Available on CMS information server **CMS IN 1999/007**

# CMS Internal Note

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

**19 February 1999**

# Data Acquisition for heavy ion physics

A. Racz

*CERN, Geneva, Switzerland*

M. Bedjidian

*Institut de Physique Nucleaire de Lyon, IN2P3-CNRS, Univ. Lyon I - Villeurbanne, France*

O. Kodolova

*Moscow State University - Moscow, Russia*

D. Kotliński

*Paul Scherrer Institut - Villigen, Switzerland*

G. Wrochna[1)]

*Soltan Institute for Nuclear Studies, Warsaw, Poland*

**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 was found that in total one can collect ∼70 events/s, which is adequate for heavy ion physics.

This paper is to replace CMS IN 1997/032 which is now obsolete.

[1)] Partially supported by the Polish Committee for Scientific Research under grant KBN 115/E-343/SPUB/P03/057/98.

# 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 $\approx 20$
- number of charged particles per $\eta$-unit $\approx 5$
- 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 (effective rate)
- number of RDPM's = 512
- 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 $\approx 7.6$ kHz
- number of charged particles per $\eta$-unit $\approx 2500$ (min. bias) — $\approx 8000$ (central)
- First Level Trigger rate $\approx 1$ kHz [1]

One expects much higher occupancy but with relatively low event 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]. They 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 has shown that occupancies in Pixel Detector [3] and in 4 outer MSGC layers [4] are low enough to perform an effective pattern recognition up to $|\eta| = 1.5$. They are quoted in the table below for the Phase II, high luminosity Tracker. They have been obtained assuming 8000 charged particles per $\eta$-unit, which is rather an upper limit. We assume to read only one side of double-sided MSGC, because at so high occupancies the stereo information is unlikely to be usefull.

| detector | width | length | radius | occupancy |
|---|---|---|---|---|
| Pixel layer 1 | 125 $\mu$m | 150 $\mu$m | 7 cm | 0.53 % |
| Pixel layer 2 | 125 $\mu$m | 150 $\mu$m | 11 cm | 0.28 % |
| MSGC layer 4 | 200 $\mu$m | 125 mm | 90 cm | 10.70 % |
| MSGC layer 5 | 200 $\mu$m | 125 mm | 98 cm | 8.60 % |
| MSGC layer 6 | 200 $\mu$m | 250 mm | 106 cm | 12.00 % |
| MSGC layer 7 | 200 $\mu$m | 250 mm | 114 cm | 9.50 % |
| MSGC disk 1 | 200 $\mu$m | 125 mm | 75 cm | 15.60 % |
| MSGC disk 2 | 200 $\mu$m | 125 mm | 83 cm | 14.60 % |
| MSGC disk 3 | 200 $\mu$m | 125 mm | 95 cm | 14.60 % |
| MSGC disk 4 | 200 $\mu$m | 125 mm | 107 cm | 9.70 % |

The occupancies are calculated for $|\eta| < 1.5$. Going beyond this region would require careful, dedicated study. Therefore for the purpose of this paper we restrict ourselves to the region of $|\eta| < 1.5$. Thus, the forward Pixel disks are not required. This restriction, however, is not valid for calorimeters. Full $\eta$ coverage ($|\eta| < 5$, including Forward Calorimeter HF) 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)
- muon system (RPC + Drift Tubes + CSC)

This configuration is a baseline, which is rather conservative. If the occupancies will be significantly lower one can think of using other detectors in the following order of priority.

1. two inner MSGC barrel layers
2. barrel Silicon Detector
3. remaining MSGC endcap disks
4. forward Silicon Detectors
5. stereo layers

Similar approach can be taken for lighter ions.

# 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 for DAQ

## *Pixel barrel*

1 module = 2 rows $\times$8 chips $\times$(53$\times$52) pixels = 44 k pixels
**Layer 1**
(30$\times$8) modules $\times$44 k pixels = 10.6 M channels
$\times$0.53% occupancy = 56 k hits
**Layer 2**
(46$\times$8) modules $\times$44 k pixels = 16.2 M channels
$\times$0.28% occupancy = 45 k hits

The full readout option, without zero suppression, is not provided by hardware.

