Simulation of the pixel detector Module Control Circuit

David Calvet[1]

Centre de Physique des Particules de Marseille
163 avenue de Luminy, case 907, F-13288 Marseille cedex 9

Abstract

A tool to simulate the effect of the pixel detector Module Control Circuit on the pixel detector data flow has been written. Some efficiency studies are shown and a parameterization of this efficiency has been performed in order to study the influence of the circuit performances on the physics results.

---

[1]calvet@cppm.in2p3.fr

Introduction

The ATLAS pixel detector is made of hybrid modules containing a sensor, 16 Front-End (FE) integrated circuits and 1 Module Control Circuit (MCC), which is the interface between the FE circuits and the ReadOut Drivers (ROD) [1]. The aim of this note is to describe a tool (MCCSim) designed to simulate the internal architecture of the MCC and to study the effect of this circuit on the pixel detector data flow. A performance study has been performed using MCCSim and results are shown. Finally, a parameterization of the MCC effect on the pixel detector data has been implemented in the ATLAS detector simulation software (ATLSIM), to make possible the inclusion of the circuit performances in the physics analyses.

1 MCCSim description

The overall structure of the MCCSim tool is shown on figure 1. The input for MCCSim is a data file which simulates the output of the 16 FE circuits of a single pixel detector module. This file is read according to the trigger "signal" generated by a dedicated module and the data are fed into the MCC simulator, where event building is performed before writing the data into an output file.

1.1 Trigger generation

Every new Beam CrossOver (BCO), the level-1 trigger signal (LV1) is generated in MCCSim in four steps :

1. the LHC beam structure [2] is reproduced and when this new BCO corresponds to a beam gap (no proton bunch), the LV1 signal is false and the procedure stops,

2. if there were less than 5 BCOs since the last LV1 true, the LV1 signal is false and the procedure stops, in order to reproduce the systematic trigger deadtime [3],

3. the so-called "complex" dead-time algorithm [3] is implemented : each time a new LV1 true signal is generated, the content of a bucket is increased by a fixed value (400). This bucket is leaking at a constant
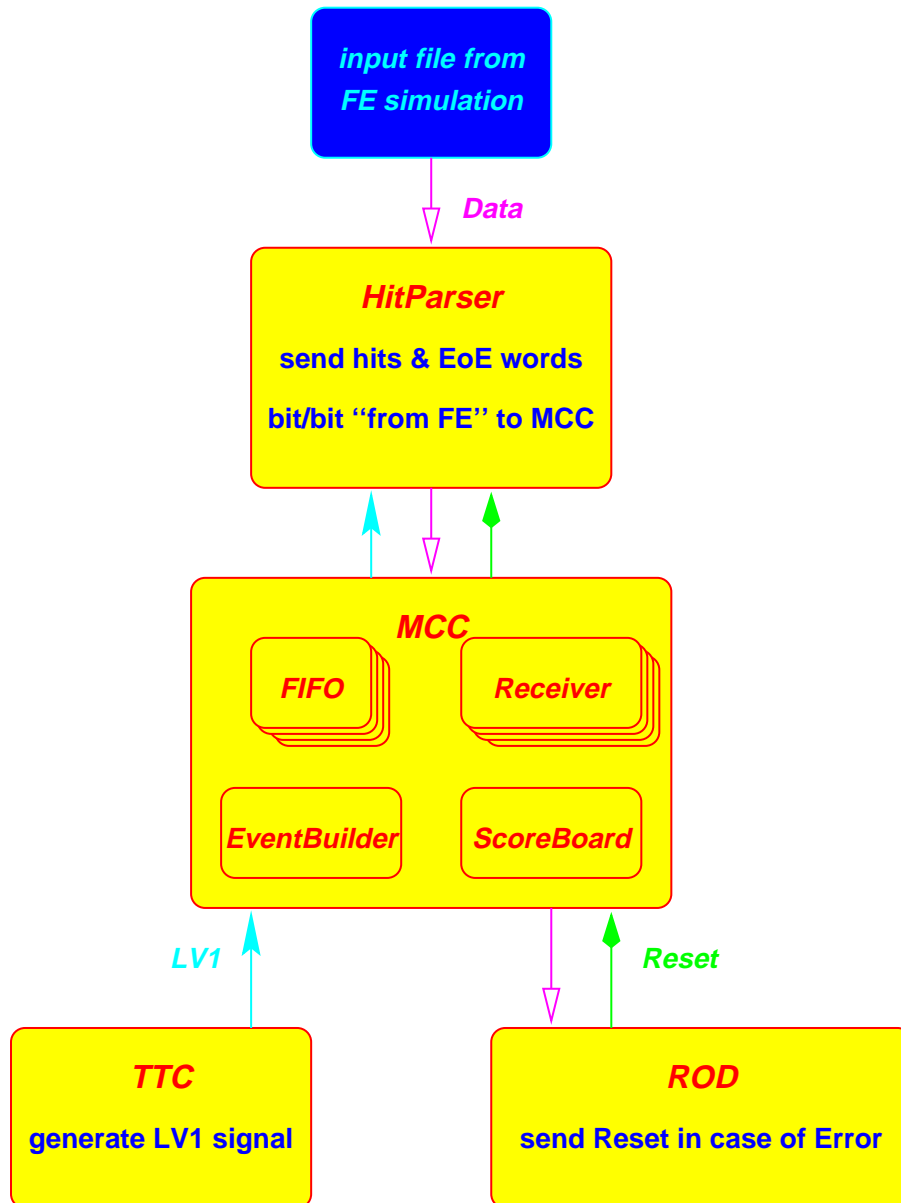
Figure 1: MCCSim software structure.

rate of 1 per BCO. When the bucket content is greater than 3200, the LV1 signal is false and the procedure stops,

4. a new LV1 signal is generated using a random generator; the probability for this signal to be true is fixed and set according to the parameter called *LV1 Target Rate*.

Due mainly to the LHC beam structure and also to the generated deadtimes, the average LV1 rate is less than the *LV1 Target Rate*. For example, a *LV1 Target Rate* of 100 kHz leads to an average LV1 rate close to 75 kHz.

1.2 LV1 filtering by MCC

Each time a new LV1 true signal is generated, it is sent to the MCC. Then, the MCC sends a trigger signal to the FE circuits, or it can reject it, in two cases :

- a pending event counter keeps track of how many events are still to be built by the MCC : the trigger signals where sent to the FE circuits but the events have not been sent to the ROD yet. If this pending event counter is greater or equal to the limit (a parameter called *Max Pending LV1* in MCCSim), then the new LV1 true signal is rejected. This limit is 16 in the present design.
- when the MCC is in ERROR state (see section 1.6), the new LV1 true signal is rejected.

