SDC-92-247

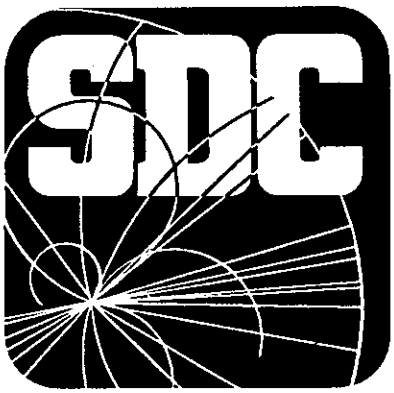

SDC
SOLENOIDAL DETECTOR NOTES

FRONT END AND DCC SIMULATIONS FOR THE SDC STRAW TUBE SYSTEM

April 21, 1992

A. Holscher, G. Stairs and P.K. Sinervo
University of Toronto

Front End and DCC Simulations for the SDC Straw Tube System

A. Hölscher, G. Stairs and P. K. Sinervo
University of Toronto
Toronto, Canada

April 21, 1992

contact: A.Hölscher
e -mail: holscher@cerebus.physics.utoronto.ca

This document is a collection of three notes, which describe the buffer and bandwidth requirements at various stages of the SDC straw tracker front end system.

1 Simulations of Front End Board Occupancies for the SDC Straw Tube Tracker

This note describes the results of a study of the front end board occupancies for the different superlayers of the SDC straw tracker.

2 Front End Buffer Requirements for the SDC Straw Tube Tracker

This note describes estimates made of the buffer requirements for the L1 and L2 buffer.

3 Buffer Occupancies for the SDC Straw Tube DCC System

This note describes a study of the buffer and bandwidth requirements on the front end boards and at the DCC level.

Simulation of Front End Board Occupancies for the SDC Straw Tube Tracker

A. Hölscher, P. K. Sinervo and G. Stairs
University of Toronto
Toronto, Canada

Penny Estabrooks
~~Carlton~~ Carleton University
Ottawa, Canada

March 11, 1992

Abstract

In this note we present simulations of the occupancies of the Front end boards for the straw system using SDCSIM generated events.

1 Introduction

The straw system consists of about 110,000 straws, which are arranged in 5 superlayers. For an exact description of the geometrical arrangement we refer to [1]. In this note we determine occupancies of different layers and of the electronic front end boards. We use physics events generated by the SDC simulation program. No attempt has been made to generate electronic noise.

2 Superlayer occupancies

First we show the occupancies of different event types alone. However, at the nominal luminosity of the SSC of $L = 10^{33} cm^{-2} s^{-1}$, there are 1.6 minimum bias events per crossing. For the straw system the detector integrates over three bunch crossings per each trigger in order to catch all signals. This implies that for each physics events one triggers on, one has roughly 5 underlying minimum bias events, which of course add to the occupancies. We

generated 5 minimum bias events, 5 events containing a Higgs decaying in the following decay chain: Higgs $\rightarrow$ ZZ $\rightarrow 4\mu$ and 5 dijet events with a p_t in the range of 1 - 100 GeV/c. Table 1 shows the occupancy per channel per trigger, first for the different event types alone and then for the event types of interest plus the underlying 5 minimum bias events. Figure 1 shows these distributions for the higgs + minbias, dijet + minbias and 1 minbias cases represented graphically.

	Layer 1	Layer 2	Layer 3	Layer 4	Layer 5
1 minbias	0.83 %	0.35 %	0.28 %	0.16 %	0.16 %
1 Higgs	6.0 %	3.0 %	1.9 %	1.75 %	1.75 %
1 dijet	2.7 %	1.3 %	0.85 %	0.64 %	0.63 %
1 Higgs + 5 minbias	10.0 %	4.7 %	3.3 %	2.5 %	2.5 %
1 dijet + 5 minbias	6.8 %	3.1 %	2.3 %	1.4 %	1.4 %

Table 1: Layer occupancies for different types of events

3 Front end board occupancies

We are also interested in correlations in specific Front end boards. As shown in Figure 2 each superlayer consist of 6 or 8 layers, which are bundeled together into modules and read out by one Front end board (FEB). It is assumed that each front end board contains 160 straw channels. This implies that there are about 850 FEBs for the whole straw system. Each FEB consists of five so called FMUX's, each reading out 32 channels [2]. Therefore there will be about 4250 FMUX's in the whole straw system. The FMUX's are assumed to be radially arranged.

