Time structure analysis of the LHCb Online network

INTRODUCTION

The LHCb Online Network is a real time high performance network, in which 350 data sources send data over a Gigabit Ethernet LAN to more than 1500 receiving nodes. The aggregated throughput of the application, called Event Building, is more than 60 GB/s. The protocol employed by LHCb makes the sending nodes transmit simultaneously portions of events to one receiving node at a time,which is selected using a credit-token scheme. The resulting traffic is very bursty and sensitive to irregularities in the temporal distribution of packet-bursts to the same destination or region of the network.

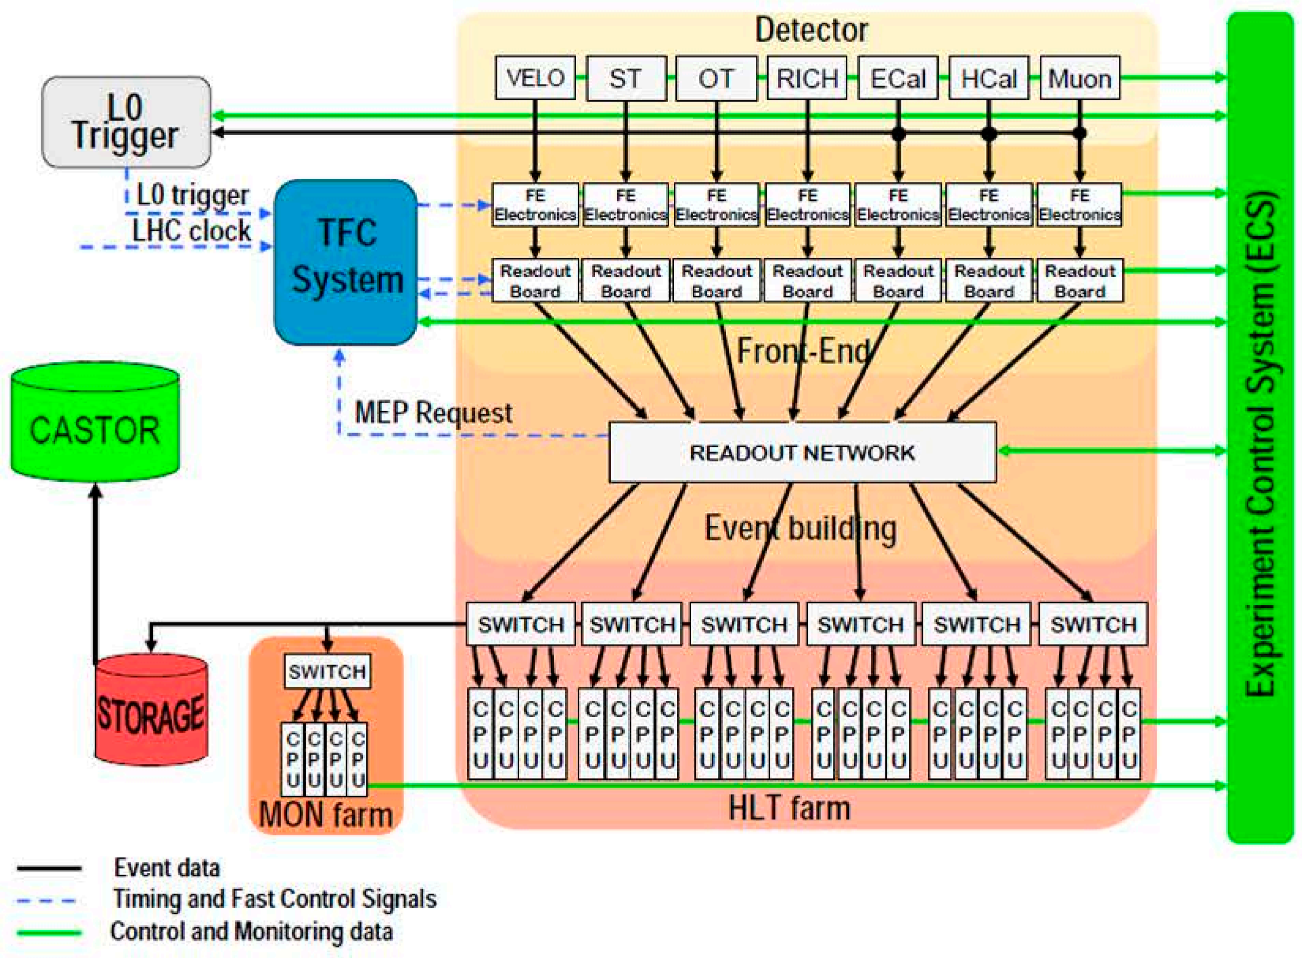

In order to study the relevant properties of such a dataflow, a non-disruptive monitoring setup based on a networking capable FPGA (NetFPGA) has been deployed. In our system, the NetFPGA takes care of filtering and timestamping traffic of interest, over four 1 Gb listening links. Packets are timestamped at the very earliest possible moment to minimise jitter caused by FIFOs.