1.3 Hit parsing from "FE" to MCC

As soon as the trigger signal is received by the FE circuits, they start to send data to the MCC. In MCCSim, a module called HitParser is responsible for the FE-to-MCC data flow handling. When a new LV1 true signal is received by HitParser, it loads data for a complete new event from the data input file and sends the hit and the End of Event (EoE) words to the MCC input. Subsequently, these words are transferred to the MCC input at a fixed rate of one bit per BCO; the word length is adjustable and is equal to the sum of two parameters called *FE Word Length (ToT excl)* and *ToT Length*. In the present design [1], the full word length is 26 bits for a ToT (Time over Threshold) of 8 bits. The data of all the 16 FE circuits are sent to the MCC in parallel.

1.4 MCC Receiver

The data words from the 16 FE circuits are received by 16 independent Receivers inside the MCC where they are de-serialized and written to 16 independent FIFOs. When an EoE word is received, the counter corresponding to this Receiver is incremented in the MCC Score Board, so that the Score Board keeps track of how many complete events are present in every FIFO (see [1] for details).

The Receiver is also responsible for the handling of problems that may occur when the FIFO is full and data is still arriving from the FE. Two distinct cases are possible :

- WARNING condition : if the FIFO is full but the last word in the FIFO is a word belonging to the current event, then all hits from the event are removed from this FIFO (but not from the other FIFOs), and a special word (WNG#1) is written instead of the usual EoE. All hit words which may come from this FE are lost until the EoE is received for the current event.

- ERROR condition : if the FIFO is full but the last word in the FIFO is a word belonging to the previous event, no place is available to write a WNG word, since there are no hit words from the event to remove. Then, this Receiver stops working until the next MCC reset and the MCC goes into the ERROR state (see section 1.6).

The size of the FIFO (number of words it can contain) is set in MCCSim by the parameter called *FIFO Depth*. In the present MCC design, the FIFOs can contain 32 words each.

1.5 MCC Event Builder

The main task of the MCC is to group individual FE events into a complete module event. This task is performed by the Event Builder with help of the Score Board [1]. When a new complete event is ready (EoE words from the 16 FE circuits have been received and recorded in the Score Board), the Event Builder starts to write the new module event consisting of a header word (the trigger number), some data words and a trailer word (see [1] for details). In MCCSim, the length of these words is adjustable using three parameters :

4

- *MCC Word Length* for the words which are not hit words (trigger number, FE number, flags). This length is 9 in the present design, including the header/sync bit.

- the sum of *MCC Hit Length (ToT excl)* and *ToT Length* for the hit words. This length is 22 for a ToT of 8 bits in the present design, including the sync bit. The length of the trailer word is equal to the hit word length plus one.

The module event is serially sent to the ROD via an optical fibre link. The speed of this link is set in MCCSim by the parameter *Optical Link Speed* which should be a multiple of 40 Mbit/s. The MCCs in the B layer will transmit their data using two such optical links per MCC, so that the *Optical Link Speed* should be set at twice the value of other layers.

1.6 Error handling and ROD behaviour

In the case of an ERROR condition due to an unrecoverable full FIFO (see section 1.4), the MCC goes into the ERROR state where new LV1 true signals are rejected. Then, the Event Builder continues to build events which are in the FIFOs and are still arriving from the FE circuits, but skips FIFOs which are in ERROR state. Also, the Event Builder writes in each module event a special ERR flag word, in order to keep track of the data loss. When all pending events have been built and sent, the MCC stops working and waits for an external reset command from the ROD.

A problem arises from the way the ROD should react to this ERROR condition. In MCCSim, just before going to the "waiting for reset" state, the Event Builder writes a special module event consisting of a false trigger number, a special flag word (called RST) and a trailer word. The RST word is a reset request to which the ROD should respond by sending the reset command to the MCC. The time between the end of this special event and the MCC being back into data taking mode is set by a parameter called *ROD Reset Latency*, which includes the time spent by the signals in the optical fibres.

2 MCC induced data losses

Data can be lost because of the MCC in two different cases : the LV1 true signal is simply ignored or a FIFO is full.

2.1 LV1 rejection

When a LV1 true signal is rejected by the MCC (see section 1.2), **one complete module event is lost** : in the final ATLAS event, there will be no data for this particular module. For a given module, the probability to lose complete module events is a fixed probability which does not depend on the size or the type of the event, since there is no correlation between the lost event and the events which are at the origin of the rejection.

2.2 FE-event rejection

The way the Receivers are made (see section 1.4) implies that only **complete individual FE-events are lost**. When a FIFO is full, all data words of the event from this FE (a FE-event) are lost : there cannot be a partial data loss for a FE-event. In the final ATLAS event, there will be no data from this particular FE, but the other FE data are not affected.

The probability for FE-events to be lost depends on their size. There is a systematic rejection of the FE-events with a size exceeding the *FIFO Depth* : all FE-events with more than *FIFO Depth*-1 hits are lost (the EoE word is also written in the FIFO). The probability to lose smaller FE-events is proportional to the FE-event size. Sample distributions for the FE-event size are presented in figure 2; the cutoff due to the size of the FIFO is clear (*FIFO Depth* was 32).
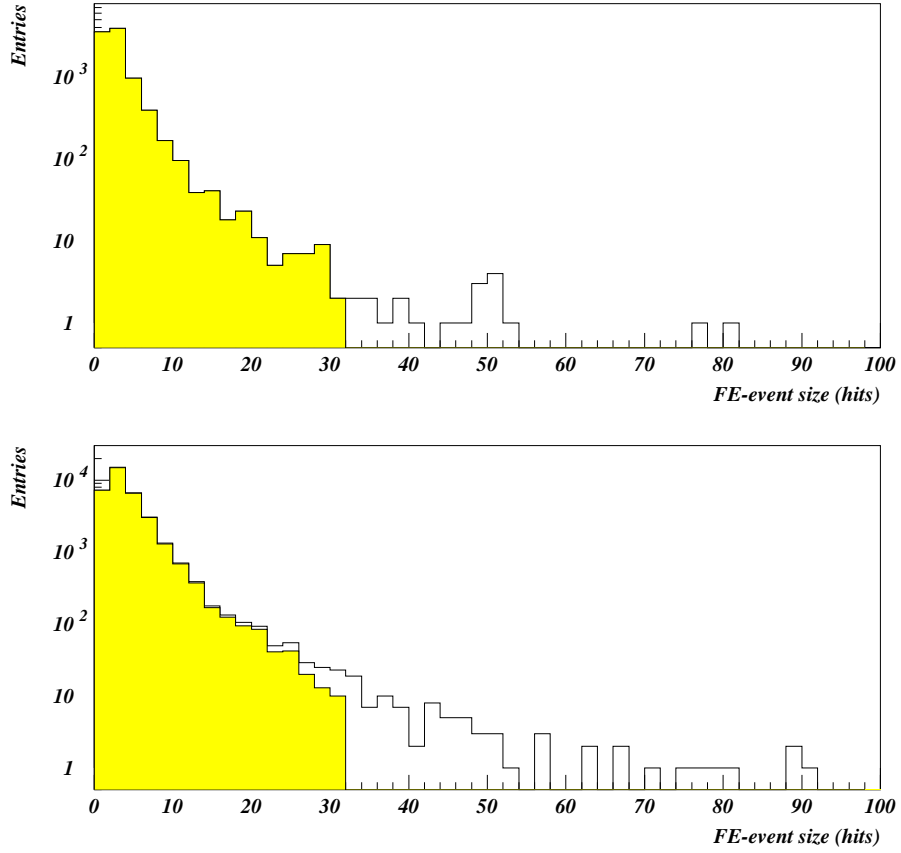
Figure 2: Examples of FE-event size distributions (in number of hits), for the central module of the layer 3 (upper plot) and of the B layer (lower plot). White events are rejected by the MCC.

