Performance of the DAQ System of the COMPASS Experiment ◇

L. Schmitt, H. Angerer, N. Franz, B. Grube, B. Ketzer, I. Konorov, R. Kuhn, W. Liebl, and S. Paul

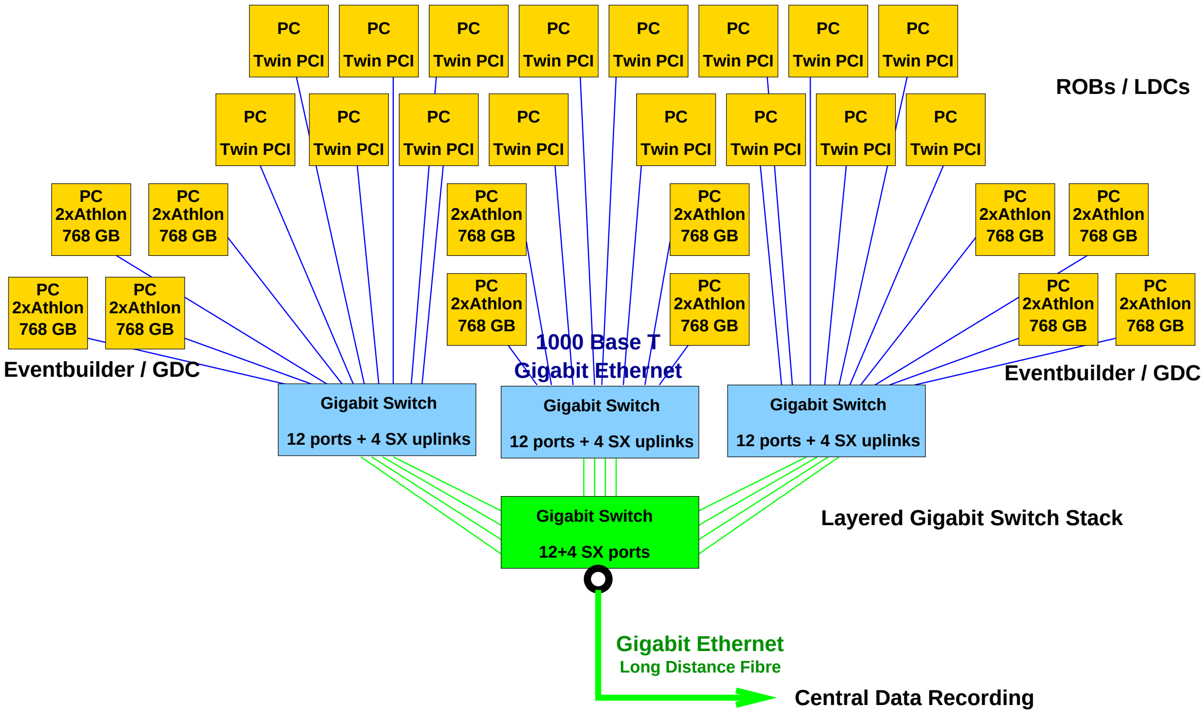

Fig. 1: The COMPASS DAQ system in 2002

In 2002 the COMPASS experiment at CERN [1] had its first 100 day long physics run. During this time its apparatus consisting of various detector types with a total of 191,000 electronics channels was read out with a small deadtime of 7 %. Our group at TUM is responsible for the development of the data acquisition system (DAQ) and components of the detector readout.

Overview. The mostly pipelined detector readout electronics feeds the data into concentrator modules, 9U VME boards called *CATCH*, which are developed at the University of Freiburg. In the same function, but specialized for the highly integrated readout of GEM [2] and Silicon [3] detectors, the *GeSiCA* board, developed at our group at TUM, is used.

¿From these boards the data are transferred via optical links to the DAQ computers. These links follow the *S-link* standard [4] developed at CERN. Several versions of optical links ranging from speeds of 100 MByte/s to 160 MByte/s are in use in COMPASS to adapt to the varying bandwidth requirements of different detectors.