In this fashion, *Multi-Event Packet* request rates, and data incoming from LHCb subdetectors sending nodes is timestamped and recorded for posterior analysis.

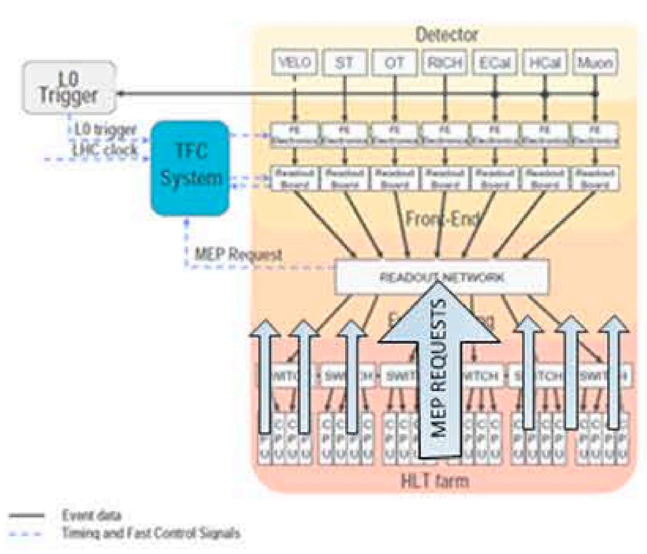

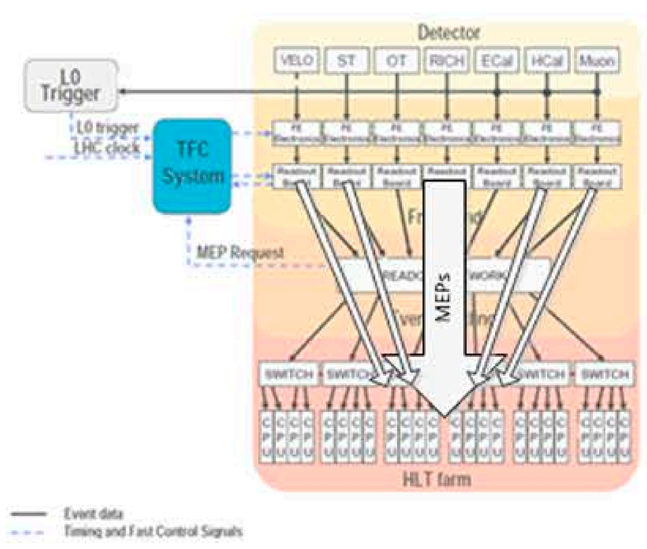

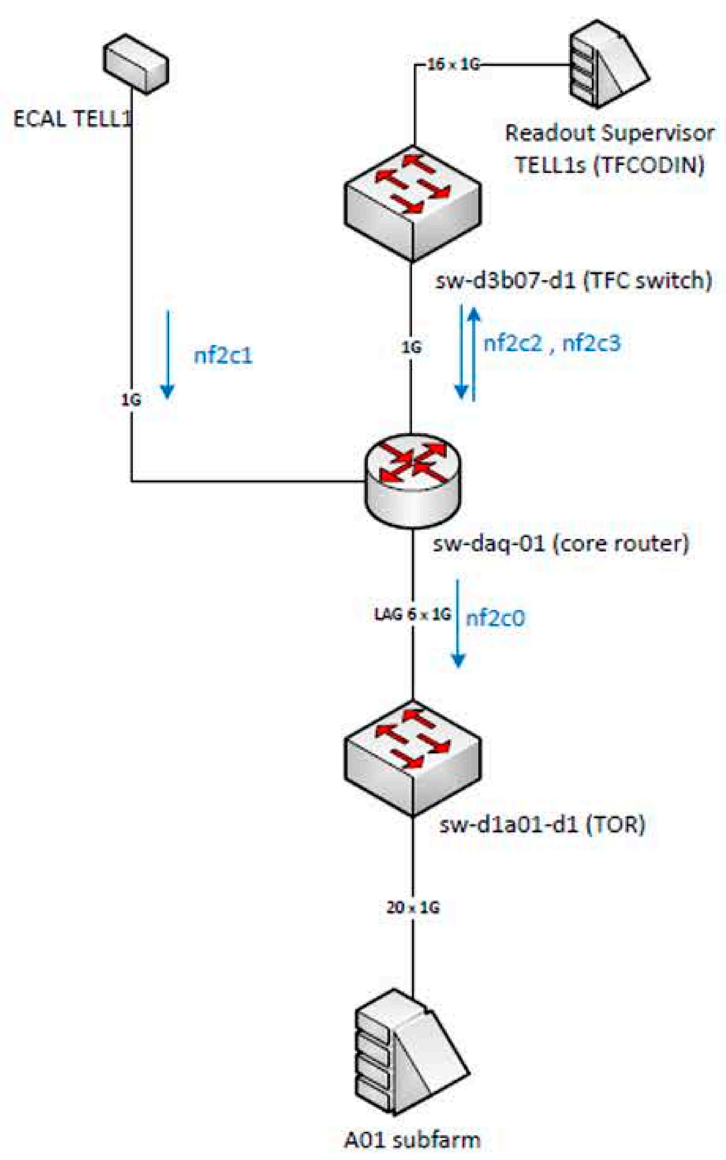

ANALYSIS

We observe a latency distribution that oscillates between 100 µs and 3 ms. For our application we have a constant rate of 60 GB/s