A.Caner, I.Tomalin
June 9, 2000

On Balancing Data Flow from the Silicon Tracker (Draft)

1 The Detector Layout

The latest design (13/5/00) of the silicon tracker is shown in Fig. 1. More detailed information can be found in Ref. [1]. The tracker is effectively divided into two halves separated by the plane $z = 0$, where each half will be read out separately. Each half consists of 10 barrel layers and 12 end-cap disks.

Let us consider one half of the detector. Then, each layer in the barrel consists of rows (called 'rods' in the outer barrel) of 6 single or double-sided modules running parallel the z-axis. The number of rows in ϕ increases from 28 in the innermost layer to 74 in the outermost layer. Each end-cap disk has been designed such that its modules can be divided up into eight petals in ϕ. It is also divided into several rings in r. The number of modules per petal in rings 1-7 is given by 3, 3, 5, 7, 5, 7 and 10 respectively.

2 The Readout Chain

- DAQ event builder switch: 512 input legs, each capable of reading data from an RDPM (i.e. memory buffer) at a rate of 200 Mbytes/sec.
- Each RDPM can read data from 1–4 FEDs, also at a total rate of 200 Mbytes/sec.
- Each FED reads data from $96 \times 2 = 192$ APVs.
- Each APV is connected to 128 silicon strips.
- APVs are mounted on single or double-sided modules, with 4–6 APVs per side.

The number of inputs and the performance of the event builder switch may yet change. Furthermore, it is possible that (at extra cost), concentrator boards may be added to the system to allow more FEDs to be connected to each RDPM. Finally, an alternative FED design with inputs for only 48 pairs of APVs is still being considered. If this design is approved, then the number of FEDs would double, so the concentrator boards would be essential.

A possible L2 trigger could reduce the data flow from the tracker into the DAQ switch by a factor of 10. If this is implemented, then the trigger signal would be used

