High throughput data acquisition with InfiniBand on x86 low-power architectures for the LHCb upgrade

Matteo Manzali
INFN - Università degli Studi di Ferrara

The LHCb experiment

- The LHCb experiment is one of the four large experiments based at CERN

- A major upgrade is scheduled in the 2018-2020 period:

    - Upgrade of the detector

    - Upgrade of the Data Acquisition system (DAQ)

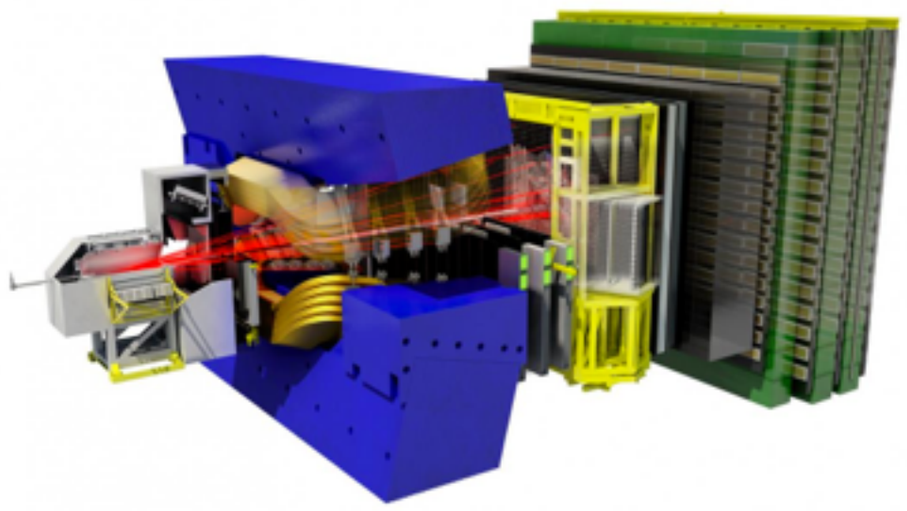

	2015		2018
Event size	65 KB	→	100 KB
Event rate	1 MHz	→	40 MHz
Aggregate bandwidth	520 Gb/s	→	32 Tb/s

