Performance of the CMS Silicon Tracker Front-End Driver

G. Iles [1], R. Bainbridge [1], D. Ballard [2], I. Church [2], E. Corrin [1], J.A. Coughlan [2], C.P. Day [2], C. Foudas [1], E.J. Freeman [2], J. Fulcher [1], W.J.F. Gannon [2], G. Hall [1], R.N.J. Halsall [2], J. Leaver [1], M. Noy [1], M. Pearson [2], M. Raymond [1], I. Reid [3], G. Rogers [2], J. Salisbury [2], S. Taghavi [2], I.R. Tomalin [2], O. Zorba [1]

[1] Imperial College, London, UK
[2] CCLRC Rutherford Appleton Laboratory, Oxfordshire, UK
[3] Brunel University, London, UK

gm.iles@imperial.ac.uk

Abstract

The CMS Silicon Tracker Front-End Driver (FED) is a 9U 400mm VME64x card which processes the raw data generated within the Silicon Tracker by the APV25 readout ASICs. The processed, zero-suppressed, data is then sent to the Data Acquisition System (DAQ). The first 2 FEDs were made at the beginning of 2003 and since then a further 15 FEDs of this type (FEDv1) have been manufactured.

All hardware modifications to the FEDv1 design have now been completed and a new iteration of the board produced, called the FEDv2, which is expected to be the final version. The firmware and software development is close to completion. The performance of a FED in the laboratory is presented.

I. INTRODUCTION

The CMS Silicon Tracker Front-End Drivers (FEDs) are 9U 400mm VME64x cards (fig. 1) [1] that each process the raw data from 192 APV25 silicon readout ASICs [2], equivalent to roughly 0.2% of the Tracker. An APV25 samples the signal from 128 silicon strips every bunch crossing and stores the data in an analogue pipeline memory. Upon receipt of a L1A (Level-1 Accept) trigger the analogue data from the corresponding bunch crossing is read out serially at 20MHz. The data are multiplexed with that of another APV25, converted to optical form via an Analogue Opto-Hybrid (AOH) [3,4] and routed out of the Tracker to a FED in the counting room via optical fibres.

Each FED receives data from 96 optical fibres. The FED converts the optical signals to electrical, digitises the data to 10bit precision at 40MHz and then processes the data in large FPGAs. The processing removes pedestals and common mode noise. The input data rate of 3.4 GB/s is then reduced by more than an order of magnitude by firmware that extracts clusters of hit strips from the data stream. Finally the data are formatted into events and sent to the CMS DAQ via S-LINK64 [5].

The first 2 FEDs, called FEDv1, were produced at the beginning of 2003 and, after initial tests confirmed no major problems existed, a further 3 were delivered shortly afterwards. A subsequent batch of 6 boards delivered in the autumn of 2003 had manufacturing errors. A further 6 boards delivered in the spring of 2004, by an alternative supplier, functioned correctly [6].

Throughout 2004 FEDs have been distributed to CERN, Pisa, and Lyon for system integration tests and module testing. In June this year a beam test took place at CERN with 4 FEDs driven by 252 optical channels. This represents about 1% of the final Silicon Tracker.

In the last month the next iteration of the FED, the FEDv2, has returned from manufacture and is under test. Initial results indicate that no problems exist and therefore it is planned to manufacture a further 20 FEDv2s by the end of this year.
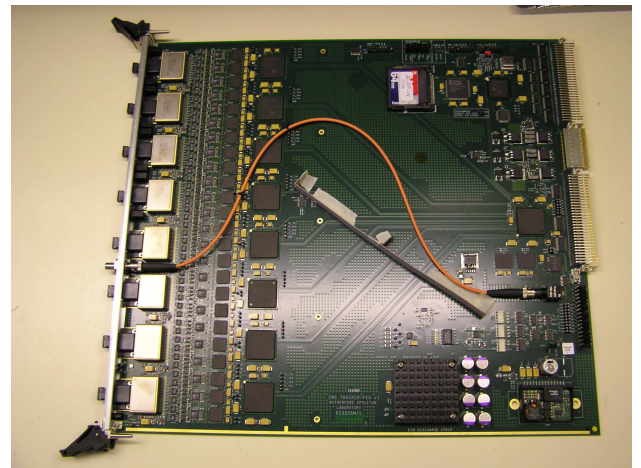

Figure 1: The FED. The 8 front-end units are visible on the left of the picture. They each contain a 12 way optical receiver (OptoRx), 12 buffers, 6 dual ADCs and a large data processing FPGA. The front-end units are numbered from 0 to 7 starting from the top of the board. The air deflector discussed in section III is visible in the centre of the image.

