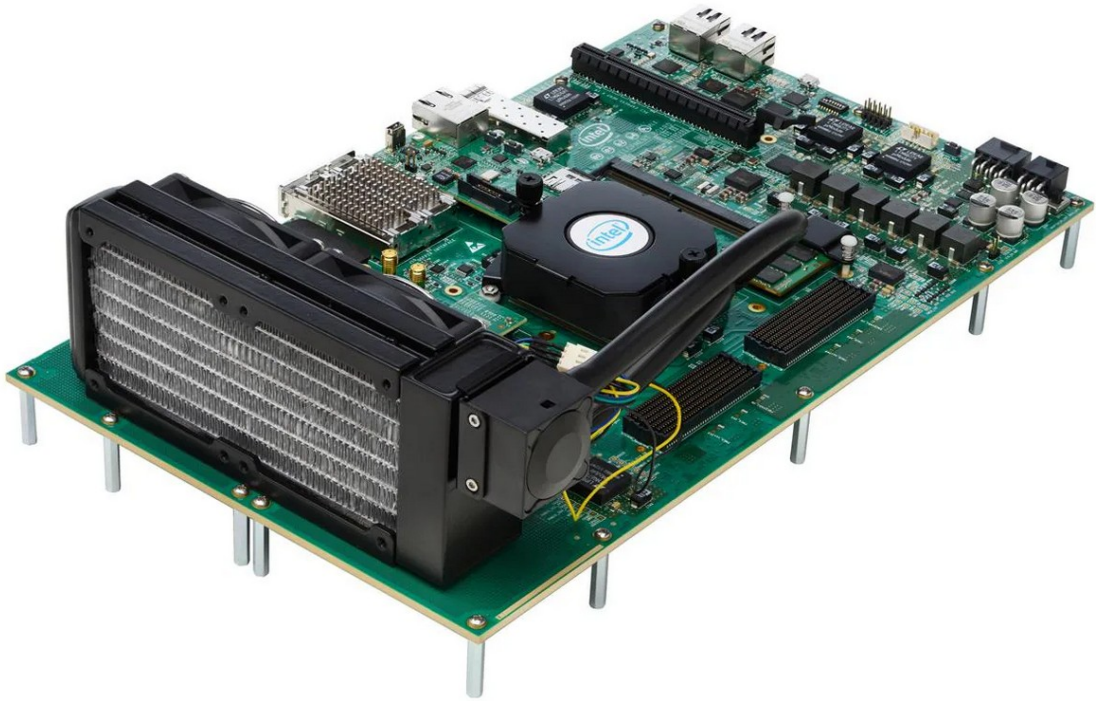

Prototyping Set-Up

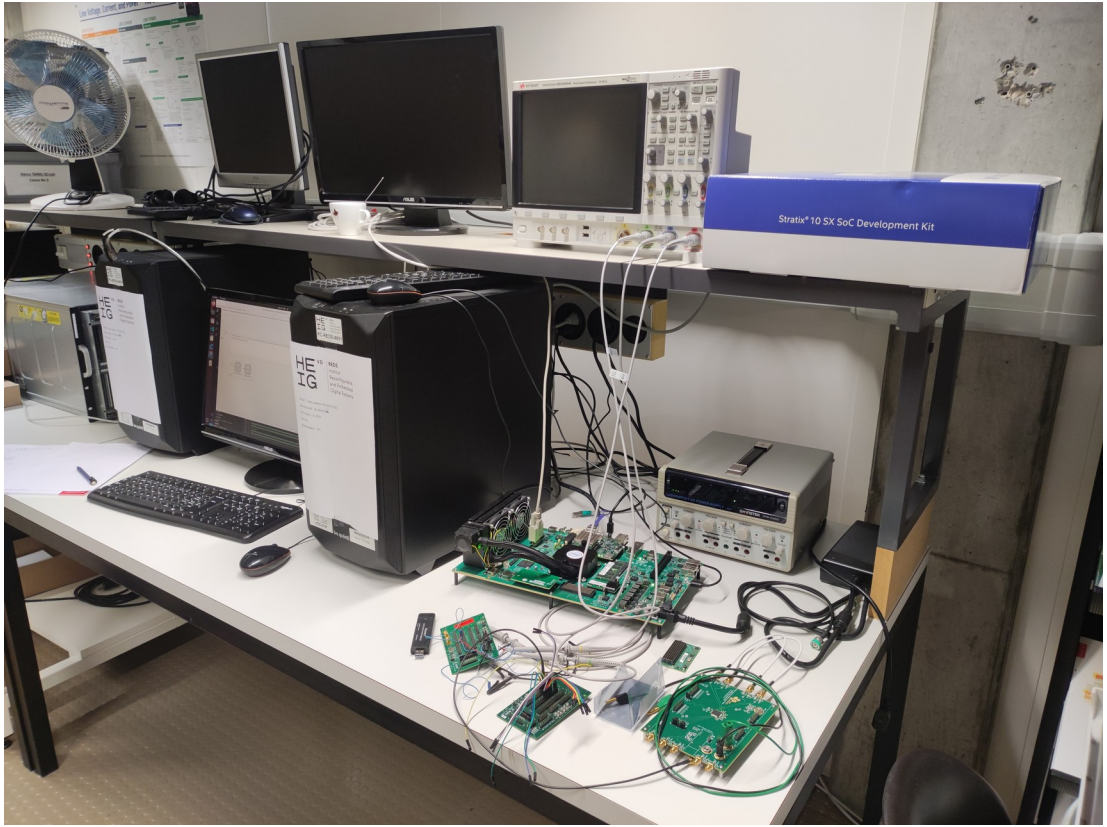

Hardware test design

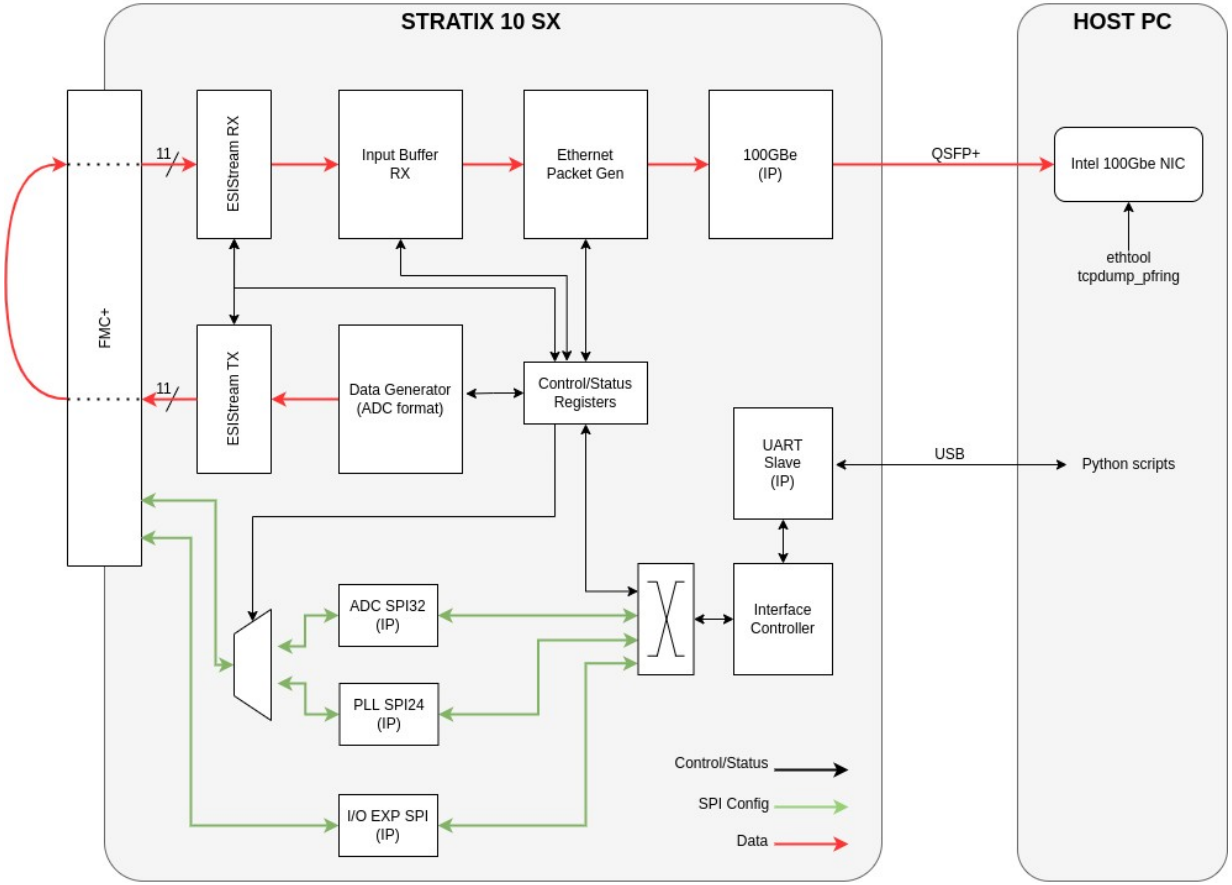

Connecting a 100Gb/s cable is…

Ethernet	10M	100M (Fast)	1G (Gigabit)	10G	100G
Time per bit	100 ns	10 ns	1 ns	0.1 ns	0.01 ns
Time for a MTU size frame 1500 bytes	1500 us	150 us	15 us	1.5 us	150 ns
Time for a 64 byte packet	64 us	6.4 us	640 ns	64 ns	6.4 ns
Packets per second	~10 K	~100 K	~1 M	~10 M	~100 M
Packets per 10 us		2 (small)	20 (small)	6 (MTU)	60 (MTU)

Christopher Lameter, Berlin 2016

DDR latency
~100 ns

Syscall
~10 us

…not enough

Linux software stack