Upgraded DAQ design

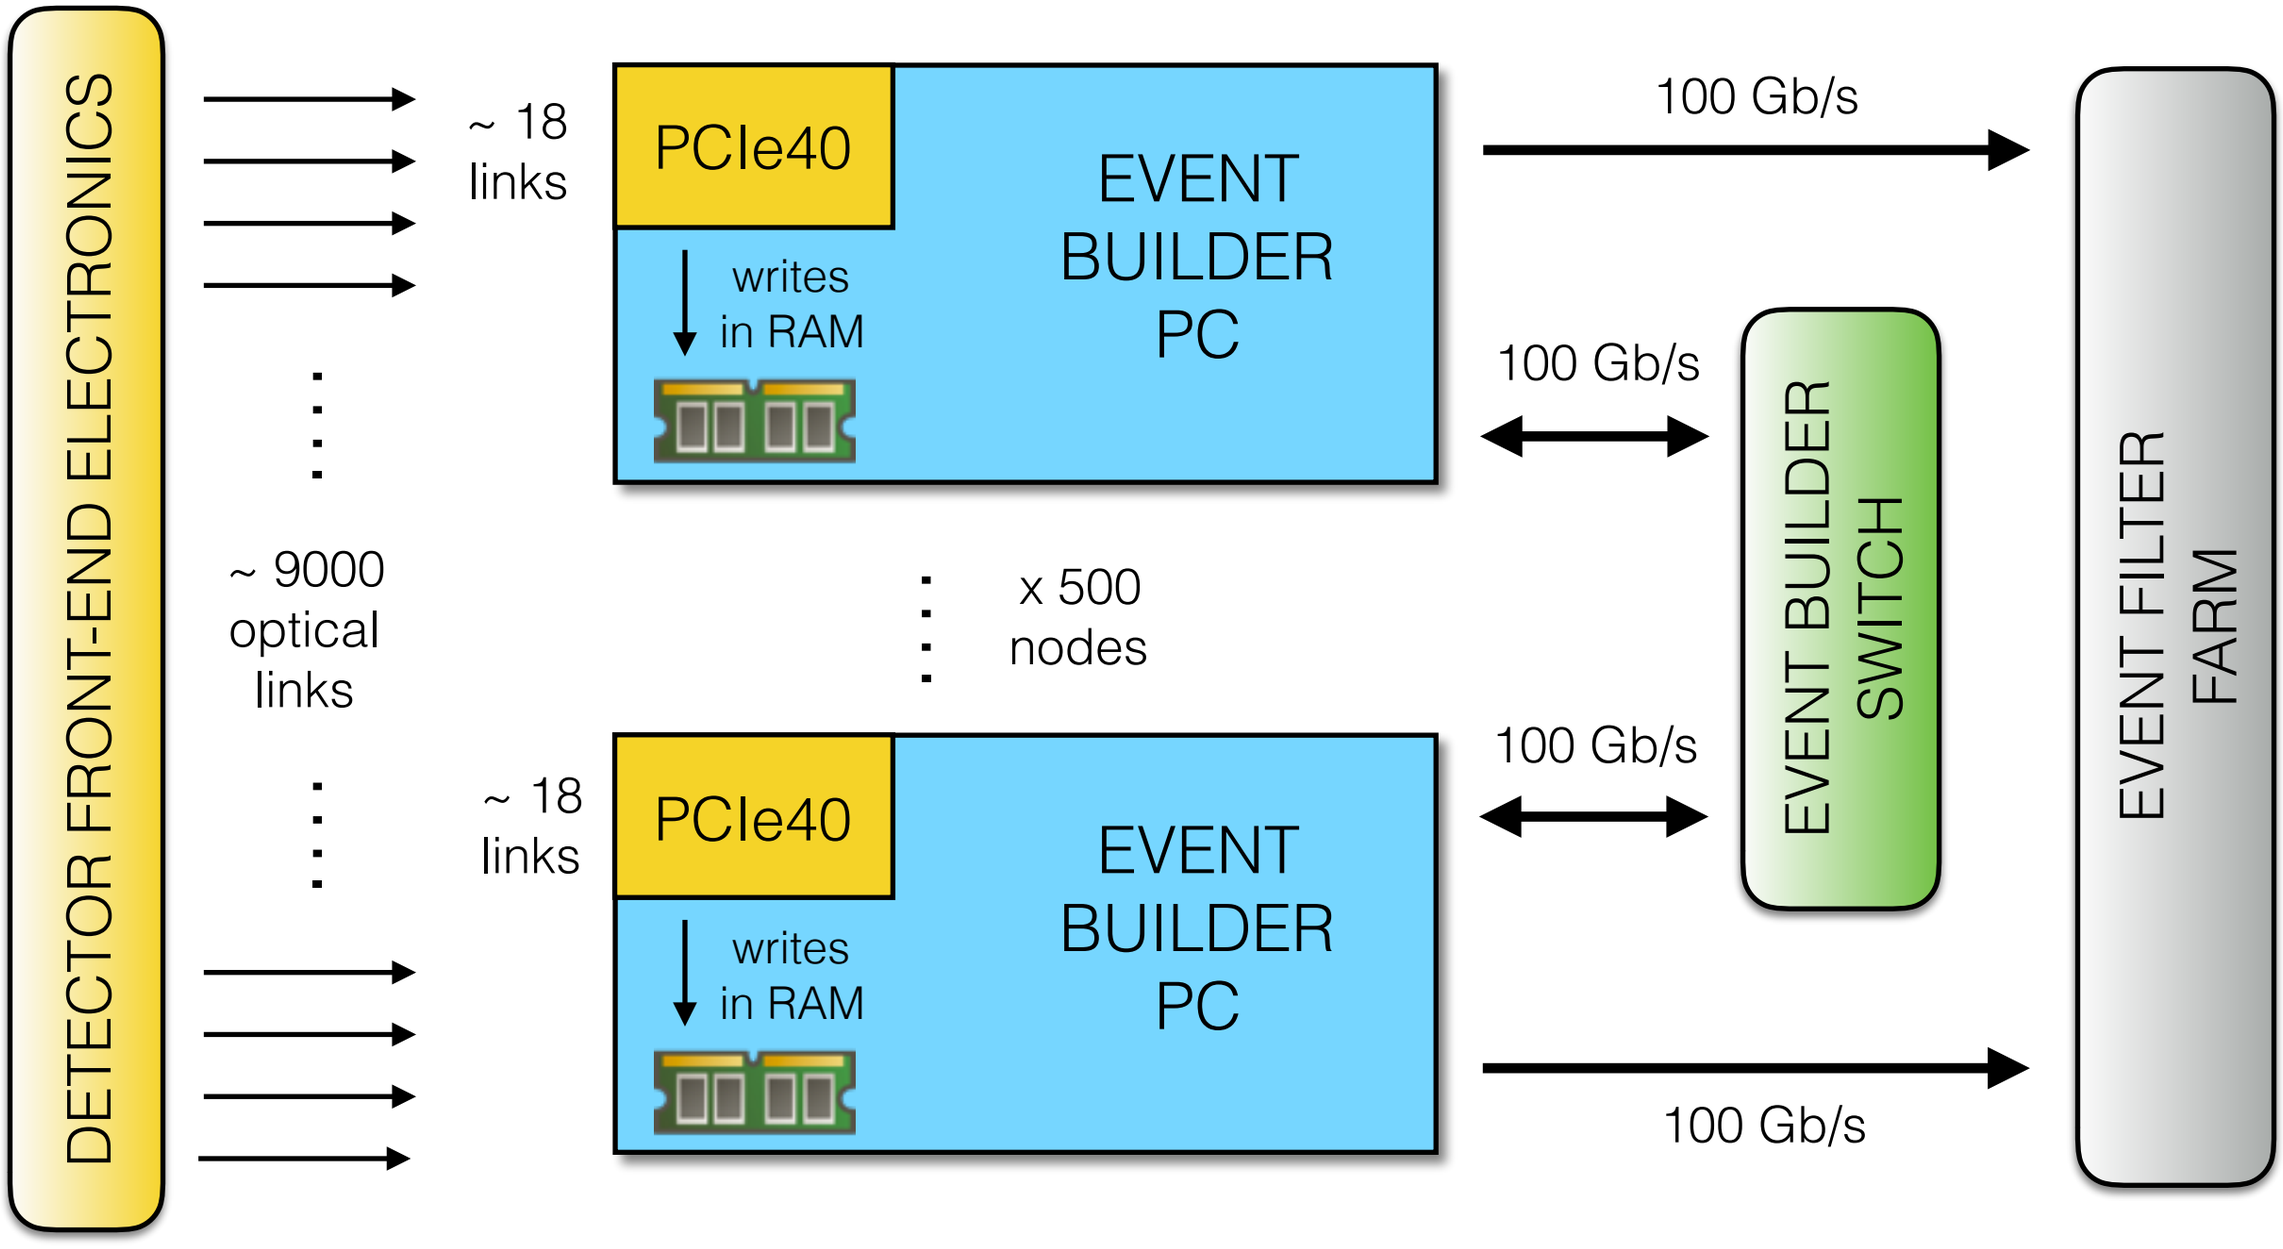

Network technologies

- Different network technologies under study by the LHCb online working group (Ethernet, InfiniBand, Intel OmniPath)
- InfiniBand standard will reach 200 Gb/s (HDR) at the end of 2017
- InfiniBand has Remote Direct Memory Access (RDMA) capabilities
  - It allows to access remote memory without involving CPU and OS