At the beginning of 2003 only the core features of the FED were available in firmware and software. During 2003/4 these have been extended so that the FED functionality is now close to that required for operation in CMS. The latest developments include: the zero suppression algorithms used to locate clusters of hit strips; remote loading of new firmware over VME; fast feedback

status signals which enable the Trigger Control System (TCS) [7] to control the L1A rate and determine the FED status; the Trigger Timing and Control (TTC) system; the S-LINK64 link from FED to DAQ. The data path from the FED to an S-LINK64 receiver card has been tested and a lower limit set on the reliability of the link. In addition to the firmware development there has also been a substantial software development with the first official release this summer [8].

The temperature of the FED has been studied in conditions that mimic as closely as possible the environment of the counting room racks. A solution to provide extra cooling to the analogue opto/electrical front-end is presented.

The S-LINK64 tests and high rate tests are still ongoing and therefore the results presented here merely show what we have currently achieved. They are not yet an exhaustive test and it will be several more months before such a test is complete.

II. LABORATORY TESTS

The laboratory test setup (fig. 2) is comprised of a FED driven by a FED Tester Ensemble (FTe). The latter is made up of 4 FED Testers [9], with one acting as the master and the remaining three as slaves. The FTe provides 96 analogue optical signals with data similar to that expected in CMS. To simplify the design and manufacture of each FED Tester, as well as keeping the cost down, each FED Tester generates only 3 unique channels. A cross-point switch allows the user to select which of these 3 channels is used to drive each of the 24 optical outputs. In a FTe there are 4 FED Testers and therefore 12 unique channels allowing all 12 channels of a FED front end unit to receive different data.

The temperature of the 8 AOHs is controlled to within 1°C on the condition that the temperature in the laboratory does not fluctuate by more than ~5°C during a test. This ensures that the optical signal remains stable and that the signal reaching the FED remains within the ADC sampling range.

An inbuilt Trigger Control System (TCS) tries to replicate that expected in CMS. It not only provides a clock and L1A, but also Bunch Counter 0 (BC0), Event Counter Reset (ECR) and Orbit Counter Reset (OCR). These would normally be supplied from the local or central TCS via the TTC system. At present these commands are not fed into a TTC system and instead only the clock and L1A are routed to the FED using custom pins on the VME backplane.

Data are transmitted from the FED to the DAQ via an S-LINK64 link. The S-LINK64 transmitter card resides on a 6U Transition card located in the back of the VME crate.

The software currently written fully calibrates the FTe AOH optical transmitters and FED timing. It can then generate fake multiplexed APV25 frame data with a hit occupancy selected by the user. At present only single strip

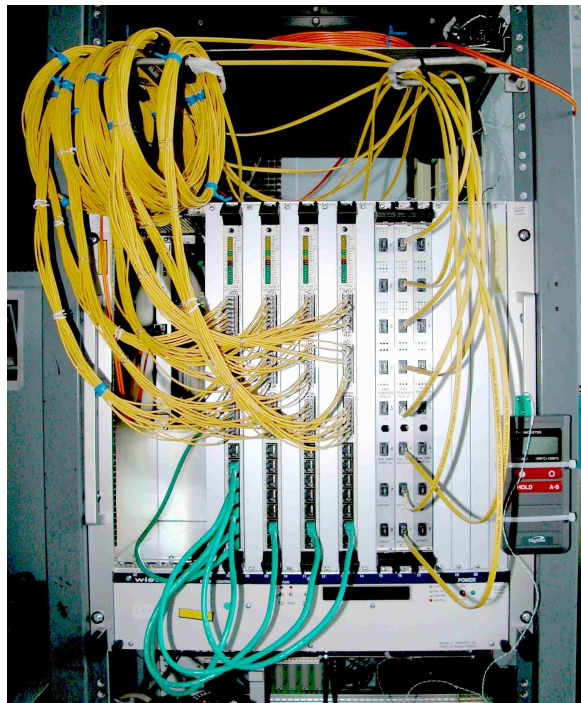clusters are created, however this will change in a future iteration of the software.

