Available on CMS information server **CMS IN 1997/032**

CMS Internal Note

The content of this note is intended for CMS internal use and distribution only

13 November 1997

Data Acquisition for heavy ion physics

G. Wrochna

CERN, Geneva, Switzerland

Abstract

CMS Data Acquisition System was designed for p-p collisions. Conditions in the case of heavy ion collisions are very different. The aim of this paper is to check whether the DAQ can work effectively in such conditions. The data flow through the entire DAQ system is examined. Data volumes of various subdetectors are calculated. It is shown that the dominant component comes from the Pixel detector. In total one can collect 40-60 events/s, with is adequate for heavy ion physics.

1 DAQ parameters

CMS Data Acquisition System (DAQ) was designed for p-p collisions at the highest designed luminosity. In this case on can expect the following conditions:

- bunch spacing = 25 ns
- luminosity $\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$
- number of p-p interactions per bunch crossing ≈ 20
- number of charged particles per η-unit ≈ 5
- First Level Trigger rate < 100 kHz

This has led to the design of DAQ with the following parameters:

- Readout Dual Port Memory (RDPM) input < 200 M Bytes/s
- number of RDPM's = 1000
- nominal Switch bandwidth = 500 G bits/s
- mass storage capacity = 100 M Bytes/s

Conditions in heavy ion collisions are very different. Let us consider an extreme case of Pb-Pb collisions:

- bunch spacing = 125 ns
- luminosity $\mathcal{L} = 10^{27}\mathrm{cm}^{-2}\mathrm{s}^{-1}$
- average interaction rate ≈ 7.6 kHz
- number of charged particles per η-unit ≈ 2500 (min. bias) — ≈ 8000 (central)
- First Level Trigger rate ≈ 1 kHz [1]

Thus one expects much bigger events but with relatively low rate. Is the CMS DAQ system suitable for such conditions? Where are possible bottle-necks? Those are the questions we are going to address in this paper.

2 Requirements for dimuon physics

Probably the most demanding physics are dimuon channels used to study formation of bound states of heavy quarks in dense matter [2]. Their require data from the muon detector to recognise muons, from the central tracker to measure precisely their momenta, and from calorimeters to estimate centrality of collisions. Recent study [3] has shown that occupancies in Pixel Detector and in 4 outer MSGC layers are low enough to perform an effective pattern recognition up to $|\eta| = 0.8$. They are quoted in the table below. They have been obtained assuming 8000 charged particles per η-unit, which is rather an upper limit.

detector	occupancy
Pixel layer 1	4 %
Pixel layer 2	2 %
MSGC layer 4	18 %
MSGC layer 5	14 %
MSGC layer 6	11 %
MSGC layer 7	8 %

An independent study [4] suggests occupancies twice lower, but for the purpose of this paper we stay with more conservative estimate.

The occupancies should not change dramatically till $|\eta| \approx 1.5$, but going beyond this region would require careful, dedicated study. Therefore for the purpose of this paper we restrict ourselves to the region of $|\eta| < 1.5$. This restriction, however, is not valid for calorimeters. Full η coverage ($|\eta| < 5.5$, including Very Forward Calorimeter) is needed to estimate centrality of the collisions. Finally, we would like to read out the following detectors:

- Pixel barrel detector
- 4 outer MSGC layers + part of forward MSGC disks up to $|\eta| = 1.5$
- all calorimeters (ECAL + HCAL + VFCAL)
- muon system (RPC + Drift Tubes + CSC) up to $|\eta| = 1.5$

3 Scope of the study

In the consecutive sections we are going to examine the CMS DAQ according to the following plan:

- calculate the data volume for each subdetector
- calculate the data flow
  - from detector Front Ends to Front End Drivers (FED)
  - from FED's to RDPM's
  - through the Switch

The aim of the exercise is not to give precise numbers. It is just a very first attempt to a rough estimation. All the numbers given below should be taken with care, because they are subject to change due to many reasons.

- The CMS detector is still under optimisation and some moderate changes in the detector layout and thus in the number of channels are possible.
- The development of the DAQ system is in the design phase. Parameters assumed here are resulting from an extrapolation of technological trends. They may change significantly depending e.g. on technology which will be finally chosen.
- There is a lot of flexibility build into the system. The system can be configured in many ways and it can work in many different modes. Concrete solutions will be adopted to current running conditions and physics needs. Therefore they cannot be determined precisely today.