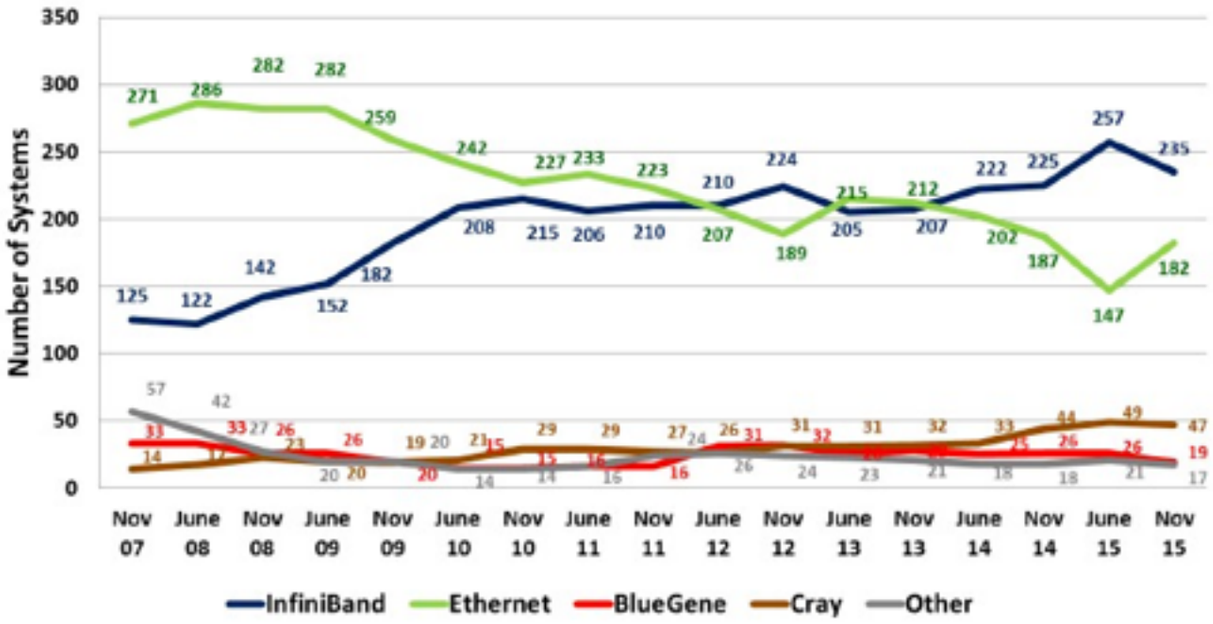

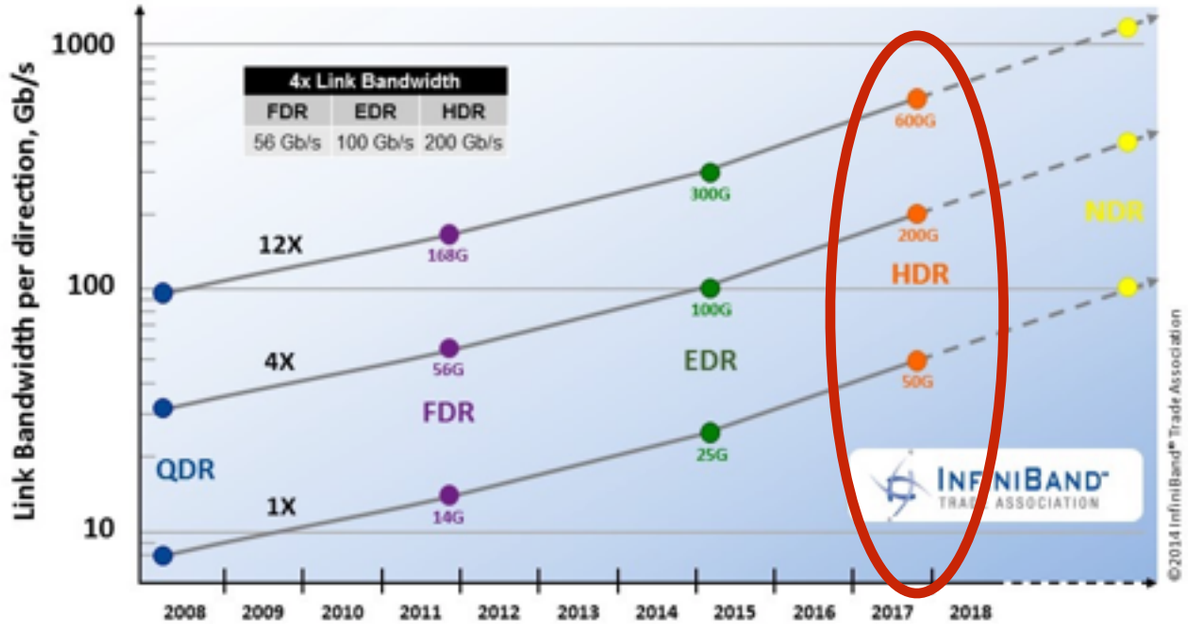

The Event Builder software

- The Large Scale Event Builder (LSEB) is an Event Builder software prototype based on the InfiniBand network technology