The data received via S-link are buffered on PCI cards called *spill buffers* to profit from the accelerator's duty cycle to reduce the sustained data rate. This card was developed in our group at TUM and has 512 MB of SDRAM memory addressed via a Xilinx FPGA as FIFO. The computers housing the spill buffers (*readout buffers*, ROBs) transmit the data via Gigabit Ethernet to eventbuilder computers, which combine the sub-detector data to a complete event blocks and transfer it to CERN's central data recording system.

Hardware. In 2002 we used 16 computers as readout buffers with each 1 GB ECC SDRAM as main memory, two 866 MHz PIII CPUs and a 3COM Gigabit Ethernet interface. They have two PCI buses to read in the data on the first 32 bit bus and then transfer it through the second 64 bit bus via Gigabit Ethernet to the eventbuilder computers. Four spill buffer cards were mounted per readout buffer.

As eventbuilders 12 computers with two Athlon MP 1900+ processors, 1 GB ECC DDR-SDRAM and an IDE-RAID with net 640 GB per machine were used. In total 7.68 TB disk space were therefore available as buffer in case of problems with tape recording giving a safety margin of up to two days. All DAQ computers are connected to a set of three switches 3COM 4900 with twelve 1000 BaseT ports and four 1000 BaseSX uplinks. These uplinks connect the frontend switches to a backbone switch 3COM 4900SX which provides the necessary crosswise connectivity for all DAQ machines. To achieve a balanced configuration, 4 eventbuilder computers are connected to each frontend switch matching the number of uplinks to the backbone. ROB computers are connected according to their average output. The total eventbuilding bandwidth amounts to 12 Gbit/s .

◇ work supported by the BMBF

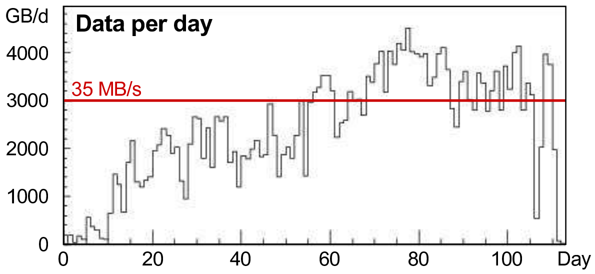

Fig. 2: Daily data rate in 2002

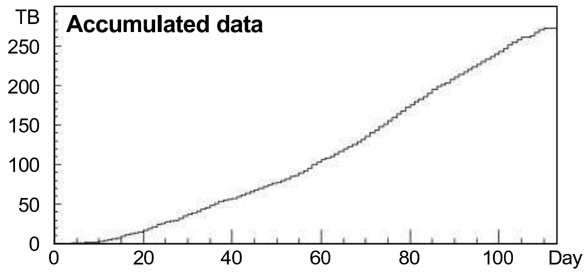

Fig. 3: Total accumulated data in 2002

Trigger Control System. To insure data integrity in the pipelined readout architecture each trigger has to be labeled by a unique event number. The synchronous transmission of the trigger and its identification to all detectors is done by means of optical signals distributed by the Trigger Control System (*TCS*) developed at TUM [5].

The TCS controller, a 6U VME board, encodes the trigger input and splits it into a synchronous trigger signal and an asynchronous event data block (event number, type etc.). The encoded pulses are transmitted by a high power laser and distributed via glass fibers and passive optical splitters to the detectors, a system developed at CERN for the LHC experiments [6].

The TCS receiver is a VME P2 transition board, which receives the optical signals and decodes the TCS information. It can be programmed via broadcasts to mask or block specific triggers from the attached CATCH module.

Apart from physical also artificial triggers can be generated and directed to individual detectors by the TCS system. They serve for calibration or noise measurements.

An additional function of the system is to provide separate deadtimes for up to eight independent trigger systems connected to the same readout electronics. This allows the sectioning of the DAQ system to test detectors independently before integration into the full system.

The TCS system provides a stable clock to the entire experiment. This is necessary for precise time measurements of particles passing through drift detectors and scintillators to correctly correlate detector hits at very high beam rate.