4 Data volumes

Pixel barrel

1 module = 2 rows $\times$8 chips $\times$(64$\times$64) pixels = 65 k pixels
Layer 1
(32$\times$10) modules $\times$65 k pixels = 21 M channels
$\times$4% occupancy = 840 k hits
Layer 2
(48$\times$10) modules $\times$65 k pixels = 32 M channels
$\times$2% occupancy = 630 k hits

Full readout
Analogue information about the signal on 1 pixel is equivalent to 1 Byte.
53 M channels $\times$1 Byte = 53 M Bytes
Zero suppressed readout
Let us assume that the readout is arranged in blocks corresponding to 1 module. Thus 2 Bytes are needed for an address within a module.
1370 k hits $\times$(1+2) Bytes = **4.1 M Bytes**

The zero suppressed readout is clearly more economic than the full one. One can, however, further reduce the data volume making use of the fact that single particle usually creates a cluster of 2-4 hits. Thus one can apply some clusterisation algorithm. As an example we describe here one proposed by Olga Kodolova [5].

A single module (physical detector) is a matrix of 128 rows $\times$512 columns of pixels. Let us introduce the following notation.

M – module number (1-800), 2 Bytes
X_i — row number (1-128), 1 Byte
Y_{ij} – column number of the first pixel in the cluster (1-512), 2 Bytes
N_{ij} – cluster size, 1 Byte
$A(k)$ – amplitude of the k-th pixel in the cluster, 1 Byte
EOR – End Of Row marker

EOM – End Of Module marker

The data format for one module can look as follows:

$$
\begin{array}{llll}
M, & X_1, Y_{11}, N_{11}, A(1), \quad \ldots, A(N_{11}), & Y_{12}, N_{12}, A(1), \quad \ldots, A(N_{12}), & \ldots, EOR \\
 & & \vdots & \\
 & X_i, Y_{i1}, N_{i1}, A(1), \quad \ldots, A(N_{i1}), & Y_{i2}, N_{i2}, A(1), \quad \ldots, A(N_{i2}), & \ldots, EOM
\end{array}
$$

Average number of clusters (particles crossing a layer) is ~ 200. Hence, for coordinates we need roughly 800 modules $\times$(1+2+1 Bytes) $\times$200 clusters = 640 k Bytes

For amplitudes one needs
1370 k hits $\times$1 Byte = 1370 k Bytes

Thus, the total data volume is $\sim$**2 M Bytes**.

MSGC barrel — option with all modules 12.5 cm long

layer	modules	chips	channels	hits
4 (stereo)	57×9 = 513	×12 = 6156	×128 = 788 k	×18% = 142 k
5 (1-side)	62×9 = 558	×8 = 4464	×128 = 571 k	×14% = 80 k
6 (1-side)	67×9 = 603	×8 = 4824	×128 = 617 k	×11% = 68 k
7 (stereo)	72×9 = 603	×12 = 7774	×128 = 995 k	×8% = 80 k
TOTAL			3 M	370 k

Full readout
Because the dynamic range of the analog signal is not yet defined I took a conservative assumption of 8 bits.
3 M channels $\times$8 bits = 24 M bits = 3 M Bytes
Zero suppressed readout
370 k hits $\times$(10+8) bits = 6.7 M bits = 832 k Bytes

Muon RPC
$<$ 1 k Byte

Muon Drift Tubes
5 k Bytes

HCAL — 100% occupancy
Full readout
14616 channels $\times$10 Bytes = 146 k Bytes

ECAL — 100% occupancy

Full precise information readout
3888 towers $\times$36 crystals $\times$2 Bytes $\times$15 time slices = 4.2 M Bytes
This solution leads to very large data volume, but probably 15 time slices are not needed (there is no pileup).

Optional – full, 1 time slice readout
3888 towers $\times$36 crystals $\times$2 Bytes = 280 k Bytes
This is already feasible, but in the case of heavy ions one can reduce readout granularity to that of trigger.

Full trigger information readout
strips of 6 crystals, only one time slice
3888 towers $\times$6 strips $\times$2 Bytes = 47 k Bytes

The dynamic range of 2 Bytes, assuming the Least Significant Bit (LSB) = 50 MeV, extends up to maximal energy of 3.2 TeV. This is by far enough.

TOTAL

	k Bytes
Pixel barrel	~ 2000
MSGC barrel	832
ECAL full trigger	47
HCAL full	146
Muon RPC	1
Muon Drift Tubes	5

+ some outer modules of forward MSGC for $|\eta| < 1.5$
+ some outer Muon Cathode Strip Chambers for $|\eta| < 1.5$
+ VFCAL

$\approx$ **3 M Bytes / event**

i.e. 3 times more than for p-p collisions at $\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Assuming mass storage of 100 M Bytes/s we can write to tape 30 events / second.
Digital compressing can give factor ~ 2, so we hope for **60 events / second**.

5 Data flow

PIXEL barrel — from Front End to FED

1 channel = analogue information equivalent to 8 bits
1 link = 4 chips $\times$(64$\times$64) pixels = 16 k channels $\Rightarrow$ 14 bits are needed for the address

Full readout
16 k channels $\times$8 bits $\times$1000 Hz = 131 M bits/s

Zero suppressed readout
16 k channels $\times$4% occupancy $\times$(8+14) bits $\times$1000 Hz = **15 M bits/s**

Here we see a possible bottle-neck. The connection is designed for **3-4 M bits/s** expected for p-p collisions at $\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The full readout option is far beyond this and therefore it is not feasible. The zero suppressed readout create 3-4 times higher data flow than that of p-p collisions. Therefore one should carefully watch the design at this point.