- It is composed of two distinct logical components, the Readout Unit (RU) and the Builder Unit (BU):

    - Each RU receives data from a generator, creates the event fragments and ship them to receiving BU in a many-to-one pattern

    - Each BU gathers event fragments together to generate full events

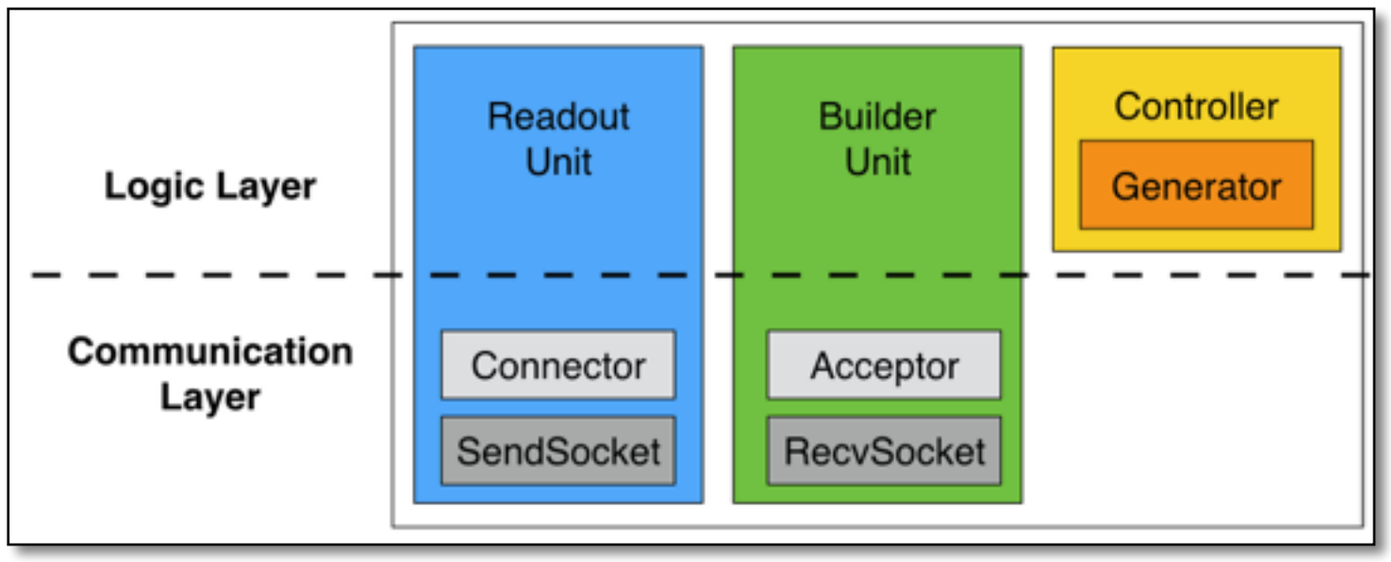

Low-power tests

- Tests performed in collaboration with the COSA project
- Each testbed is composed by two nodes connected back-to-back with InfiniBand interconnect
- Bandwidth, power consumption and temperature measurements while running the Event Builder software
- Results are compared to those obtained with a standard Intel Xeon processor (E5-2683v3)

Testbed setups

- Intel Xeon D-1540
- 8 core x86 SoC (2 threads per core)
- 1 slot PCIe 3.0 16x
- Mellanox FDR InfiniBand
- 54.3 Gb/s bandwidth