unidirectional UDP-like traffic, sensitive to losses. We observe a static ingress buffer use of 20 MB per port in the analyzed runs of proton-proton and lead-ion-proton collisions. An increase in the luminosity at later stages of the LHC would imply higher needs in the buffering of our routers.

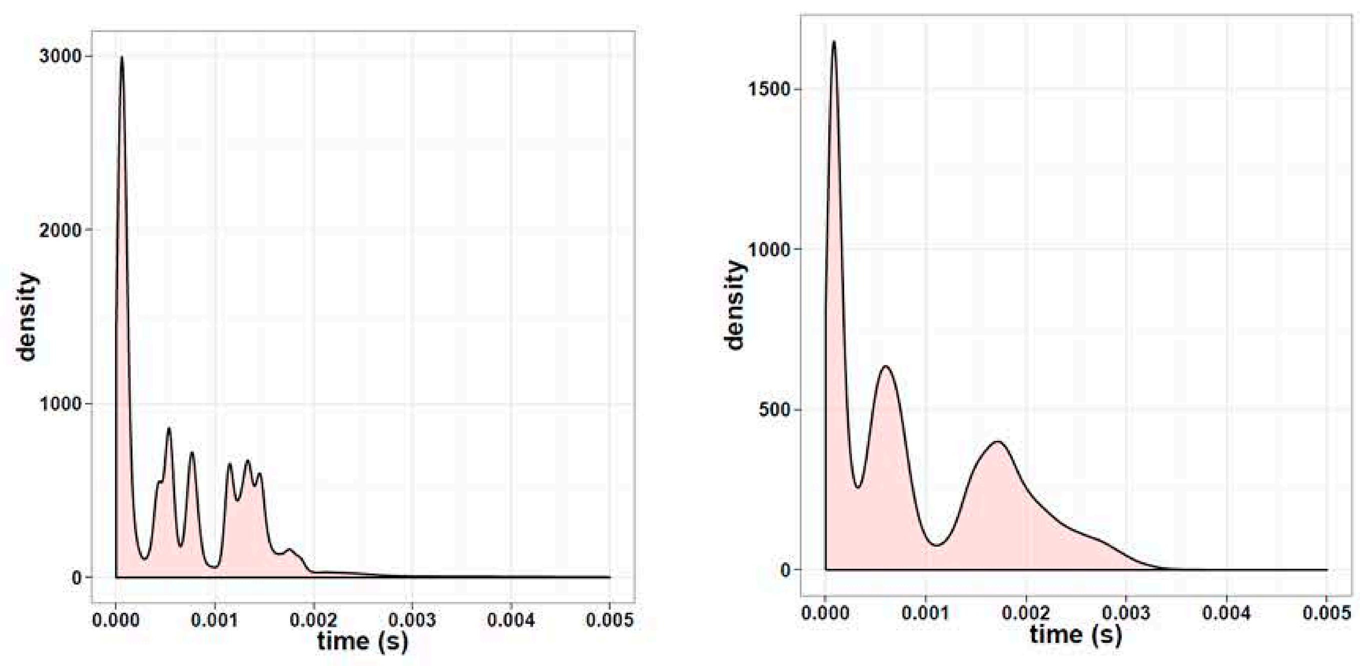

a. Latency distribution over no-run data.

b. Latency distribution over proton-proton run data.