Figure 2: Laboratory test set up. A FED Tester Ensemble (FTe) consisting of 4 FED Testers (left) driving a FED sandwiched between two neighbouring FEDs. The latter is an attempt to mimic the airflow and heat dissipation in a fully populated standard LHC crate.

III. TEMPERATURE MONITORING

In the PCB region closest to the front-panel there is a very high density of high bandwidth optical, analogue and digital components that consume the bulk of the ~80W of power dissipated by the FED. To verify that the temperature of FED components will remain within specification it was necessary to replicate the environmental conditions that can be expected at CMS. The FED under test was sandwiched between two adjacent FEDs to simulate the airflow and heat dissipated by neighbouring FEDs.

The temperature was monitored with on-board sensors located directly beneath, but on the opposite side of the PCB to the OptoRx units. Thermocouples were attached to the OptoRx units to verify the measurements reported by the FED and also placed on the ADCs and differential buffers. The FED was operated as it would be in CMS with all inputs driven. All temperatures have been corrected for an input air of 20°C.

The temperature of the OptoRx units near the top of the board reached 67°C and 59°C for fan speeds of 2460rpm and 3480rpm respectively. The latter was the maximum fan speed achievable. The nominal fan speed is 3000rpm.

The ADCs reached 82°C and 72°C for fan speeds 2460rpm and 3480rpm respectively. These temperatures were close to the design limits of 70°C for the OptoRx and 85°C for the ADC. To aid cooling in this region an air deflector was added to the centre of the board (fig. 1). This reduced the maximum temperature of the OptoRx units to 55°C at 2460rpm and 49°C at 3480rpm (fig. 3). The ADC maximum temperature was reduced to 61°C at 2460rpm and 54°C at 3480rpm.
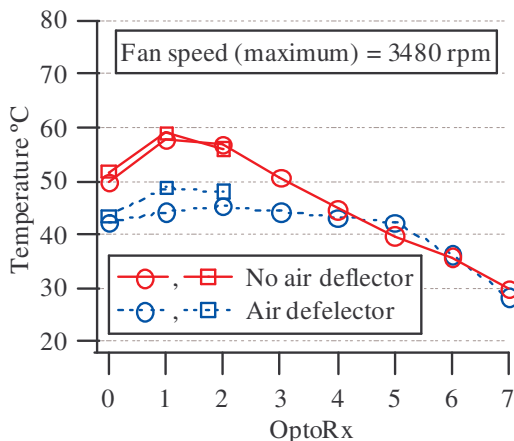

Figure 3: Temperature of the FED for a fan speed of 3480rpm. Data points indicated by circles were obtained by onboard sensors placed on the opposite side of the PCB, but directly beneath the OptoRx units. The squares were obtained by thermocouples bonded directly to the OptoRx units.

IV. S-LINK64 TESTS

In August the S-LINK64 Transition card (fig. 4) returned from manufacture. It is a simple 6U VME card which holds the S-LINK64 transmitter mezzanine [5] and buffers the S-LINK64 data to and from the FED. The transition card also converts the FED single ended fast feedback (throttle) signals to LVDS before routing them through an RJ45 connector to the TCS via Fast Merging Modules (FMM) [10] (FMM cards were not available for these tests).

The data path from the FED to the S-LINK64 receiver via the Transition card and S-LINK64 transmitter card has been verified by generating test patterns in the FED back-end FPGA and verifying them in the S-LINK64 receiver host PC. The test patterns are generated by sending a 100kHz repetitive trigger. However, because the host PC can only verify data at 28MB/s this leads to the S-LINK64 receiver asserting the back pressure signal. The FED responds by not sending any more events down the S-LINK64 and consequently its back-end buffer starts to fill. This causes it to transmit "WARN" followed by "BUSY". The TCS responds by initially lowering the L1A rate and then stopping the L1As completely. In this way the 100kHz un-throttled rate becomes in practice 17kHz. The S-LINK64 was operated with throttling for 10.8 hours during which time 1TB of data was transmitted over a 1.5m cable without error. To ensure an error rate of less than 1

per week per FED at CMS it would be necessary to run this test for 146 days.