3 MCC performance study

The MCCSim program has been used to study the MCC effect on the data flow in the pixel detector. The input data are 400 high luminosity $H \rightarrow b\bar{b}$ events, with a Higgs mass of 400 GeV (pile-up events are included in the data). The data were simulated with the Pixel TDR geometry [1] (files 28 to 31 of CERN tape Y20346). These events have been processed through a dedicated FE simulation in order to produce the MCCSim input data file. By the randomization of the modules around ϕ, the statistics was increased from 400 to 4000 events.

The nominal simulation has been performed with a *LV1 Target Rate* of 100 kHz, which gave an average LV1 rate between 75.7 kHz and 75.9 kHz. The MCC input data rate in standard conditions is shown on figure 3.

3.1 LV1 rejection origin

The LV1 rejection by the MCC can occur in two cases (see section 1.2). However, the rejection due to the pending event counter is negligible compared to the rejection due to the ERROR state. Actually, using this data set and under standard conditions, no trigger has been rejected because of the pending event counter. Of course, this is true when only one trigger signal is sent to the FE circuits for each LV1 true signal received by the MCC.

When the *Max Pending LV1* parameter is reduced to 8 instead of 16, a few triggers (less than 0.5%) are rejected because of the pending event counter : the main effect is still coming from the ERROR state rejection.

3.2 Optical link speed effect

Figure 4 shows the total MCC output data rate, with a speed of 40 Mbit/s for each optical link, the B layer having two such links per module. The B layer shows data rates relatively close to the maximum bandwidth of 80 Mbit/s, whereas in the layer 2, only a factor of 2 is available between the maximum and the used bandwidths. Safer working conditions can be obtained by increasing the optical links speed to 80 Mbit/s.

Moreover, the time spent by the MCC in ERROR state, when LV1 triggers are rejected, is directly related to the output data rate. Increasing the
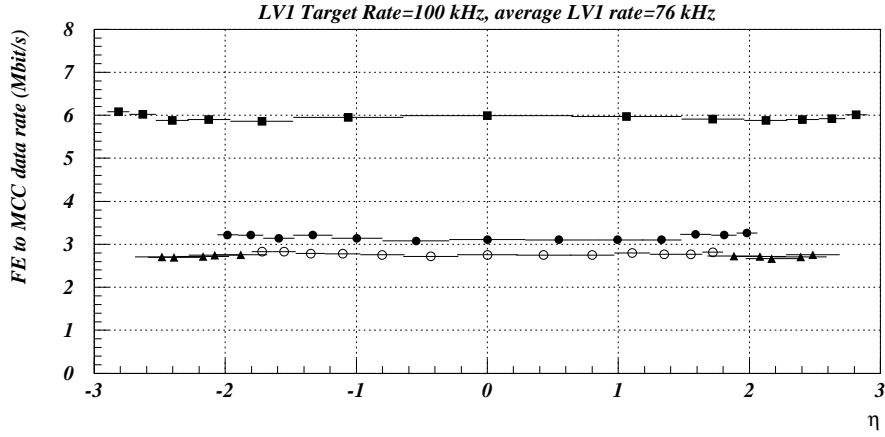
Figure 3: FE to MCC data rate map as a function of η. Each point corresponds to a module position in z. Triangles are for endcap disks, open circles for barrel layer 3, black circles for layer 2 and squares for the B layer.
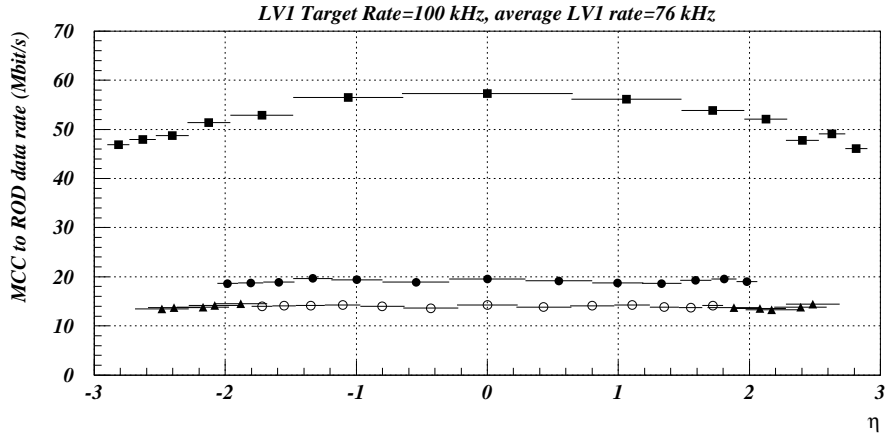

Figure 4: MCC to ROD data rate map. The *Optical Link Speed* was 40 Mbit/s/link, so that the final output speed for the B layer was 80 Mbit/s.

9

optical links speed by a factor of 2, divides also the time between the ERROR condition and the reset request by 2. Figure 5 shows the LV1 rejection and the average FIFO occupancy maps with the standard 40 Mbit/s link speed and also for 80 Mbit/s. From this figure, it is obvious that the LV1 rejection could be divided by more than 2 by using 80 Mbit/s optical links, which would improve significantly the MCC performance. Finally, the average FIFO occupancy being also directly related to the output data rate, figure 5 shows that the occupancy is much improved by increasing the links speed, leading to less WARNING and ERROR conditions and therefore to a better MCC performance.