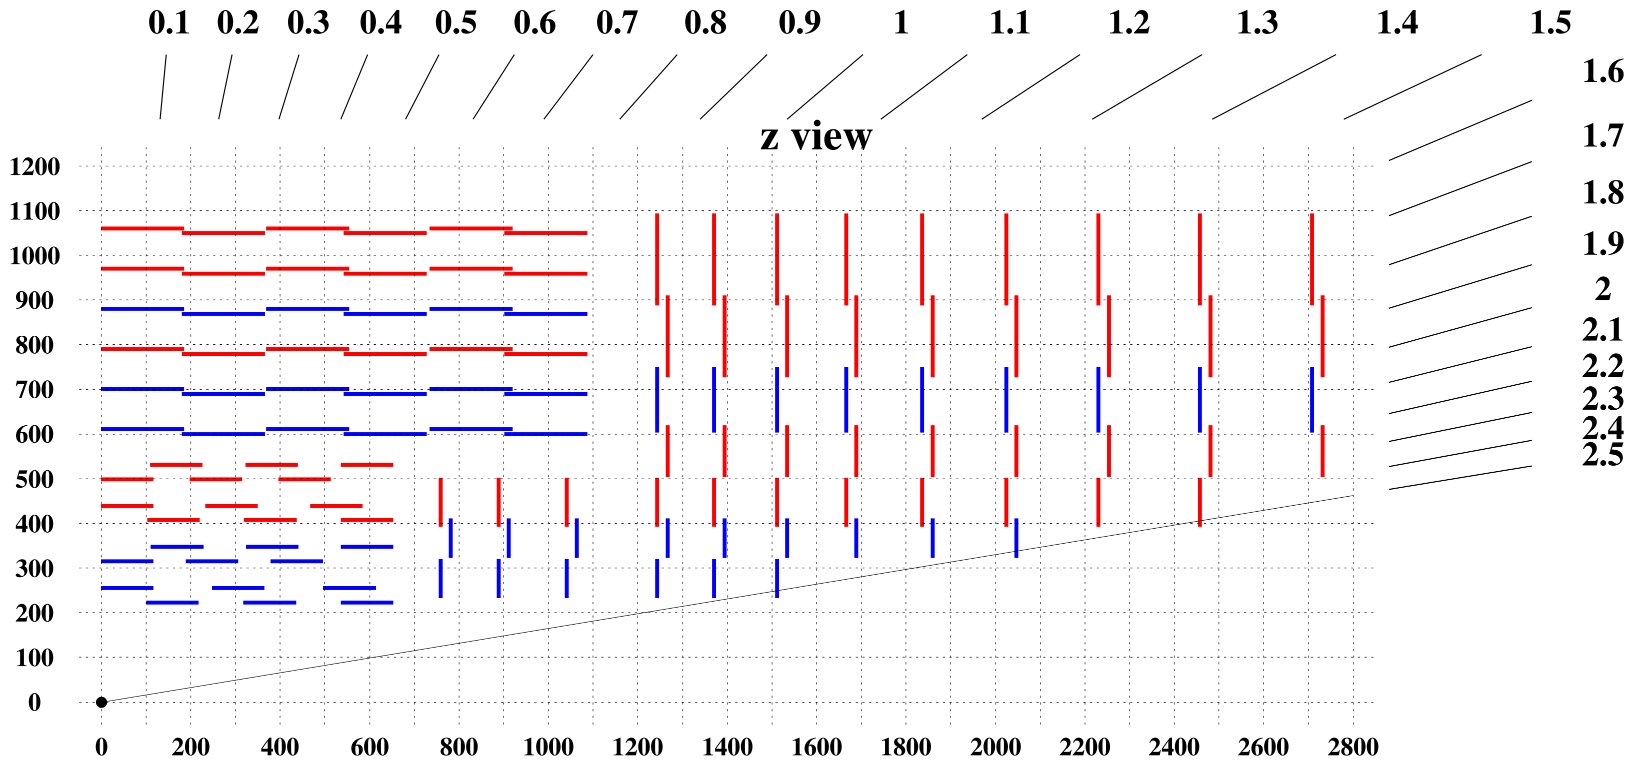

Figure 1: The latest (13/5/00) design of the all silicon tracker.

by the RDPM: only triggered events would be forwarded to the switch and other events would be deleted. To profit from this reduced data flow between the RDPMs and the switch, one would probably reduce the switch input performance to 100 Mbytes/sec and allow each switch input to read about 5 tracker RDPMs. Since a L2 trigger would not reduce the data flow into the RDPMs, its introduction would not affect any of the results presented in this note.

Trigger, timing and control information for the APVs is distributed from the FECs. Each FEC will be connected to a ring of about 10 CCUs, (or less than this if more money is available). Each CCU will feed typically 6–12 modules.

2.1 Hardware Constraints

In the barrel, fibre optic cables from each pair of APVs will run along the barrel rows. At the end of the row, they will be merged via a patch panel into twelve-way cables. Groups of eight of these 12-way cables will then be loosely bundled together, before finally being merged into a 96-way cable in a patch panel near the ECAL.

To simplify wiring near the FEDs, it is required that each 96-way cable feeds exactly one FED. Furthermore, to simplify wiring in the detector, the fibre optic cables which are merged into a 96-way cable must come from the same barrel row, or nearby rows in the same layer, or rarely from rows in neighbouring layers.

To simplify installation and debugging of the system, it is considered highly desirable that no FED should take its inputs from APVs connected to more than one CCU ring. Now, one CCU will sit at the end of each barrel row controlling all the APVs in that row. CCUs from several neighbouring rows in the same layer will form a ring. This implies two constraints: (i) Each CCU ring must contain a multiple of 192 APVs. (ii) A FED cannot take inputs from two different layers. The second of these two constraints implies that where the number of APVs in a layer/disk is not a multiple of 192, a few of the FED input channels will remain unused. This is wasteful of resources, so it may be necessary to reconsider this.

In the end-caps, fibre optic cables will run along the petals. Similar considerations to those described above for the barrel, imply that a FED must first read all the APVs in a petal, before going to the neighbouring petals. Furthermore, all its inputs must come from one end-cap disk.

3 Expected Strip Occupancy

In high luminosity p–p collisions, one expects an average of 24 superimposed minimum bias events per bunch crossing. The mean strip occupancy under these conditions is 0.011. It depends strongly on radius as shown in Fig. 2, but has little dependence on the z-coordinate. Further details can be found in [2]. The strip occupancy is defined as the number of strips associated with a reconstructed cluster, divided by the total number of strips in the detector.