**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. Analogue information about the signal on 1 pixel is equivalent to 4-6 bits. Taking more conservative value of 6 bits for the analog pixel information and taking into account some overhead due to module headers one can assume 3 Bytes per pixel in total.
(56 + 45) k hits $\times$3 Bytes = **300 k Bytes**

The zero suppressed readout is clearly more economic than the full one. One can, however, try to reduce further 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. Let us consider a simple example.

A single module (physical detector) is a matrix of 106 rows $\times$416 columns of pixels. Let us introduce the following notation.

$M$ – module number (1-608), 10 bits
$X_i$ – column number (1-416), 9 bits
$Y_{ij}$ – row number of the first pixel in the cluster (1-106), 7 bits
$N_{ij}$ – cluster size, 3 bits
$A(k)$ – amplitude of the k-th pixel in the cluster, 6 bits
$EOC$ – End Of Cluster marker (7 bits, to be distinguished from the next $Y_{ij}$)
$EOM$ – End Of Module marker (9 bits, to be distinguished from the next $X_i$)

The data format for one module can look as follows:

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

The clusters are characterised by the following average values.

| | layer 1 | layer 2 |
|---|---|---|
| clusters (particles) per module | 150 | 80 |
| pixels per module | 470 | 250 |
| pixels per cluster | 3.1 | 3.1 |
| columns per cluster | 2.1 | 2.1 |

Thus an average cluster will contain 3 pixels in 2 rows and 2 columns:

$$\begin{array}{llllllllll} X, & Y_1, & N_1, & A_1, & A_2, & Y_2, & N_2, & A_1, & EOC & \\ 9+ & 7+ & 3+ & 6+ & 6+ & 7+ & 3+ & 6+ & 7 & = 54 \text{ bits} \end{array}$$

Hence, in total we need roughly
Layer 1: (30$\times$8) modules $\times$54 bits $\times$150 clusters = 250 k Bytes
Layer 2: (46$\times$8) modules $\times$54 bits $\times$80 clusters = 200 k Bytes

The total data volume is $\sim$**450 k Bytes**, which is higher than in the case of single pixels. This means that the average cluster size is too small to compensate the overhead of the format.

## *MSGC*

In the endcap only 6 disks on each side are used.

| | modules | channels | hits |
|---|---|---|---|
| layer 4 | $59\times18 = 1062$ | $\times512 = 544$ k | $\times10.7\% = 58$ k |
| layer 5 | $64\times18 = 1152$ | $\times512 = 590$ k | $\times8.6\% = 51$ k |
| layer 6 | $69\times10 = 690$ | $\times512 = 353$ k | $\times12.0\% = 42$ k |
| layer 7 | $74\times10 = 740$ | $\times512 = 379$ k | $\times9.5\% = 36$ k |
| disk 1 | $48\times12 = 576$ | $\times512 = 295$ k | $\times15.6\% = 46$ k |
| disk 2 | $52\times12 = 624$ | $\times512 = 320$ k | $\times14.6\% = 47$ k |
| disk 3 | $60\times12 = 720$ | $\times512 = 369$ k | $\times14.6\% = 54$ k |
| disk 4 | $56\times12 = 672$ | $\times512 = 344$ k | $\times9.7\% = 33$ k |
| TOTAL | | 3.2 M | 367 k |

**Full readout**
Because the dynamic range of the analog signal is not yet defined we took a conservative assumption of 8 bits.
3.2 M channels $\times$8 bits = 25 M bits = 3.2 M Bytes
**Zero suppressed readout**
367 k hits $\times$(10+8) bits = 6.6 M bits = 826 k Bytes

## *ECAL* — 100% occupancy

About 3 Bytes per cristal are needed — 18 bits amplitude + 6 bits for the format
There are 18 $\times$4 trigger towers in $\phi$ and 2 $\times$24 trigger towers in $\eta$.