The system was operated successfully and at a very low time jitter of about 43 ps RMS between TCS receivers.

Performance. After tailoring and debugging of network and S-Link PCI drivers and tuning disk I/O of the system full production stability could be reached. In this phase the DAQ exceeded the maximum anticipated data rate by up to 50 % (see fig. 2). In average 25000 triggers were recorded per SPS cycle (5 s beam per 16.8 s) with data sizes between 40 and 50 kB. This corresponds to a sustained data rate of 60 MByte/s, an onspill rate of 200 MByte/s and a disk writing rate of 5 MBytes/s. During the beamtime in 2002 COMPASS recorded 260 TByte of data corresponding to about 5×10^9 events (fig. 3). In test runs a sustained rate of 120 MB/s to disk could be reached, without disk writing even 220 MB/s were possible, corresponding to 22 kHz trigger rate.

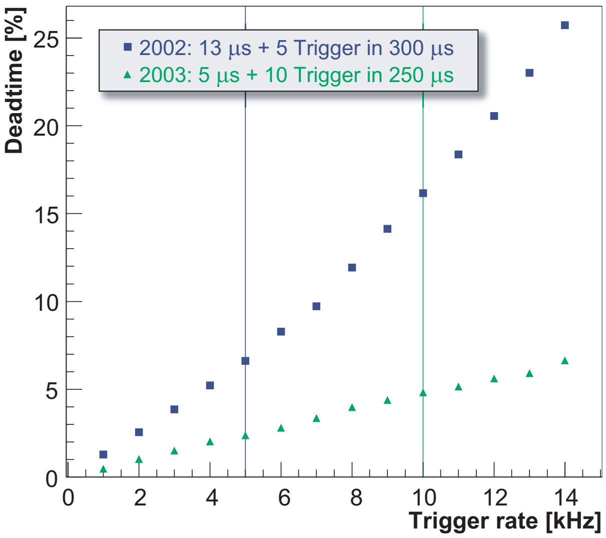

Fig. 4: Deadtime of the COMPASS DAQ, 2002 vs. 2003

Outlook for 2003. Further improvements in readout and stability will make a maximum trigger rate of 10 kHz onspill possible at a deadtime of only 4.5% compared to 7% at 5 kHz in 2002 (see fig. 4). However, since by now the main bottleneck of the system is the limited bandwidth for tape writing, this amount of data has to be filtered online and reduced to the nominal tape writing speed. Therefore at TU Munich an online filter algorithm is in development which is supposed to read data via a named pipe from the eventbuilder process. Since the eventbuilder computers were running at a CPU load of only about 20 % the remaining power can be used to reduce e.g. an initial trigger rate of 50000 events per spill this value to about half with a maximum time per event of approx. 3 ms.

References

[1] G. Baum *et al.* [COMPASS Collaboration], "COMPASS: A Proposal for a Common Muon and Proton Apparatus for Structure and Spectroscopy", CERN-SPSLC-96-14.

[2] B. Ketzer *et al.*, "Triple GEM Tracking Detectors for COMPASS", IEEE Trans. Nucl. Sci. **49** (2002) 2403

[3] R. de Masi *et al.*, "Present Status Of Silicon Detectors in COMPASS", Contribution to the 9th European Symposium on Semiconductor Detectors, Schloss Elmau, June 2002, proceedings submitted to Nucl. Instr. and Meth. A

[4] H.C. van der Bij, R.A. McLaren, O. Boyle and G. Rubin, "S-LINK, a data link interface specification for the LHC era", IEEE Trans. Nucl. Sci. **44** (1997) 398

[5] I. Konorov *et al.*, "The Trigger Control System for the COMPASS Experiment", contribution to the IEEE Nuclear Science Symposium, San Diego 2001, Conference Record (2002).

[6] B. G. Taylor [RD12 Project Collaboration], "TTC distribution for LHC detectors", IEEE Trans. Nucl. Sci. **45** (1998) 821