In Pb–Pb collisions, it is necessary to read out the outer five detector layers in the

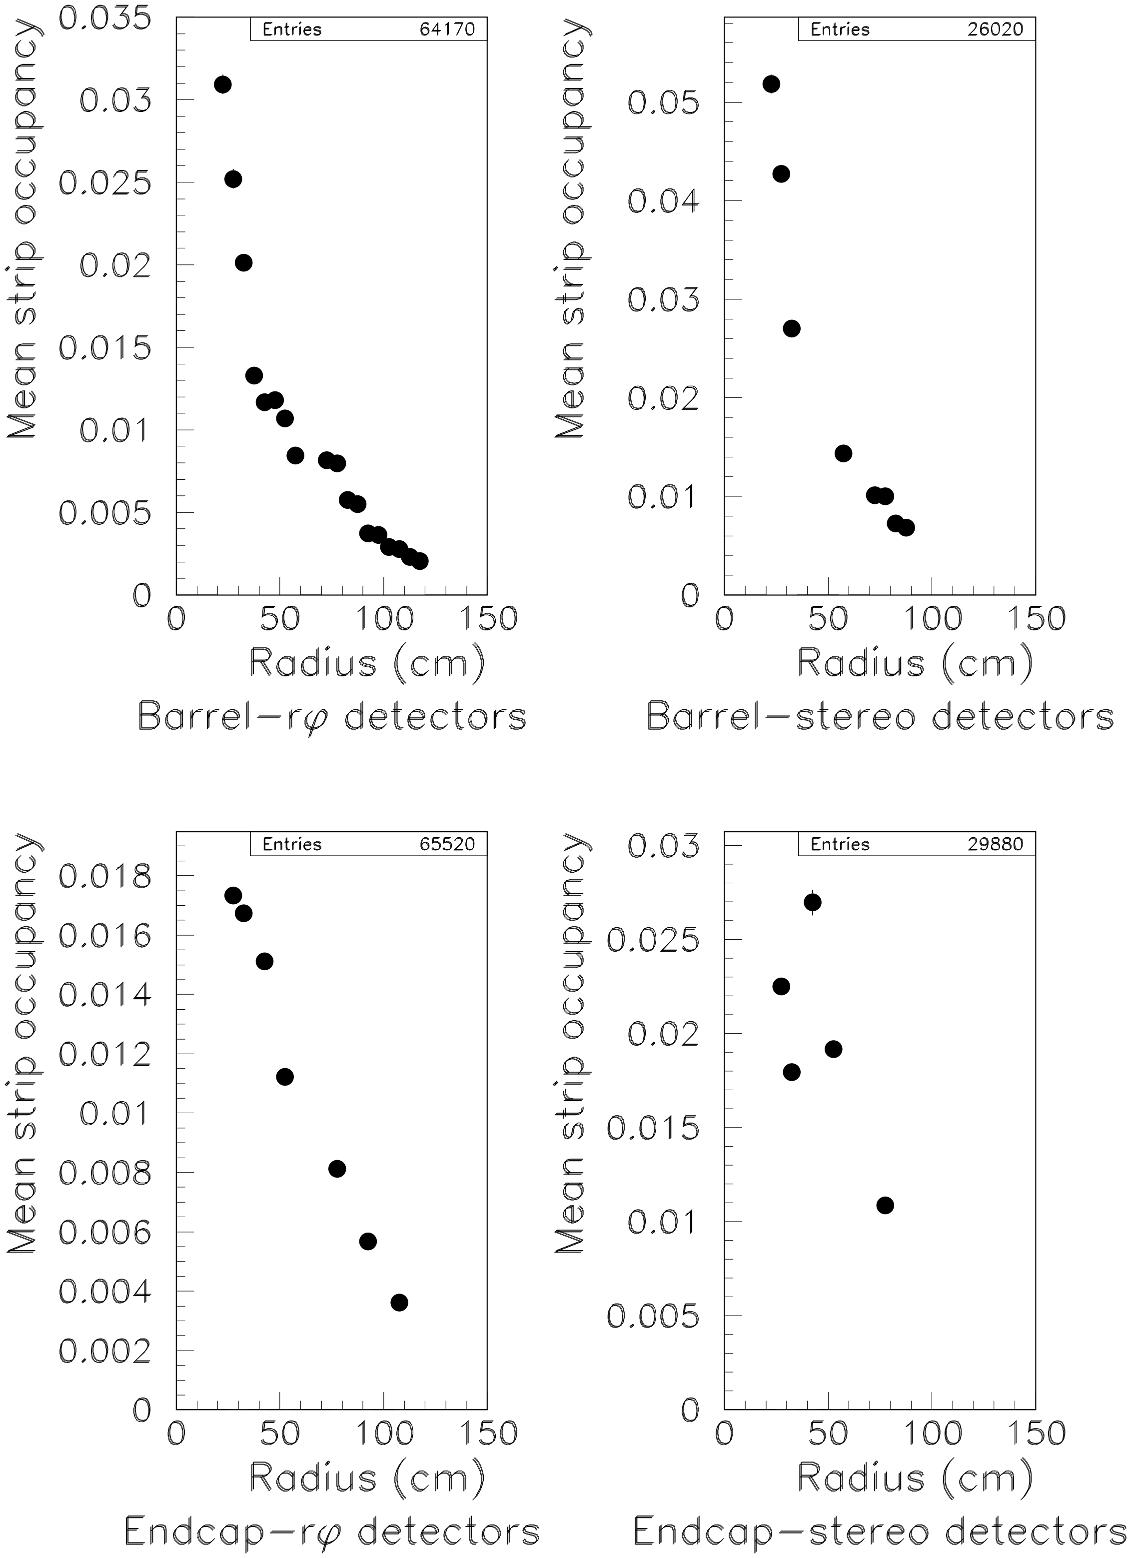

Figure 2: Dependence of strip occupancy on radius for barrel/end-cap and $r-\phi$/stereo detectors.

rapidity region $|\eta| < 1.5$. It would be good to also read out additional detector layers, but this may be impractical due to the high occupancies. The strip occupancy in barrel layers 6–10 are expected to be 0.16, 0.12, 0.08, 0.052 and 0.034 respectively [3].

The numbers for both p–p and Pb–Pb collisions should be treated with some caution for several reasons:

- The length and strip pitches of the detector has evolved since the study was done. This renders it inaccurate at about the 25% level. In particular, the occupancy in the innermost barrel layer is now likely to be roughly 0.04 for both $r - \phi$ and stereo layers.
- Clusters were reconstructed using the TDR cluster finding algorithm [4]. It is believed that a much simpler algorithm will be used in the FED. This will be less effective at rejecting out of time clusters, leading to an increase in the strip occupancy by an estimated factor of 1.2. Furthermore, it will be less effective at rejecting noise, leading to an additional increase in strip occupancy of 0.003 independent of position in the detector. These corrections, which were derived using a toy Monte Carlo program, will be applied to all occupancy/data flow numbers given below in this note. They are however, only approximate, and will depend upon the particular choice made for the FED clustering algorithm.
- The physics simulation of PYTHIA may not be perfect.
- If the LHC accelerator performs very well, achieving a beam-beam tune spread of 0.015 (as has been done in the SPS), then the luminosity could be 2.3 times higher than that usually assumed during the high-luminosity running [5].

4 Data Flow

It is foreseen that the FED will use a simple cluster finding algorithm and only output to the DAQ, information for those strips associated with clusters. Assuming 2 bytes of information are needed per strip (pulse height plus strip address), and that an 80 byte event header is also output, the data rate R in Kbytes/s from each FED is given by:

$$R = T[80 + 2 \times S \times 128 \times (96f \times 2)] \tag{1}$$

where T is the L1 trigger rate in KHz and S is the absolute strip occupancy. In p–p collisions, the L1 trigger rate will be 100 KHz, whilst in Pb–Pb collisions, it will be 1 KHz [3]. The variable f is the fraction of the 96 FED input channels which are in use, and should be close to unity.

It is possible that a data compression algorithm will be employed in the FED which may reduce these rates by a factor of 1.2–1.5. This possibility will be neglected in numbers given below.

Tables 1a-b show the number of FEDs needed per barrel-half layer or per end-cap disk. Where the number of APVs in a layer/disk is not a multiple of 192, the fraction of the FEDs' input channels which are used is slightly less than unity, as is also shown in these tables. An important exception occurs in the innermost barrel layer: here, in principle, one only needs 11 FEDs in each barrel-half to read all the APVs. However, this would lead to an unacceptably high data rate per FED of 214 Mbytes/s. Hence it is proposed to increase the number of FEDs to 16, and use only a fraction of the available FED input channels.

Tables 1a-b also show give the mean strip occupancies and corresponding data flow rates per FED expected in the barrel layers and end-cap disks. These have both been corrected for the effect of the FED clustering algorithm mentioned in Sect. 3.

One finds a total number of FEDs in the tracker of $2 \times (122+95) = 434$. On average, 94% of the FED input channels are used. Where the FEDs in a given layer/disk do have unused input channels, these should be evenly divided among the those FEDs.

In p–p collisions, the total data rate from the tracker FEDs is $2 \times (10.1 + 8.0) \approx 36$ Gbytes/s, corresponding to a mean event size of 0.36 Mbytes. Averaging over the entire tracker, the corrected mean occupancy is 0.016.