**Full precise information readout**
3456 towers $\times$25 crystals $\times$3 Bytes $\times$10 time slices = 2.6 M Bytes
This solution leads to very large data volume, but probably 10 time slices are not needed. Without the pileup 1 time slice gives a precision of 1-2 %, which is enough for heavy ions. The possibility of recording e.g. 3 slices for lighter ions (O-O, Ca-Ca) is being considered. Currently it is not forseen to implement a digital filter reducing the number of time slices at the Front End, but one can apply a filter at the Online Farm.

**Full granularity, 1 time slice readout**
3456 towers $\times$25 crystals $\times$3 Bytes = **260 k Bytes**
This is already feasible, but in the case of heavy ions one can consider reduction of readout granularity to that of the trigger.

**Full trigger information readout**
towers of 25 crystals, only one time slice
3456 towers $\times$3 Bytes = 10 k Bytes
This is very low value, but this granularity is probably not adequate for studying electrons. Therefore, in case of a possible future need we consider further the full granularity, 1 time slice option.
In any case the dynamic range of 18 bits, assuming the Least Significant Bit (LSB) = 20 MeV, extends up to maximal energy of 5 TeV. This is by far enough.

## *HCAL (HB+HE+HF)* — 100% occupancy

**Full readout**
It is assumed to readout the energy calculated online by local DSPs. 2 Bytes are needed for the amplitude.
14 k channels $\times$2 Bytes = 28 k Bytes

## *Muon RPC*
$<$ 1 k Byte

## *Muon Drift Tubes*
5 k Bytes

### *Muon Cathode Strip Chambers*

6 k Bytes per muon segment in one station are needed.
Conservative assumption of 2 muons crossing 4 endcap stations gives
8 $\times$6 k Bytes = **48 k Bytes**

### *TOTAL*

| | k Bytes |
|---|---|
| Pixel barrel | 300 |
| MSGC barrel | 826 |
| ECAL - 1 time slice | 260 |
| HCAL full | 28 |
| Muon RPC | 1 |
| Muon Drift Tubes | 5 |
| Muon CSC | 48 |
| in total | $\sim$1448 |

The total event size $\approx$ **1.5 M Bytes** is only 1.5 time higher 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 about **70 events / second**.

## 5 Data flow

### *PIXEL barrel — Front End*

Single Pixel Front End chip covers an array of 52 columns, 53 pixels each. If at least one hit occured in a column the time-stamp is recorded and the full information (amplitude and position) from the hit pixels is stored in a buffer. In the case of very high rate some data might be lost due to buffer overflow. Let us consider the most demanding case of layer 1 at $\eta = 0$.

| | |
|---|---|
| Pixels in 1 module | 470 |
| Pixels in one chip | 30 |
| Pixels in a hit column | 2.1 |
| Hit columns in one chip | 14 |
| Pixel hit rate | 60 Hz |
| Column rate | 3.5 kHz |
| Chip rate | 8 kHz |

The rates mentined above are based on physics simulation and do not include any noise contribution. The readout error rate depends on the buffer size.

| Buffer size | 8 | 16 | 24 | 32 |
|---|---|---|---|---|
| Error rate | $4.7 \cdot 10^{-5}$ | $1.0 \cdot 10^{-5}$ | $5.2 \cdot 10^{-6}$ | $1.8 \cdot 10^{-7}$ |

For the p-p collisons the pixel buffer size was chosen to be 24, this seems also sufficient for the Pb-Pb collisions.

Another possible mechanism of spoiling the data is an overlap of clusters created by different particles. Again the simulation shows the values which are tolerable.

| | |
|---|---|
| Pixel overlaps | 0.75 % |
| Cluster overlaps | 1.7 % |

### *PIXEL barrel — from Front End to FED*

For 1 pixel channel one needs to read the analog amplitude and the position information consisting of the chip, column and row addresses. This information is send from the Front End to the FED using an analog link, with the digitial information being octal coded as analog signals (see Tracker TDR). Overall about 8-10 analog samples per pixel are transfered.