The following results have been produced with an optical link speed of 80 Mbit/s, which gives a bandwidth of 160 Mbit/s for the B layer.

3.3   Comparison of LV1 and FE-event rejection

The right upper plot from figure 5 shows the fraction of module events which are lost due to the LV1 rejection mechanism. The middle plot from figure 6 shows the fraction of module events which are lost due to the FE-event rejection mechanism. It can be noticed that the probability to lose a full module event due to a LV1 rejection is one order of magnitude higher than due to a multiple FE-event rejection.

The most likely case is to lose some FE-events and have a partial loss in a module event. The upper plot of figure 6 shows the fraction of events (for accepted LV1s) which are completely readout. The few percent of events (between 4 and 17.5% in the B layer) which are partially readout is mostly due to big FE-events, which are systematically rejected, since the FE-event efficiency for events smaller than the *FIFO Depth* is close to 1 (more than 97% in the B layer, see lower plot of figure 6).

3.4   FIFO depth effect

It seems obvious that increasing the size of the Receiver FIFOs can only increase the global MCC efficiency. The first effect is a direct effect on the maximum size of FE-events allowed by the MCC (see section 2.2). Moreover, the FIFOs are also less often full and, consequently, the MCC is less often in ERROR state. But, when the FIFOs are very large and the MCC is in ERROR state, it takes more time to empty the FIFOs, and the MCC spends
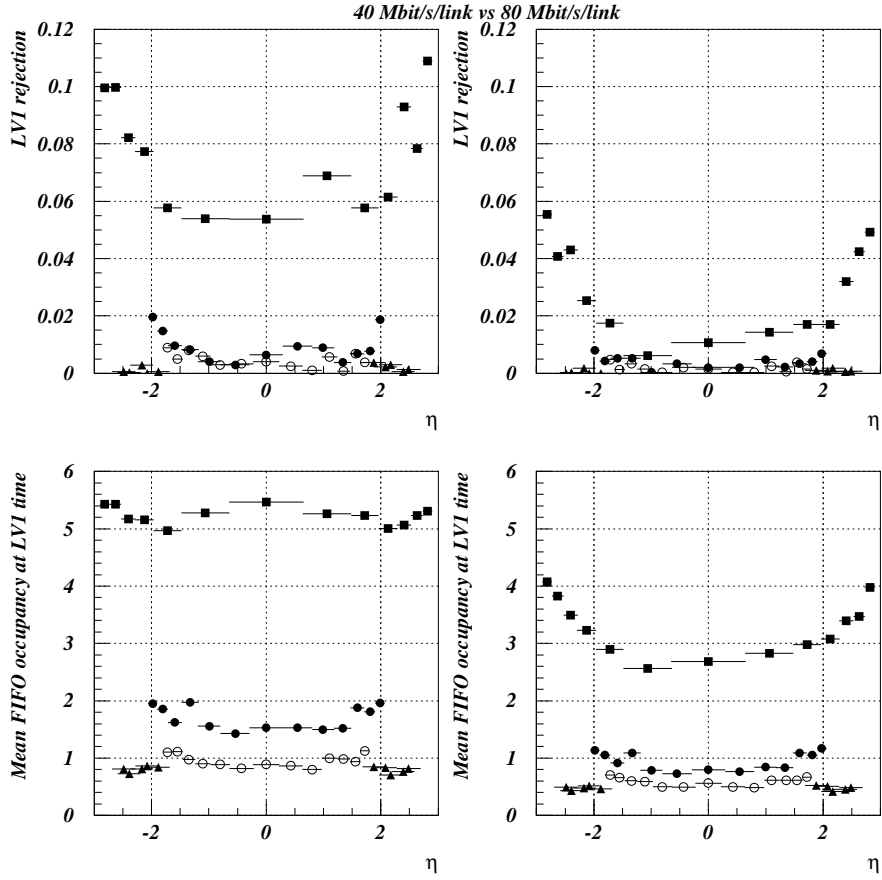
Figure 5: LV1 rejection and average FIFO occupancy maps. Left hand side plots are for 40 Mbit/s/optical link, right hand side plots are for 80 Mbit/s/link.
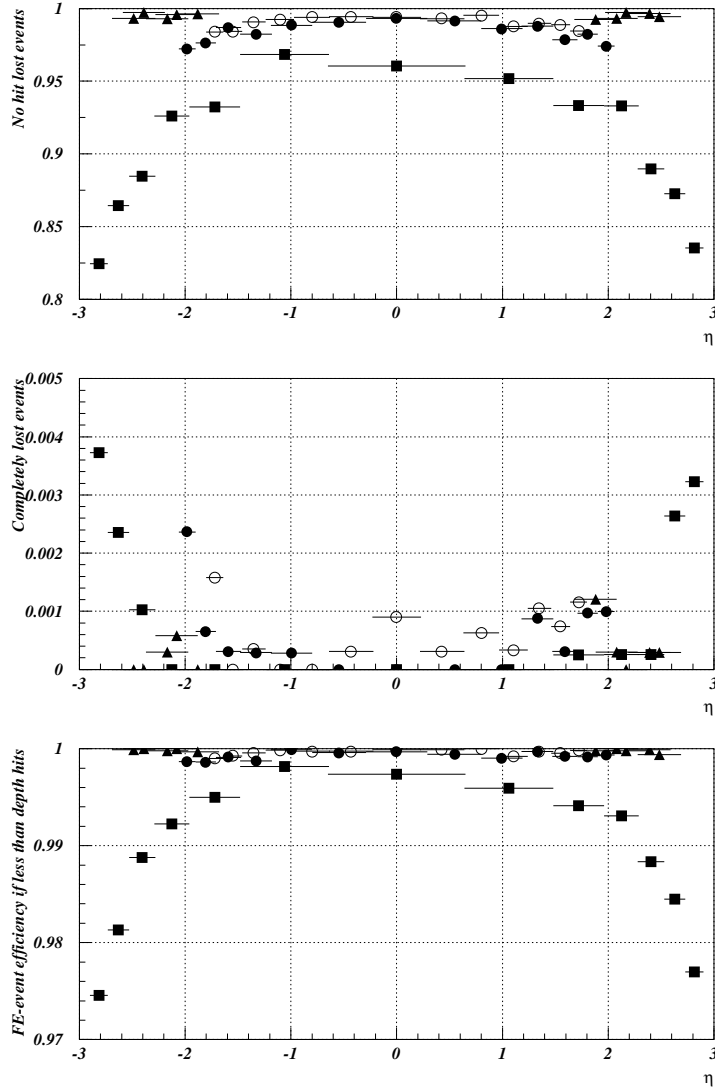
Figure 6: Upper plot : map of the fraction of module events with no lost hits. Middle plot : map of the fraction of module events completely lost. Lower plot : map of the FE-event efficiency for events smaller than the *FIFO Depth*. LV1 rejection is not included in these plots.