To test the maximum rate at which the FED could send data down the S-LINK64 most of the verification checks were removed. In this way the host PC was able to process the S-LINK64 fragments far more quickly and FED and S-LINK64 were able to handle the 100kHz rate without asserting "WARN" or "BUSY". The event size was slowly increased until the FED started asserting "BUSY". By probing the S-LINK64 back pressure signal we were able to determine that the FED was throttling because the S-LINK64 was asserting back pressure rather than because the FED was delaying transmission of an event packet. The maximum data rate achieved was 469MB/s, whereas the theoretical limit for the S-LINK64 running at 80MHz is 640MB/s. Note that in CMS the S-LINK64 receiver is only intended to operate at up to an average rate of 200MB/s or peak rate of 400MB/s, where the average rate is over timescales of a second [5].
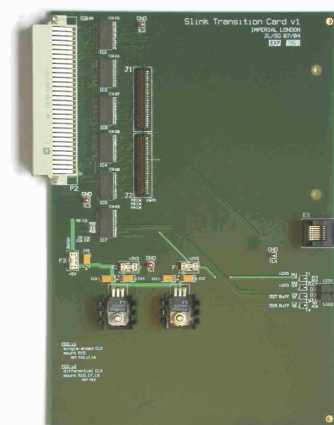

Figure 4: The S-LINK64 Transition card without the S-LINK64 transmitter daughter card mounted.

V. HIGH RATE TESTS

One of the objectives for laboratory testing is to operate the FED in its most challenging modes and to verify that the data it sends to the S-LINK64 receiver matches what we expect. At present the verification software is not yet complete. However, we are able to drive the FED with the type of data expected when the LHC is operating at high luminosity.

As a first step towards this goal we have been checking to see how the FED buffers and internal high speed links perform as we increase the mean strip occupancy in the fake data and thus the size of the zero suppressed events.

The position of both the data buffers and high speed links within the Tracker are shown in fig. 5. The first buffer is on the APV25 readout chip. All the APV25 readout chips operate in synchronisation and therefore the

APV25 buffers can be modelled outside the Tracker by the APV25 Emulator (APVE) [11]. Should they become close to full the APVE will assert "WARN" followed by "BUSY" on the fast feedback system and the local or central TCS will respond by lowering the L1A rate. Data from a multiplexed pair of APVs are sent down an analogue optical link. At the FED the data are digitised at 40MS/s to 10bit precision, the APV frames detected and processed. The cluster 8bit data from a multiplexed pair of APVs are then buffered in a 2kB memory internal to the FPGA before being sent down an 80MB/s link to the back end FPGA. There is an 80MB/s link for each of the front-end units. A single FED channel shares this link with the another 11 channels in a front-end unit. At the back-end the data are again buffered, but this time by a 2MB Quad Data Rate SRAM memory external to the FPGA which can be read and written simultaneously at up to 640MB/s. The data are then sent to the DAQ over the S-LINK64. The latter has a theoretical maximum data rate of 640MB/s when operated at 80MHz, however in CMS it is expected to operate this link at 200MB/s on average and 400MB/s peak.

Our goal was to evaluate the buffering and high speed operation of the FED as we increase the mean strip occupancy. However, the S-LINK64 receiver is unable to match the capability of the FED. The FED was therefore instructed to ignore the S-LINK64 back pressure signal and as a consequence the data must be discarded.

Ideally all FED buffers would assert "WARN" followed by "BUSY" as they become full. However, if the TCS were slow to respond to the request to lower or stop sending L1As or perhaps events were still pending in the

APV25 buffers then the FED buffers might still face the prospect of buffer overflow. To retain synchronisation in CMS it is required that if FED buffers are close to overflow the FED should drop the event payload and buffer only the event header with a flag to indicate the data was dropped. There should be sufficient remaining space in the buffer for these smaller events, which will be quickly despatched over the S-LINK64, thus allowing the FED to recover. Note that should the S-LINK64 data rate be permanently reduced to below that of empty events the FED will inevitably suffer buffer overflow and thus loss of synchronisation. This buffer protection is not yet fully implemented, although the back-end buffer will already assert "WARN" followed by "BUSY". Consequently, we have had to use a different approach.