SKACH

Worst possible situation, single very fast transmission:
- No Flow steering (Intel at least)
- No NIC multi-queue support
- No per-Core processing
- No BPF help
- NO XDP help

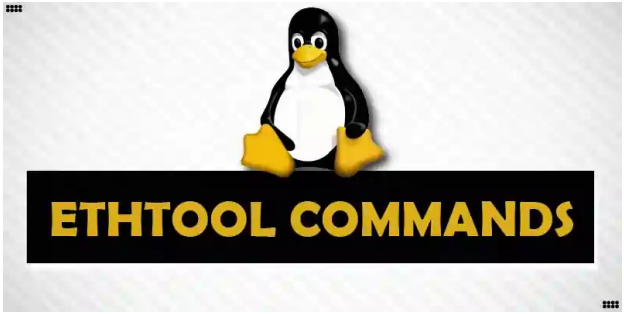

DPDK
DATA PLANE DEVELOPMENT KIT

```
# /usr/local/sbin/tcpdump_pfring --version
tcpdump_pfring version 4.9.3
libpcap version 1.9.1 (with TPACKET_V3)
```

PF_RING™

High-speed packet capture, filtering and analysis.

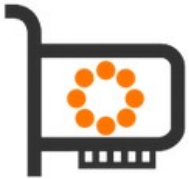

Measurements

Max data rate achieved ~97 Gbps (with syntetic data generated by the FPGA)

ADC interface:

2 ADCs X 12.5 Gbps X 4 bits/sample = 100 Gbps

->

All the stream can not be transmitted on 100GbE.

Proposed solution is to reduce the lane rate to 11.5 Gbps.

Data rate achieved:
Stream of ADC A or B: 46.3 Gbps
Stream of ADC A and B : 92.7 Gbps

What's next

Optical interface Boards:

- Routing to be completed
- Design Review
- Production:
  - Optical components lead time 7ws
  - Board production and assembly 5ws
- Test

PC Software:

- Acquisition Control
- Full system test

FPGA Design:

- Double ADC test on-going
- ADC Board integration test
- Optical boards test
- Full system test


HEIG VD
REDS
Institute
Reconfigurable
and Embedded
Digital Systems

SKACH