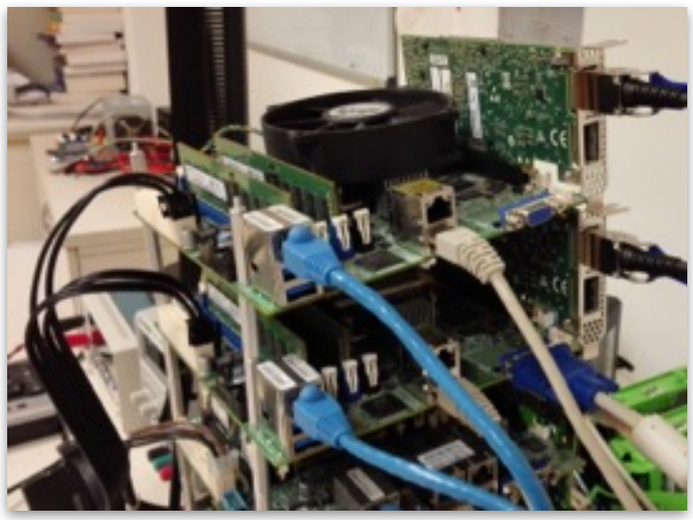

- Intel Atom C2750
- 8 core x86 SoC (no multithreading)
- 1 slot PCIe 2.0 8x
- QLogic QDR InfiniBand
- 27.2 Gb/s bandwidth

System settings

- The **CPUFreq** governor allows the clock speed of the processor to be adjusted on the fly:
    - Most commonly used are the ondemand and the performance
    - After comparison benchmarks we chose the performance

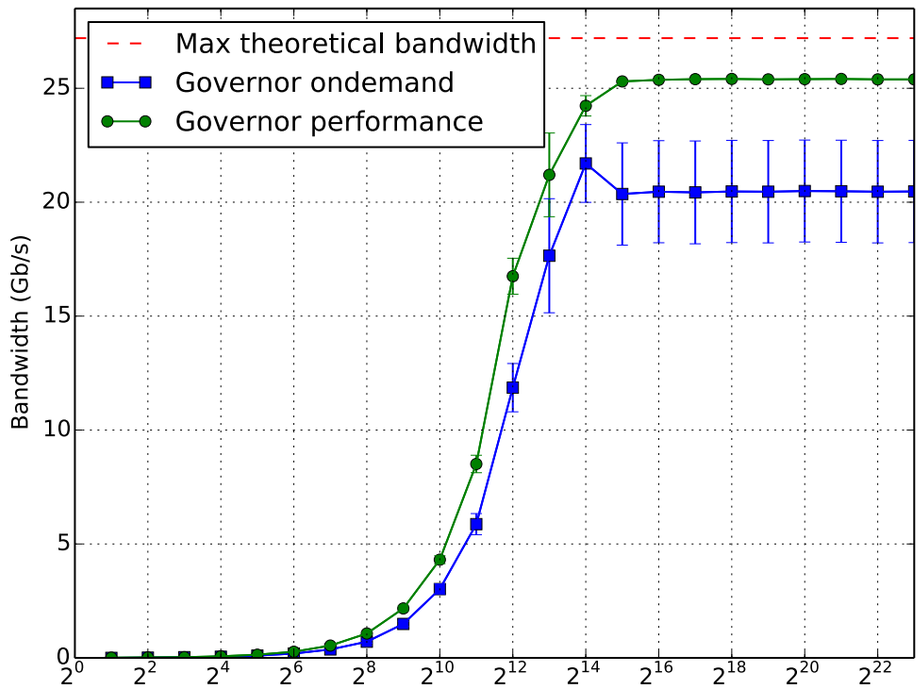

- **c-states** allow systems to save power by partially deactivating CPU parts that are not in use:
    - more CPU units are stopped and more energy is saved
    - but more time required for the CPU to be again 100% operational
    - we decided to keep the c-states enabled

Intel Xeon D-1540

- Pure computation process running on 4 cores in the second half of the test
- The performances are fully comparable
- The Xeon D consumes about 1/3 with respect to the standard Xeon