For each test up to 100k L1As, with a 100kHz Poisson distribution, are sent to the FED. The FED status register is constantly monitored for error conditions such as buffer overflow. When this occurs we record the number of L1As sent, reset the FED and then repeat the test 5000 times for each mean strip occupancy setting. Thus a FED L1A lifetime is determined as a function of occupancy (fig. 6). The FED is operated in zero suppression mode (i.e. automatically finding APV25 event frames and extracting the hit strips in the form of strip clusters). The fake data generated by the FTe for a given mean strip occupancy consists of single strip clusters. This creates the largest event sizes because of the way in which the data are formatted. To simplify analysis the number of hit strips per APV25 is not allowed to randomly fluctuate, but is held constant. We believe that this is a valid approximation
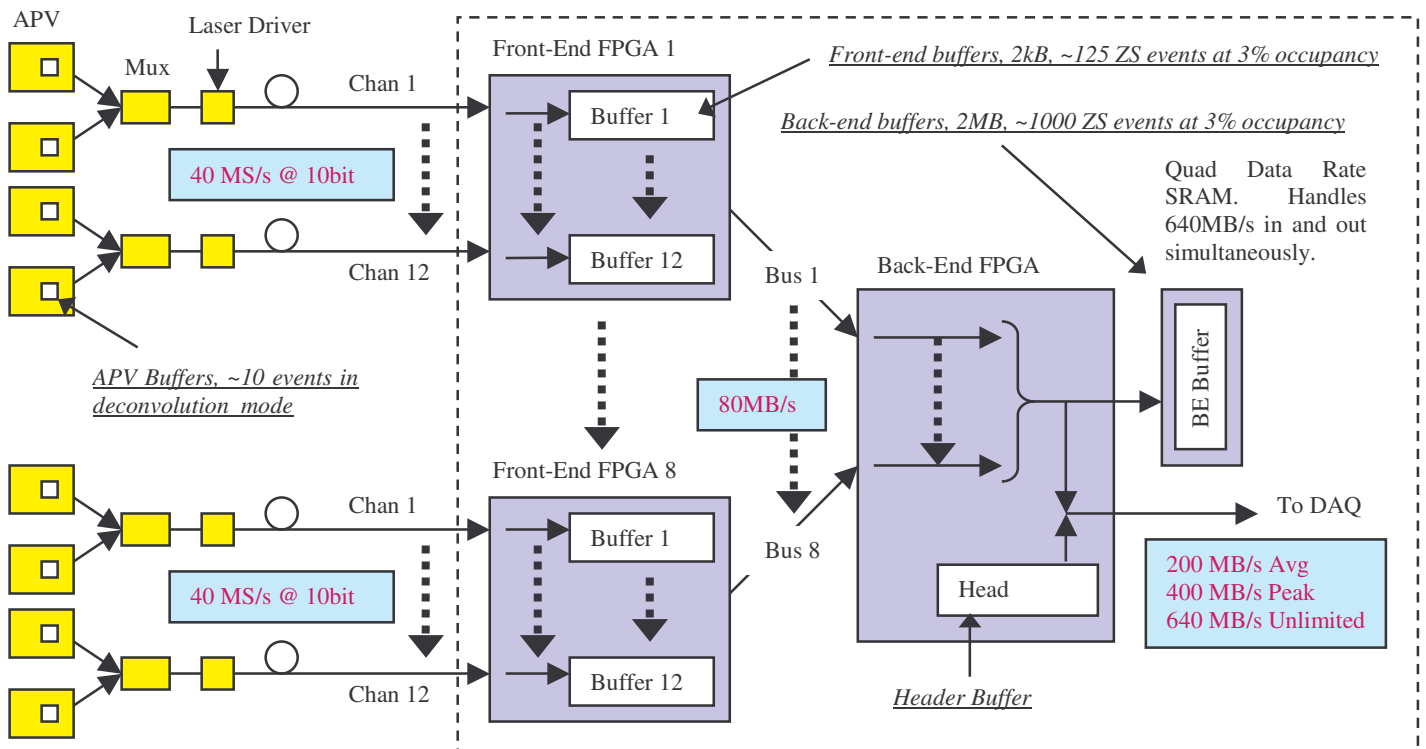