12

more time before requesting a reset to the ROD. The overall effect can be seen on figure 7, were the global MCC hit efficiency (taking into account the LV1 efficiency) is plotted as a function of the FIFOs size, for the B layer, where the most important effect is observed. The absolute value of raw MCC hit efficiency should be considered with caution : large FE-events are systematically lost and, in the same time, their size implies that they are considered with a high weight in the calculation of the hit efficiency.

A problem occurs related to the way the FIFO size could be increased. Obviously, bigger FIFOs would lead to bigger MCC circuits which would have a lot of disadvantages. Reducing the number of bits per word in the FIFOs is also a way of increasing the FIFO size and keeping the MCC area almost constant. It could be achieved by reducing the number of ToT bits and by removing the LV1 number which is presently stored in the FIFO.

3.5   ToT length effect

The length of the ToT information has a direct effect on input and output data rates. Therefore, decreasing the number of bits of the ToT information increases the MCC efficiency. Figure 8 shows the average FIFO occupancy as a function of the *ToT Length*, for the B layer, where the most important effect is observed.

The ToT information could be reduced because not all the 8 bits will be used during ATLAS data taking : the LV1 trigger latency viewed by the FE circuits will be around 110 BCOs. Therefore, any hit with a ToT bigger than 110 will not be readout, due to the FE circuit readout architecture. This means that 7 bits will be the maximum ToT length which will be available during ATLAS data taking.

Reducing the number of bits for the ToT information has also an indirect effect : the ToT information is stored inside the MCC in a FIFO whose total area is related to the number of bits it can store. Therefore, increasing the FIFO size from 32 to 34 words would lead to a total area approximately equivalent to the area of a FIFO of 36 words, but with a ToT information of 6 bits instead of 8. The decreasing of the *ToT Length* to 6 bits and the increasing of the *FIFO Depth* to 36 leads to a noticeable improvement (see figure 9), while the MCC circuit size is only slightly increased (around 10%).
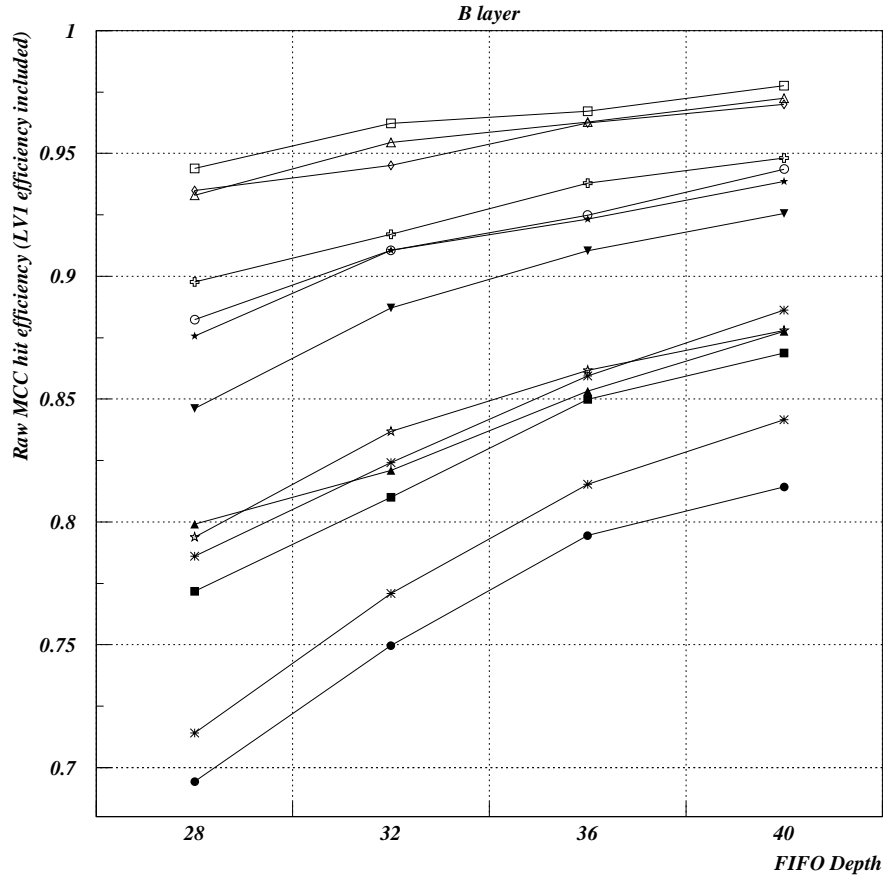
13

Figure 7: Raw MCC hit efficiency (with LV1 rejection effect included) as a function of the *FIFO Depth* for the 13 modules of the B layer. The present MCC circuits have a FIFO depth of 32.
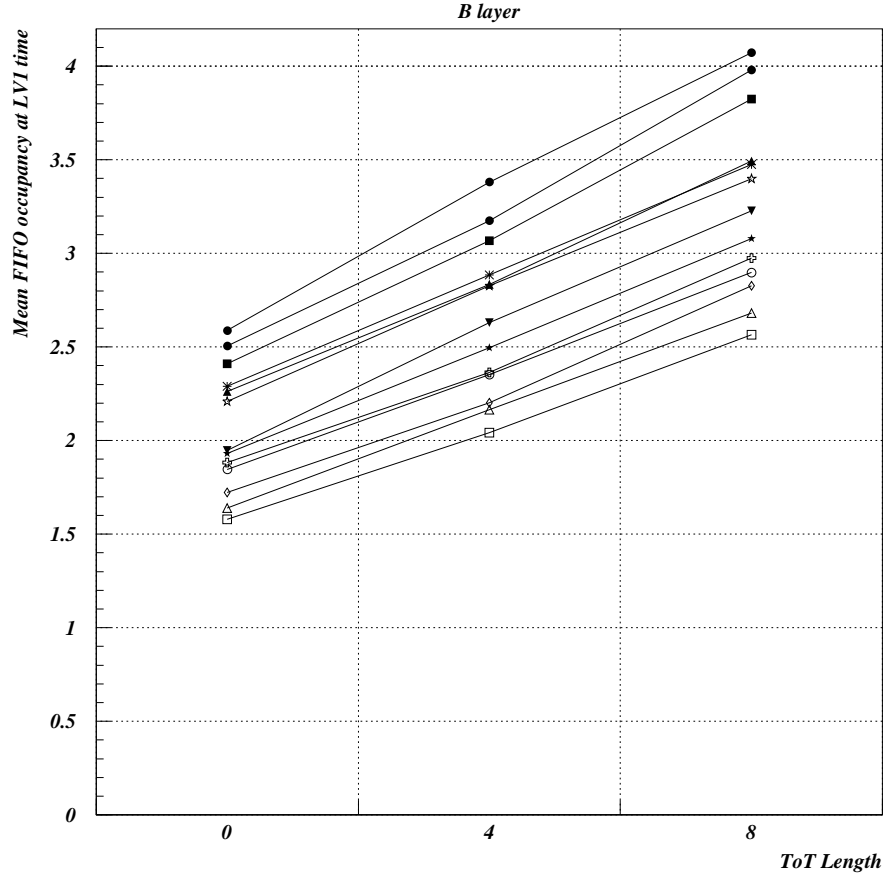
14