The distribution of the load of MEP Requests is an accurate indirect indicator of the load of the router in the time. The latency spikes obseved in the latency distribution over time agree with the router load spikes. Irregularities in the traffic are observed related to the different run type states of the detector. Inside a run, the traffic shape is similar to the decreasing collision rate.

Latency spikes are an indicator of the load of our switches. A higher load can imply dead-time, when there are no more credits to accomodate the event load.

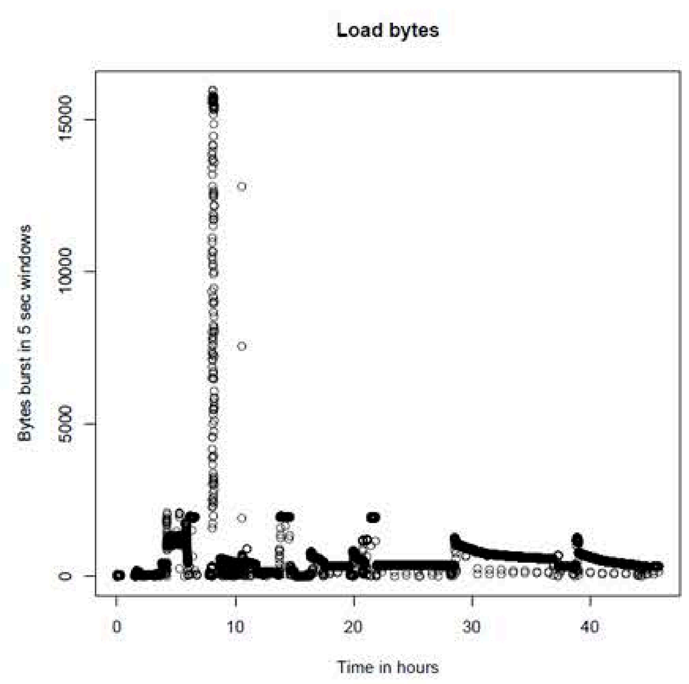

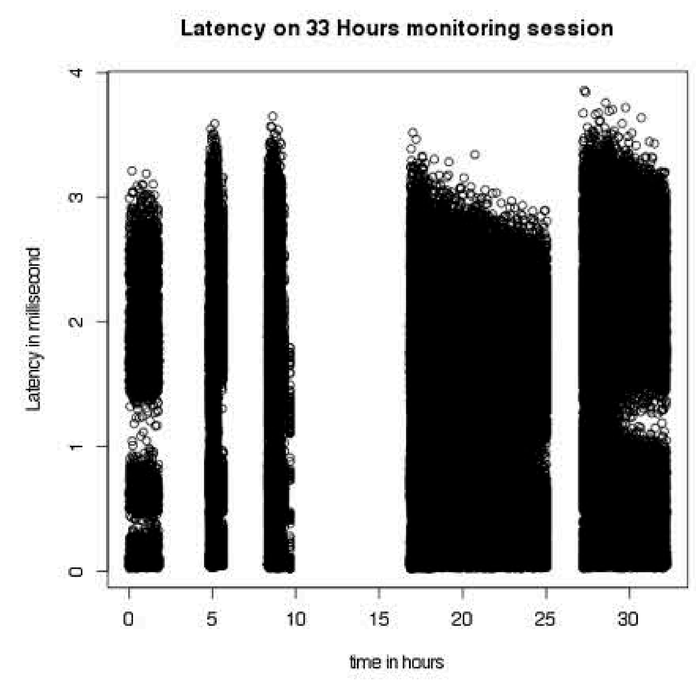

CONCLUSION

Our Netfpga monitoring setup accounts for precise network shaping in our DAQ network. Micro-burst and jitter effects are observed, which may be disruptive at higher data frequencies. Drops or backpressure effects originating from variable latencies can be observed with our monitoring system, which will play a decisive role at later stages of the LHC.

Gianni Antichi[1], Marc Bruyere[2,3,4], Daniel Hugo Cámpora Pérez[6]
Guoming Liu[6], Niko Neufeld[6], Stefano Giordano[1], Philippe Owezarski[2,3], Andrew Moore[5]
dcampora@cern.ch
1 University of Pisa, 2 CNRS, 3 Université de Toulouse, 4 DELL Inc., 5 University of Cambridge, 6 CERN