5 Connecting FEDs to RDPMs

From Sect. 4, it is clear that data rates from Pb–Pb collisions will be much lower than from p–p collisions. Hence as far as balancing mean data flow rates is concerned, we can concentrate on the latter. Given the uncertainties in data flow rates referred to in Sect. 3, it is important to redo these studies using a Monte Carlo based on the final silicon geometry and using the FED cluster finder. Even then, a substantial safety margin should be left. Nonetheless, one might imagine RDPMs reading the FEDs given in Tables 2a-b. In total, the number of RDPMs required is then $2 \times (66 + 54) = 240$. In producing these tables, an effort was made not only to balance data rates, but also to arrange that the FEDs read by each RDPM corresponded to a fairly small region in ϕ. This may be useful for a L2 trigger. However, other configurations may perform equally well, and the main purpose of these tables is just to illustrate what might be done.

The data rates quoted in Tables 2a-b neglect the additional factor of 1.2–1.5 compression which might be achieved using the FED data compression algorithm mentioned in Sect. 4. Assuming that some part of this can be achieved, one one would obtain a safety factor of about 1.5, with respect to the nominal RDPM input rate of 200 Mbytes/s. Ideally, one would have liked something closer to a factor of two ...

References

[1] *"The Tracker Layout Working Page"*, http://duccio.home.cern.ch/duccio/layout/lay.html.

[2] Tracker TDR Appendum, CERN/LHCC 2000-016.

[3] A. Racz *et al.*, *"Data acquisition for heavy ion physics"*, CMS IN 2000/027.

[4] Tracker TDR, CERN/LHCC 98-6.

[5] Private communication from F. Ruggiero, SL-AP group leader.

Table 1a: The number N of FEDs and the fraction f of their input channels in use, and tracker occupancies and FED data flow rates expected in the layers of each half of the barrel. Results are given for p–p and Pb–Pb collisions.

Layer	FEDs		p–p Collisions		Pb–Pb Collisions	
	N	f	Mean Occupancy	Data Rate/FED (Mbytes/s)	Mean Occupancy	Data Rate/FED (Mbytes/s)
1	16	0.66	0.044	151	—	—
2	15	0.95	0.031	153	—	—
3	6	0.97	0.017	89	—	—
4	7	1.00	0.016	87	—	—
5	14	0.94	0.015	77	—	—
6	15	1.00	0.014	77	0.19	9.4
7	7	0.97	0.010	56	0.15	7.2
8	15	1.00	0.007	42	0.10	5.0
9	13	0.95	0.006	36	0.07	3.3
10	14	0.99	0.006	37	0.04	2.0

Table 1b: The number N of FEDs and the fraction f of their input channels in use, and tracker occupancies and FED data flow rates expected in each end-cap disk. Results are given only for p–p collisions, as insufficient information exists for Pb–Pb collisions.

Disk	FEDs		p–p Collisions	
	N	f	Mean Occupancy	Data Rate/FED (Mbytes/s)
1	4	0.95	0.025	125
2	4	0.95	0.025	125
3	4	0.95	0.025	125
4	11	0.94	0.016	82
5	11	0.94	0.016	82
6	11	0.94	0.016	82
7	9	0.98	0.015	80
8	9	0.98	0.015	80
9	9	0.98	0.015	80
10	8	0.91	0.014	71
11	8	0.91	0.014	71
12	7	0.93	0.013	67

Table 2a: Proposed RDPM-FED links in each barrel-half. The RDPMs have been subdivided into categories depending on the number of FEDs they read from each layer.

RDPM category	Number of RDPMs of this category	Reads FEDs from layers	Data Rate/RDPM (Mbytes/s)
1	16	1	151
2	15	2	153
3	6	3,7	145
4	7	4,6	164
5	8	5,8,10	156
6	6	5,9,10	150
7	7	6,8,9	155
8	1	6,7	133

Table 2b: Proposed RDPM-FED links in each end-cap. The RDPMs have been subdivided into categories depending on the number of FEDs they read from each disk.

RDPM category	Number of RDPMs of this category	Reads FEDs from disks	Data Rate/RDPM (Mbytes/s)
9	4	1	125
10	4	2	125
11	4	3	125
12	9	4,7	162
13	2	4,12	149
14	11	5,6	164
15	9	8,9	160
16	8	10,11	142
17	2	12,12	134
18	1	12	67