A particle transversing one superlayer will leave several hits in the particular front end board. Figure 3 and 4 [1] show the occupancies of the FMUX's for 1 minbias event and 1 dijet event alone. One clearly sees peaks at an occupancy of 6, which indicates that the transversing particle has a hit in all of the layers of one superlayer. These 6 hits mostly belong to one FEB, which is why there is a clear peak in the occupancy distribution of the FEB. The 6 hits may also belong to one FMUX, but here the correlation is not so strong. In comparison we show the occupancy distribution assuming that the modules would be arranged in the ϕ direction, where one does not have this correlation effect. Many more

[1] Note that the mean of the occupancy distributions was only calculated for FEB's or FMUX's, which contain at least one hit.

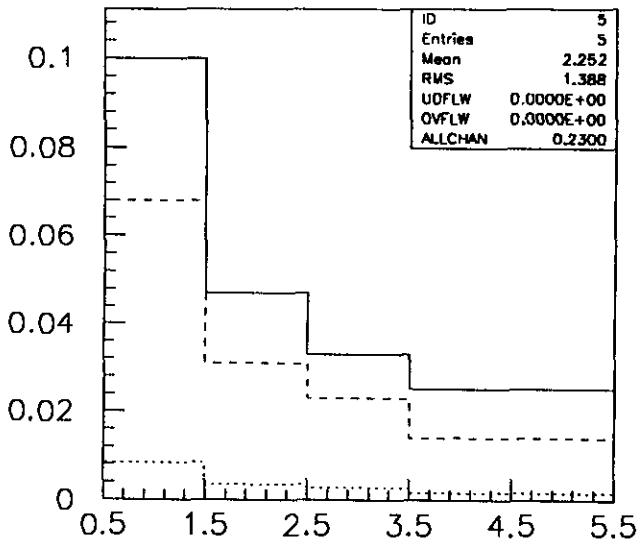

Figure 1: Occupancies of the different superlayers of the straw system for a higgs event (solid histogram) and a dijet event (dashed histogram). These events include 5 underlying minimum bias events. For comparison the occupancy for one minimum bias event (dotted histogram) is shown.

modules contain a hit, but the average occupancy of a hit module or FMUX is much smaller. Note that the average FEB and FMUX occupancies are very similar for the minbias and dijet events. However, in the minbias case fewer modules actually have one or more hits. Figure 5 finally shows the occupancy distribution of the Higgs or dijet events plus 5 underlying minbias events. The shapes don't change very much since the additional particles from the minbias events mostly transverse other modules. This is also illustrated in table 2, which shows the percentage of FMUX's or front end boards that don't contain any hit for all the different events categories. Including the minbias events clearly decreases the number of empty FMUX's and FEB's, inplying that they hit different modules. We remind the reader however, that we didn't generate any electronic noise, which might change this picture.

The occupancy distributions, however, show long tails. For the generated 5 events of each category an occupancy of 105 for one FEB was observed and an occupancy of 41 for one FMUX.

	empty FEB's	empty FMUX's
1 minbias	95 %	98 %
1 higgs	68 %	85 %
1 dijet	84 %	93 %
1 higgs + 5 minbias	55 %	77 %
1 dijet + 5 minbias	66 %	84 %

Table 2: Fraction of FEB boards and FMUX's that do not contain any hits.

4 Conclusions

The maximal inner superlayer occupancy was found to be 10%. This translates into the following average datarates from the front end boards, assuming that each hit contains 4 bytes and an L2 trigger rate of 10 kHz and assuming that each front end board contains 160 straws. Again this table ignores the electronic noise. These datarates, however, are

Layer 1	Layer 2	Layer 3	Layer 4	Layer 5
0.16 hits/μs	0.08 hits/μs	0.05 hits/μs	0.04 hits/μs	0.04 hits/μs
0.64 Mbyte/s	0.32 Mbyte/s	0.2 Mbyte/s	0.16 Mbyte/s	0.16 Mbyte/s

Table 3: Datarate from one FEB for the Higgs + 5 minbias events, assuming 4 bytes/hit and excluding electronic noise.

the average datarates. For any event the hits are correlated. When a module is hit by a transversing particle, it usually contains at least 6 hits. It can contain many more hits, when it is hit by a jet. In most cases, however, the Front end boards are empty.

References

[1] SDC detector notes, SDC-91-125, SDC modular straw outer tracking system conceptual design report.

[2] G.Stairs et al., *Front end Collector Bus Design Description*, UofT-DCC-03, Jan. 1992 (unpublished).