Figure 8: Average FIFO occupancy as a function of the *ToT Length* for the 13 modules of the B layer. The present FE and MCC circuits have a ToT of 8 bits.
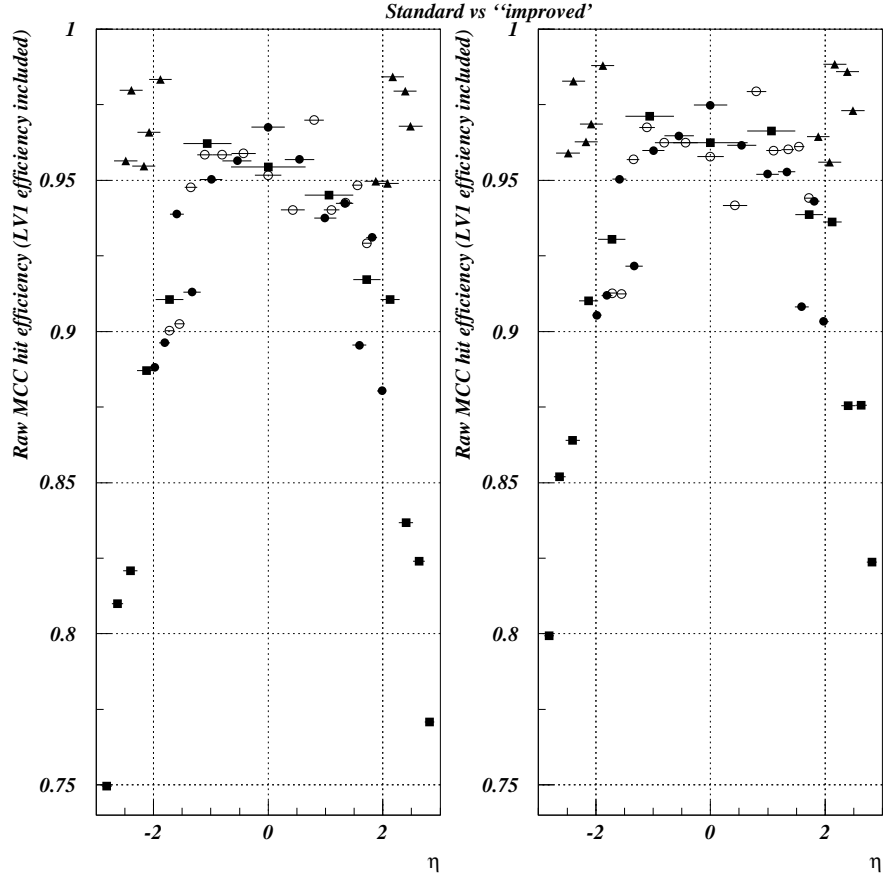
Figure 9: Raw MCC hit efficiency (with LV1 rejection effect included) maps. Left hand side is the standard configuration (*ToT Length* of 8 and *FIFO Depth* of 32), right hand side is the improved configuration (*ToT Length* of 6 and *FIFO Depth* of 36).

3.6 ROD behaviour effect

It is very difficult presently to know what will be the real ROD behaviour in case of ERROR condition. This behaviour has a strong effect on the time spent by the MCC in ERROR state, and therefore on the number of rejected triggers. As stated in section 1.6, a possible mechanism for the reset sequence has been implemented in MCCSim, but it is one of the most optimistic schemes. Moreover, for the previous results, the *ROD Reset Latency* was 70 BCOs, which is very close to the minimum latency which could be achievable by the ROD. However, it would be possible to increase this number to values which remain small compared to the time lapses between the declaration of the ERROR condition and the sending of the reset request to the ROD (Stopping time). This time is only slightly higher in the B layer than in the other layers and in the disks because the pixel occupancy increase is compensated by the higher output rate. The average Stopping time ranges rathly between 200 and 800 BCOs.

Figure 10 shows the deterioration of the LV1 rejection when increasing the *ROD Reset Latency* parameter to maybe more realistic values. This is shown for the B layer, which is the most sensitive to this effect. Obviously, the final ROD mechanism will have to be very efficient if we want to have useful pixel data in ATLAS. In particular, waiting for the next ATLAS reset, which could occur only every 1 to 10 ms, is impossible.