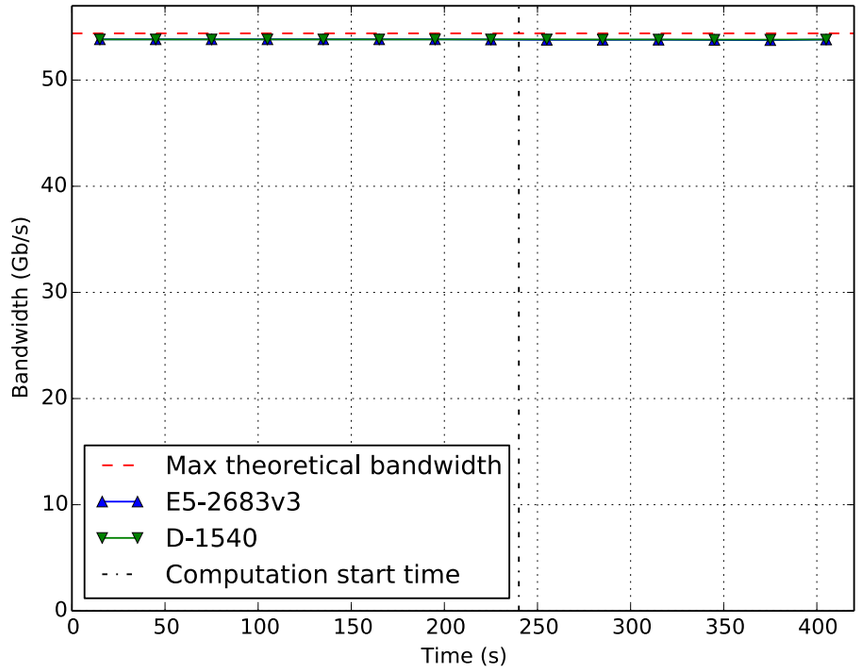

	E5-2683v3	D-1540
Idle power consumption	80.78 W	28.23 W
EB power consumption	161.00 W	49.02 W
EB power consumption with computation	176.54 W	79.12 W
Max temperature	56.0 C	59.0 C
Average bandwidth	53.82 Gb/s	53.82 Gb/s

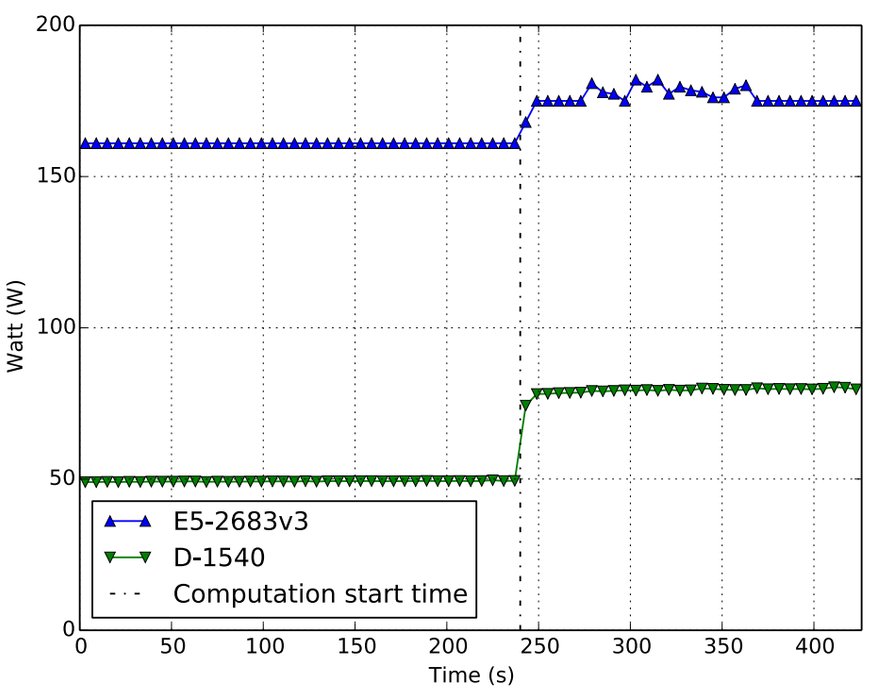

Intel Atom C2750

- The C2750 cannot reach the performance of the E5-2683 (~66% of the bandwidth)
- Power consumption of the C2750 is really low (~20% with respect to the E5-2683)
- Bandwidth becomes really unstable running a computation process on the C2750