Figure 5: Diagram showing the data links and buffers within the Silicon Tracker

given the very large size of the buffers in the FED in zero suppression mode, which will tend to absorb any fluctuations in event size.
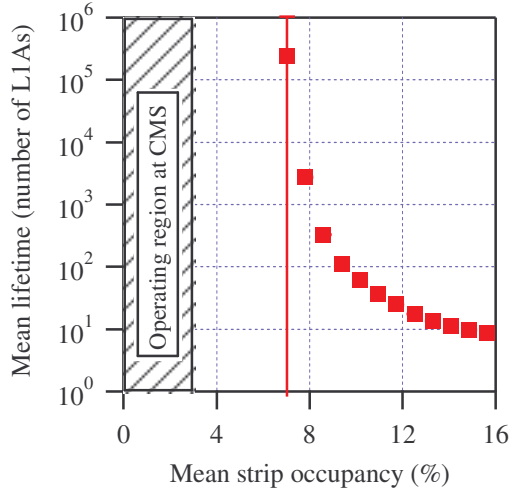

Figure 6: The mean FED lifetime as a function of mean strip occupancy. The fake data used to generate this plot uses single strip clusters. The zero suppressed events currently contain a substantial amount of debug information. Both these features create larger events than would be expected in CMS thus increasing the data rate and making the FED lifetime shorter (i.e. this is the worst case scenario).

The FED did not lock up below an occupancy of 6.25%, which is in agreement with previous studies [12]. In practice the limit will be set by the data acquisition system. This limits each FED or pair of FEDs sharing a Front-end Readout Link (FRL) to 200MB/s whereas each FED has an intrinsic bandwidth of 640MB/s. Furthermore, there are only ~256 FRLs available for the entire Tracker [13] and thus a large fraction of the 440 FEDs in CMS must share a FRL. Previous studies [13, 14] have shown that at CMS the mean strip occupancy is below 3% and the data rate from the Tracker can be handled by the DAQ envisaged.

VI.   FUTURE

It is planned to manufacture a further 20 FEDv2s by the end of this year. A tender exercise is already underway for the procurement of approximately 500 FEDs including spares for the final system [6]. A FED test system based on a LabVIEW graphical interface, which does not require optical inputs should be completed shortly for in-house industrial testing [6]. FED software is close to completion [8] and a substantial amount of the system integration software already exists.

Testing will continue in both the laboratory and with another test beam at CERN in October. It is planned to perform more thorough high rate tests in the future with the FTe where not only do we operate at high rates in zero-suppressed mode, but also perform a detailed cross-check between the fake data delivered to the FED and that received through the S-LINK64, however this is still some months from completion.

VIII.   REFERENCES

[1]   J. Coughlan et al., "The CMS Tracker Front-End Driver", Ninth Workshop on Electronics for LHC Experiments, CERN-LHCC-2003-055

[2]   M. Raymond et al., "The CMS Tracker APV25 0.25μm CMOS readout chip", Sixth Workshop on Electronics for LHC Experiments, CERN-LHCC-2000-041.

[3]   J. Troska et al., "Optical Readout and Control Systems for the CMS Tracker", IEEE Transactions on Nuclear Science, Volume: 50 Issue: 4 , Aug. 2003, 1067-1072.

[4]   M. Friedl et M. Pernicka, "Experience with Large-Scale Industrial Production Considering the CMS Tracker Analog Optohybrids", These proceedings.

[5]   V. Brigljevic et al., "FEDkit : A Design Reference for CMS Data Acquisition Inputs", Ninth Workshop on Electronics for LHC Experiments, CERN-LHCC-2003-055

[6]   J. Coughlan et al., "The Manufacture of the CMS Tracker Front-End Driver", These proceedings.

[7]   J. Varela, "CMS L1 Trigger Control System", CMS Note, 2002-033

[8]   J. Fulcher et al., "Software and DAQ for the CMS Silicon Tracker Front End Driver", These proceedings.

[9]   G. Iles et. al., "The CMS Tracker Front-End Driver", Ninth Workshop on Electronics for LHC Experiments, CERN-LHCC-2003-055

[10]   A. Racz et al., "The Final prototype of the Fast Merging Module (FMM) for readout status processing in CMS DAQ", These proceedings.

[11]   G. Iles et al., "The APVE emulator to prevent front-end buffer overflows within the CMS silicon strip tracker", Eighth Workshop on Electronics for LHC Experiments, CERN-LHCC-2002-034

[12]   E. Corrin, "Development of Digital Readout Electronics for the CMS Tracker", Thesis, Imperial College, London, March 2003

[13]   I. Tomalin, "Expected Data Rates from the Silicon Strip Tracker", CMS Note, 2002-047

[14]   I.Tomalin, "On Calibration, Zero Suppression Algorithms and Data Format for the Silicon Tracker FEDs", CMS Internal Note, 2001-02