PIXEL barrel — from FED to RDPM

1 module = 2 rows $\times$8 chips $\times$(64$\times$64) pixels = 65 536 channels

Layer 1: 16 modules $\rightarrow$ 1 FED $\rightarrow$ 1 RDPM
16 $\times$65 536 = 1 M channels
Full readout
1 M channels $\times$1 Byte $\times$1000 Hz = 1 G Bytes/s
Zero suppressed readout
1 M channels $\times$4% occupancy $\times$(1+2) Bytes $\times$1000 Hz = **126 M Bytes/s**

Layer 2: 32 modules $\rightarrow$ 1 FED $\rightarrow$ 1 RDPM
32 $\times$65 536 = 2 M channels
Full readout
2 M channels $\times$1 Byte $\times$1000 Hz = 2 G Bytes/s
Zero suppressed readout
2 M channels $\times$2% occupancy $\times$(1+2) Bytes $\times$1000 Hz = **126 M Bytes/s**

Here also the full readout option is excluded, whereas the zero suppressed readout is feasible.

MSGC barrel

64 links $\times$256 channels = 16 k channels $\rightarrow$ 1 FED
8 FED's $\rightarrow$ 1 RDPM

Full readout
8 $\times$16 k channels $\times$8 bits = 128 k Bytes
$\times$1000 Hz = 128 M Bytes/s

Zero suppressed readout
8 $\times$16 k channels $\times$18% occupancy $\times$(10+8) bits = 52 k Bytes
$\times$1000 Hz = 52 M Bytes/s

The conditions are reasonable for both options.

ECAL

64 towers $\rightarrow$ 1 RDPM

Full precise information readout
64 towers $\times$36 crystals $\times$2 Bytes $\times$15 time slices = 69 k Bytes
$\times$1000 Hz = 69 M Bytes/s

Optional – full, 1 time slice readout
64 towers $\times$36 crystals $\times$2 Bytes = 4.6 k Bytes
$\times$1000 Hz = 4.6 M Bytes/s

Full trigger information readout
strips of 6 crystals, only one time slice
64 towers $\times$6 strips $\times$2 Bytes = 768 Bytes
$\times$1000 Hz = 0.8 M Bytes/s

Thus all the options are feasible.

HCAL, VFCAL, MUONS
very low bandwidth required

SWITCH

event size = 3 M Bytes
trigger rate = 1000 Hz
$\Rightarrow$ required bandwidth = **3 G Bytes/s**

nominal bandwidth = 500 G bits/s = 62 G Bytes/s
In practice only 50% of the nominal bandwidth can be used due to traffic problems.
Hence the effective bandwidth is $\sim$ **30 G Bytes/s**. This ensures a large safety margin (factor 10).

6 Conclusions

There is no problem with data volume and flow with all detectors, but Pixels. The link from the Front End to the FED is a possible bottle-neck. The Front End $\rightarrow$ FED rate (15 M bits/s) is significantly higher then that for p-p (3-4 M bits/s).

Significant reduction of the data volume can be achieved by an on-line clusterisation of Pixel data. This can be performed by the *Processor Farm* which was originally designed as *Event Filter*. In this case the name *Data Filter* would be more appropriate. This is an impressive example of high flexibility of the CMS Trigger and DAQ System. In fact the whole study presented in this paper is a good illustration of this flexibility. The system is able to work in conditions which are very far from those for which it was designed.

The expected event size for Pb-Pb collisions is $\sim$**3 M Bytes** before a digital compression. With the First Level Trigger rate of 1 kHz and a mass storage of 100 M Bytes/s one can write to tape $\sim$**40-60 events / s**, depending on the compression factor. This seems to be adequate for heavy ion physics.

Acknowledgement

Many people contributed to this paper by providing useful information and by stimulating discussions. Among them are Marc Bedjidian, Horst Breuker, Sergio Cittolin, Daniel Denegri, Olga Kodolova, Danek Kotliński, Alexander Nikitenko, Attila Racz, and João Varela. I would like to express my gratefulness to all of them.

References

[1] **CMS Note 1997/089**, G Wrochna, *Muon trigger for heavy ion physics*.

[2] M. Bedjidian, *Status of Dimuon Reconstruction*, presented at *II CMS Heavy Ion Meeting*, 4-6 February 97, Dubna.

[3] O.L. Kodolova, *CMS Tracker Occupancy, Pattern Recognition and Vertex Finding in Heavy Ion Collisions*, presented at *II CMS Heavy Ion Meeting*, 4-6 February 97, Dubna.

[4] D. Kotliński, private communication.

[5] O.L. Kodolova, private communication.