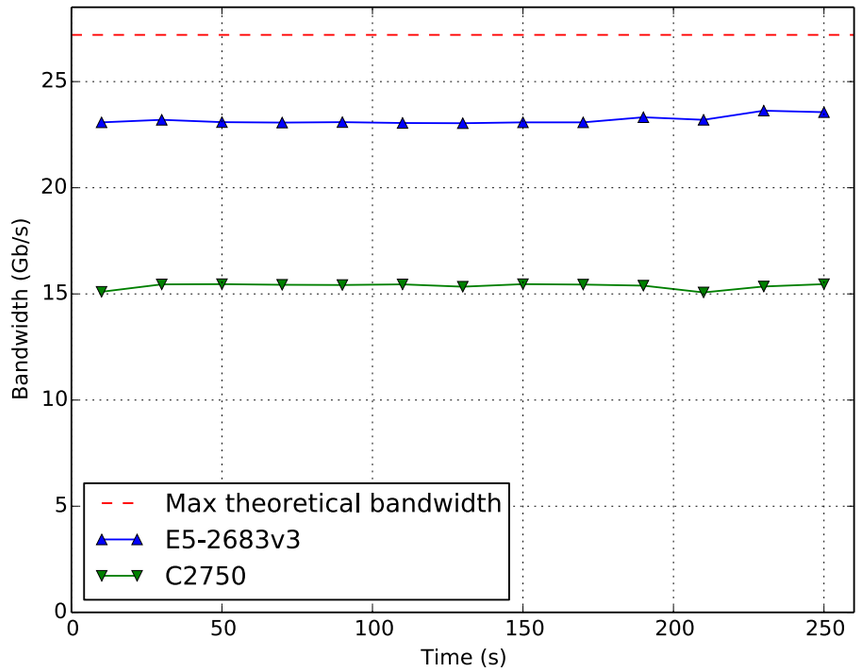

	E5-2683v3	C2750
Idle power consumption	77.46 W	18.20 W
EB power consumption	154.44 W	28.93 W
Max temperature	52.0 C	37.0 C
Average bandwidth	23.19 Gb/s	15.37 Gb/s

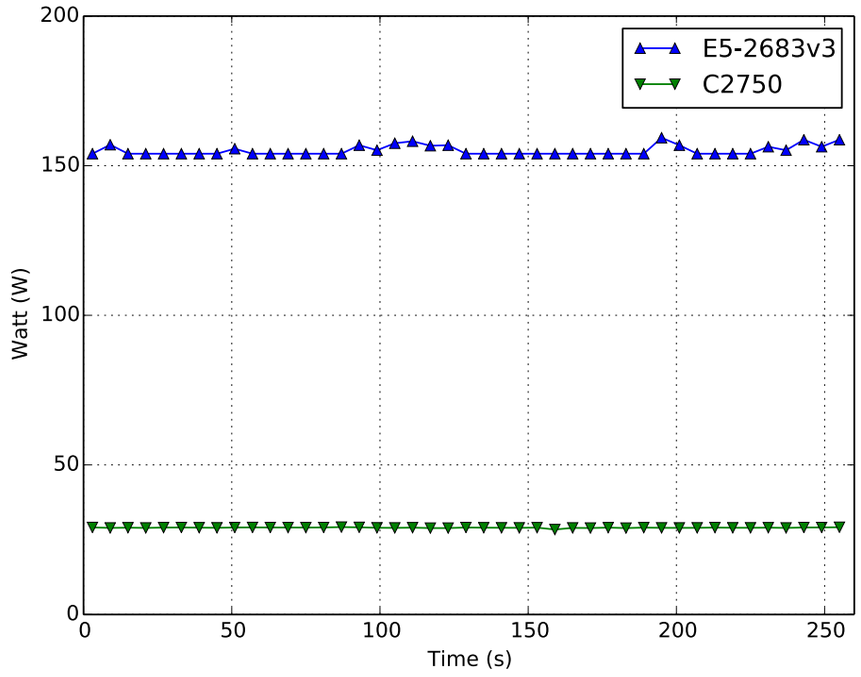

Conclusions

- The Intel Xeon D-1540 seems a really interesting processor for high-throughput data acquisition purposes.
    - It brings all the functionalities of the XEON family processors but reducing costs and power consumption.

- The Intel Atom C2750 can't compete with the high-performance XEON family processors.
    - It is still interesting for data acquisition purposes in case the requirements are not so strict.

- Further investigations will be made also with the Intel OmniPath technology.