3.7 LV1 rate upgrade effect

All the previous results have been obtained with an average LV1 rate of 76 kHz, which is the higher rate allowed by the level-1 trigger logic. However, this upper limit could be upgraded to 100 kHz [3], which corresponds to a *LV1 Target Rate* around 135 kHz. This upgrade will not concern the MCC, therefore the MCC must be already designed for this absolute maximum rate of 100 kHz. Figure 11 shows the effect on the MCC performances of the LV1 rate upgrade. The upper plot must be compared to figure 4 and the lower plots to figure 5. The degradation of the MCC performances is significant. Better performances could be obtained by raising the speed of the optical links to 160 Mbit/s (figure 12 to be compared to figures 11 and 5). A further improvement of the MCC performances could be obtained by using the configuration described in section 3.5 (see figure 13).
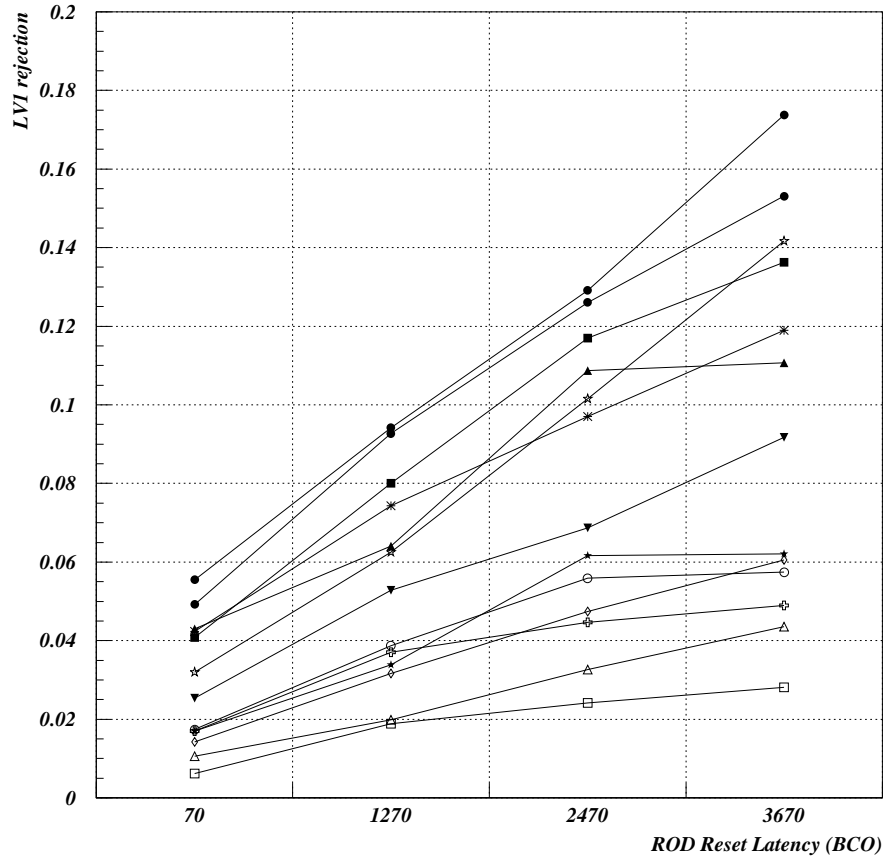
Figure 10: LV1 rejection as a function of the *ROD Reset Latency*, expressed in number of BCOs, for the 13 modules of the B layer.
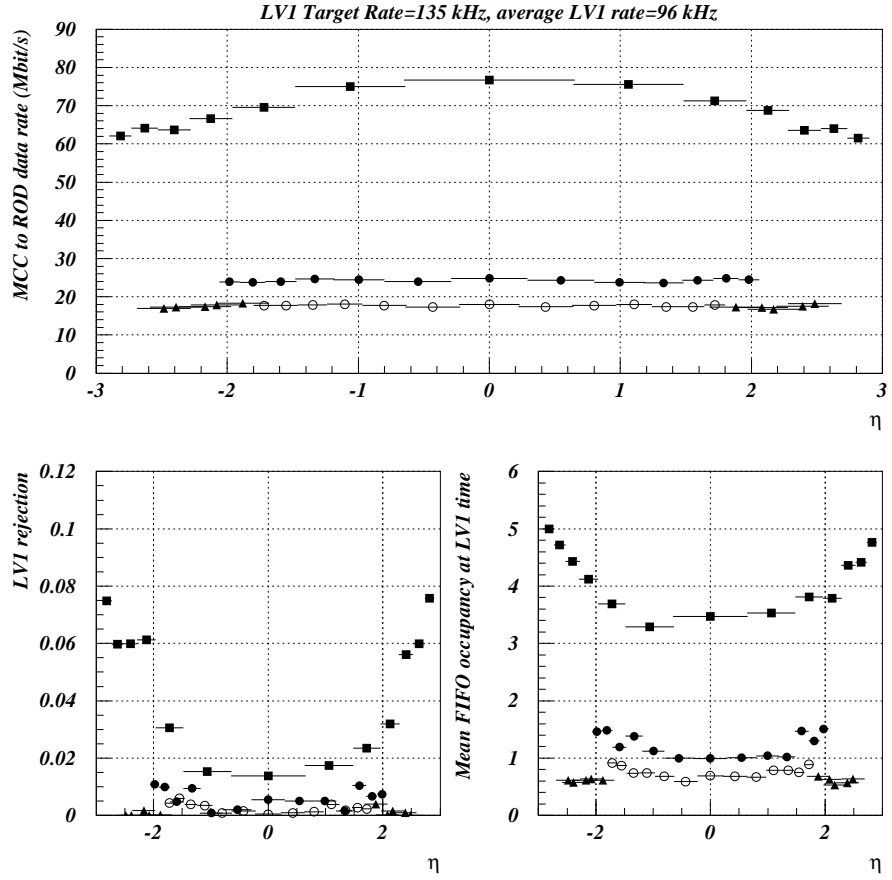
Figure 11: MCC to ROD data rate, LV1 rejection and average FIFO occupancy maps. For these plots, the average LV1 rate was set to the absolute upper limit (96 kHz).
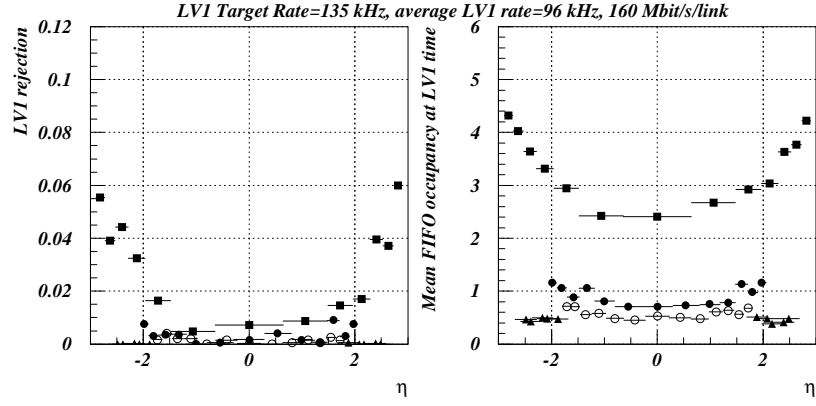
Figure 12: LV1 rejection and average FIFO occupancy maps. For these plots, the average LV1 rate was set to the absolute upper limit (96 kHz) and the *Optical Link Speed* to 160 Mbit/s/link.
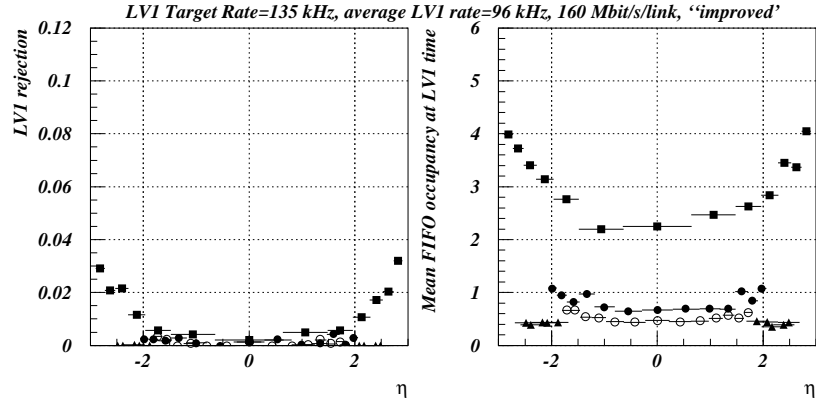

Figure 13: LV1 rejection and average FIFO occupancy maps. For these plots, the average LV1 rate was set to the absolute upper limit (96 kHz), the *Optical Link Speed* to 160 Mbit/s/link, the *ToT Length* to 6 bits and the *FIFO Depth* to 36.