**Layer 1:** 1 link = 4 chips $\times$(53$\times$52) pixels = 11 k channels
**Layer 2:** 1 link = 8 chips $\times$(53$\times$52) pixels = 22 k channels
14-15 bits are needed for the address — 2-3 for the chip number + 6 for the column number + 6 for the row number

**Full readout** is not possible.

**Zero suppressed readout**
**Layer 1:** 11 k channels $\times$0.53% occupancy $\times$(8-10) samples $\times$1000 Hz = **0.58 M samples/s**
**Layer 2:** 22 k channels $\times$0.28% occupancy $\times$(8-10) samples $\times$1000 Hz = **0.62 M samples/s**

This is lower than the rate expected for p-p collisions at $\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ which is 3.8 M samples/s and 4.3 M samples/s for the layer 1 and 2 respectively.

### *PIXEL barrel — from FED to RDPM*

1 module = 16 chips $\times$(53$\times$52) pixels = 44 096 channels

**Layer 1:** 3 $\times$16 modules $\rightarrow$ 3 FED's $\rightarrow$ 1 RDPM
3 $\times$16 $\times$44 k = 2.1 M channels
**Zero suppressed readout**
2.1 M channels $\times$0.53% occupancy $\times$3 Bytes $\times$1000 Hz = **34 M Bytes/s**

**Layer 2:** 3 $\times$32 modules $\rightarrow$ 3 FED's $\rightarrow$ 1 RDPM
3 $\times$32 $\times$44 k = 4.2 M channels
**Zero suppressed readout**
4.2 M channels $\times$0.28% occupancy $\times$3 Bytes $\times$1000 Hz = **36 M Bytes/s**

Thus the zero suppressed readout is feasible.

### *MSGC barrel*

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

**Full readout**
2 $\times$16 k channels $\times$8 bits = 32 k Bytes
$\times$1000 Hz = 32 M Bytes/s

**Zero suppressed readout**
2 $\times$16 k channels $\times$15.6% occupancy $\times$(10+8) bits = 11 k Bytes
$\times$1000 Hz = **11 M Bytes/s**

The conditions are comfortable for both options.

### *ECAL*
68 towers $\rightarrow$ 1 RDPM

**Full precise information readout**
68 towers $\times$25 crystals $\times$3 Bytes $\times$10 time slices = 51 k Bytes
$\times$1000 Hz = **51 M Bytes/s**

**Optional – full, 1 time slice readout** (currently not considered at the Front End level)
68 towers $\times$25 crystals $\times$3 Bytes = 5.1 k Bytes
$\times$1000 Hz = 5.1 M Bytes/s

**Full trigger information readout**
towers of 25 crystals, only one time slice
68 towers $\times$3 Bytes = 204 Bytes
$\times$1000 Hz = 0.2 M Byte/s

Thus all the options would be feasible.

### *HCAL, MUONS*
very low bandwidth required

### *SWITCH*
event size = 1.5 M Bytes
trigger rate = 1000 Hz
$\Rightarrow$ required bandwidth = **1.5 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 20).

## 6 Conclusions

There is no problem with the data volume and the data flow for any of the subdetectors. The system is able to work in conditions which are very different from those for which it was designed, which is a good demonstration of its flexibility.

The expected event size for Pb-Pb collisions is $\sim$**1.5 M Bytes** without any 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 $\sim$**70 events / s**. 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 Horst Breuker, Sergio Cittolin, Daniel Denegri, John Elias, Alexander Nikitenko, and João Varela. We would like to express our 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, D. Kotliński, http://www1.psi.ch/~kotlinski/cms_pix_hi.html.

[4] **CMS Note 1999/004**, O. L. Kodolova, M. Bedjidian, S. Petrouchanko, *Dimuon reconstruction in heavy ion collisions using a detailed description of CMS geometry*.