4 Efficiency parameterization

A full simulation of the MCC, like in the MCCSim software, is not possible in the framework of ATLSIM for technical reasons (mainly the CPU time consumption). Therefore, the MCC efficiencies have been parameterized and included in this way in the PIXBDIG and PIXEDIG modules of ATLSIM, which perform the digitization of the pixel detector hits modules.

As stated in section 2, the data can be lost by the MCC because of two different mechanisms : LV1 true signal rejection and FIFO overflow conditions.

4.1 LV1 rejection

The LV1 rejection parameterization is very simple : independently for each pixel module, a fixed probability is used to reject events, using a random number generator. The probabilities are extracted, by example, from the plots on figure 5.

4.2 FE-event rejection

For an event which survives the LV1 rejection, the hits are treated as FE-events and no longer as a module-event :

- FE-events with a number of hits greater than or equal to the *FIFO Depth* are systematically rejected.

- FE-events with a number of hits lower than the *FIFO Depth* are rejected with a probability calculated with the formula :

$$P = e^{\frac{hits - d_0}{d_1}}$$

  where *hits* is the FE-event size. The two parameters d_0 and d_1 are extracted from a fit of the FE-event efficiency distribution. Fits are shown on figure 14.

4.3 ATLSIM implementation

A new bit has been added to the digits NOISE word (bit 3). A value of 1 for this bit signals that the corresponding digit has been lost by the MCC.
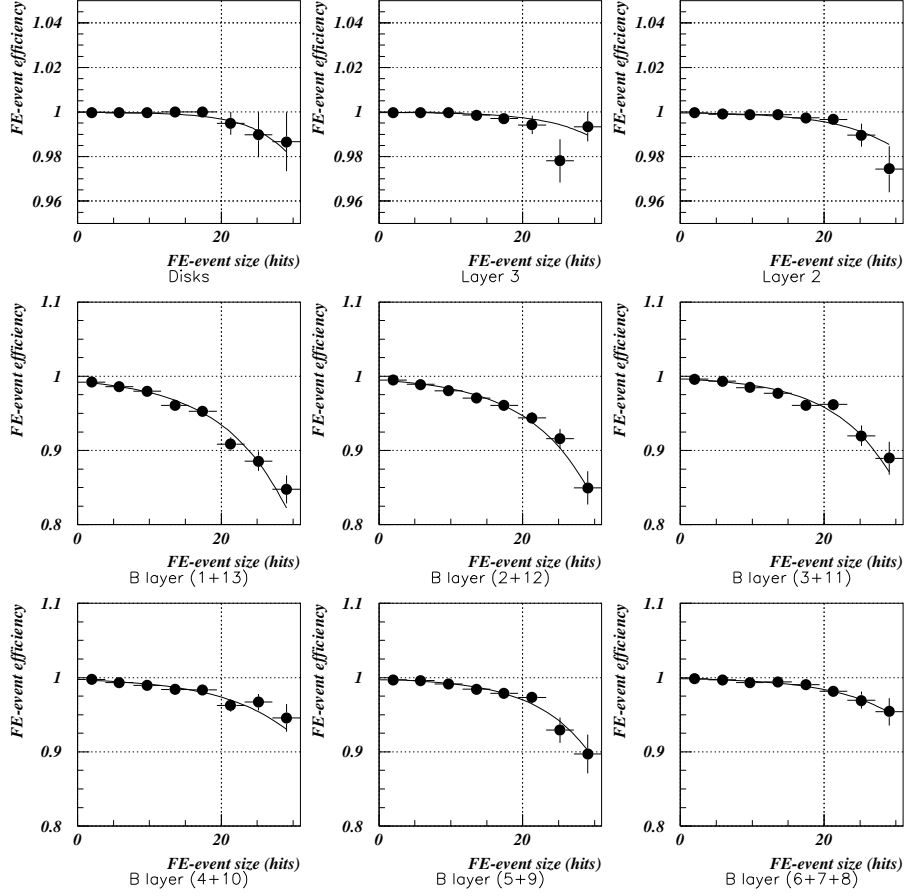
Figure 14: FE-event efficiency as a function of the FE-event size for the disks, layer 3, layer 2, and different η positions in the B layer. The superimposed curves are the parameterizations used in ATLSIM.

22

This bit is meaningful only for digits above the threshold or unmasked. For digits below the threshold or masked, the bit is 0.

The parameters used for the MCC efficiency parameterization (FIFO size, LV1 rejection, d_0 and d_1) are stored in the PIXB/PBMC and PIXE/PEMC banks. The configuration used by PIXBDIG or PIXEDIG can be selected by changing the value of the MCCVers parameter in the PIXB/PBEL or PIXE/PEEL banks.

Conclusion

The MCCSim program described in this note has been used to study the performances of the pixel detector Module Control Circuit. From these results, two conclusions can be extracted :

- the optical links should be used at 80 Mbit/s (at least), with two links for the B layer modules.
- some mechanism should be found in order to reduce the impact of the ERROR conditions. A possible new algorithm is the following :
  - when an ERROR condition occurs in a Receiver, a counter (Lost Events Counter) is started,
  - each time a new EoE word is received by this Receiver, the Lost Events Counter is incremented,
  - as soon as a free position is available in the FIFO, a special warning word is written (it could be WNG#1) and the Lost Events Counter is decremented,
  - there is not anymore an ERROR state for the MCC.

A parameterization of the MCC performances has been added to the PIXBDIG and PIXEDIG modules of ATLSIM. Therefore, physics analyses can now include the effect of this circuit, using the present architecture.

Acknowledgements

References

[1] ATLAS Pixel Detector Technical Design Report, ATLAS Pixel Detector Community, CERN/LHCC 98-13.

[2] LHC Project Web Site, CERN.

[3] ATLAS Level-1 Trigger Technical Design Report, ATLAS Level-1 Trigger Group, ATLAS TDR-12.