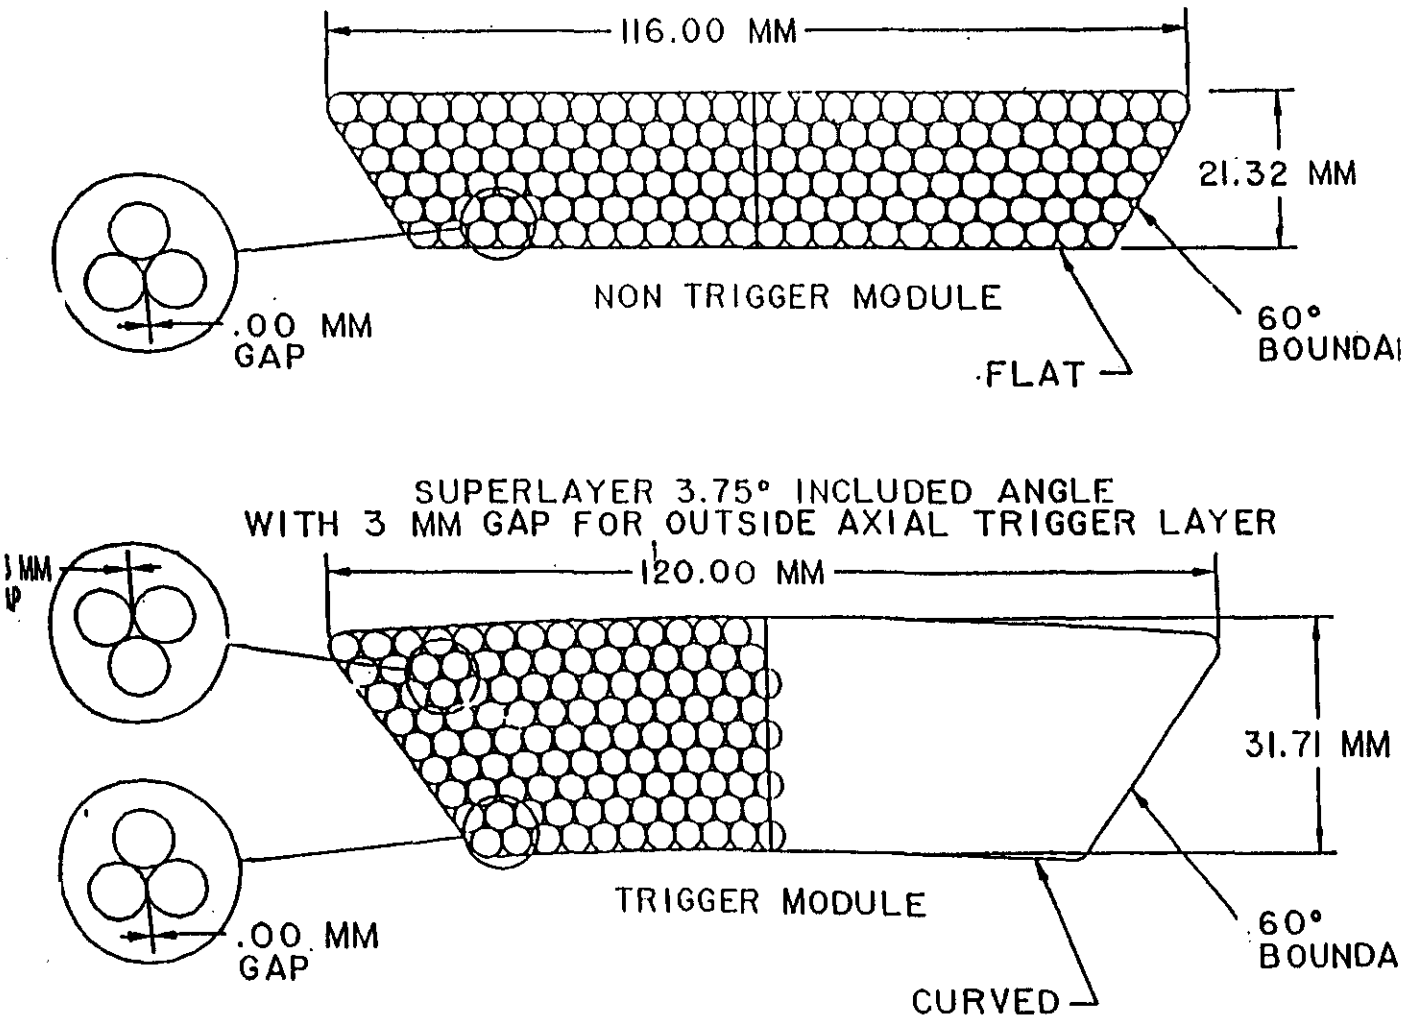

Fig. IV.2. The basic module design is shown here schematically. The two basic designs are shown; a 6 layer module for non-trigger layers and stereo layers, and a 9 layer trigger module that has all straws on a radial line. The present design will probably use an 8 layer trigger module.

Figure 2: Module design within the superlayers. For details see [1].

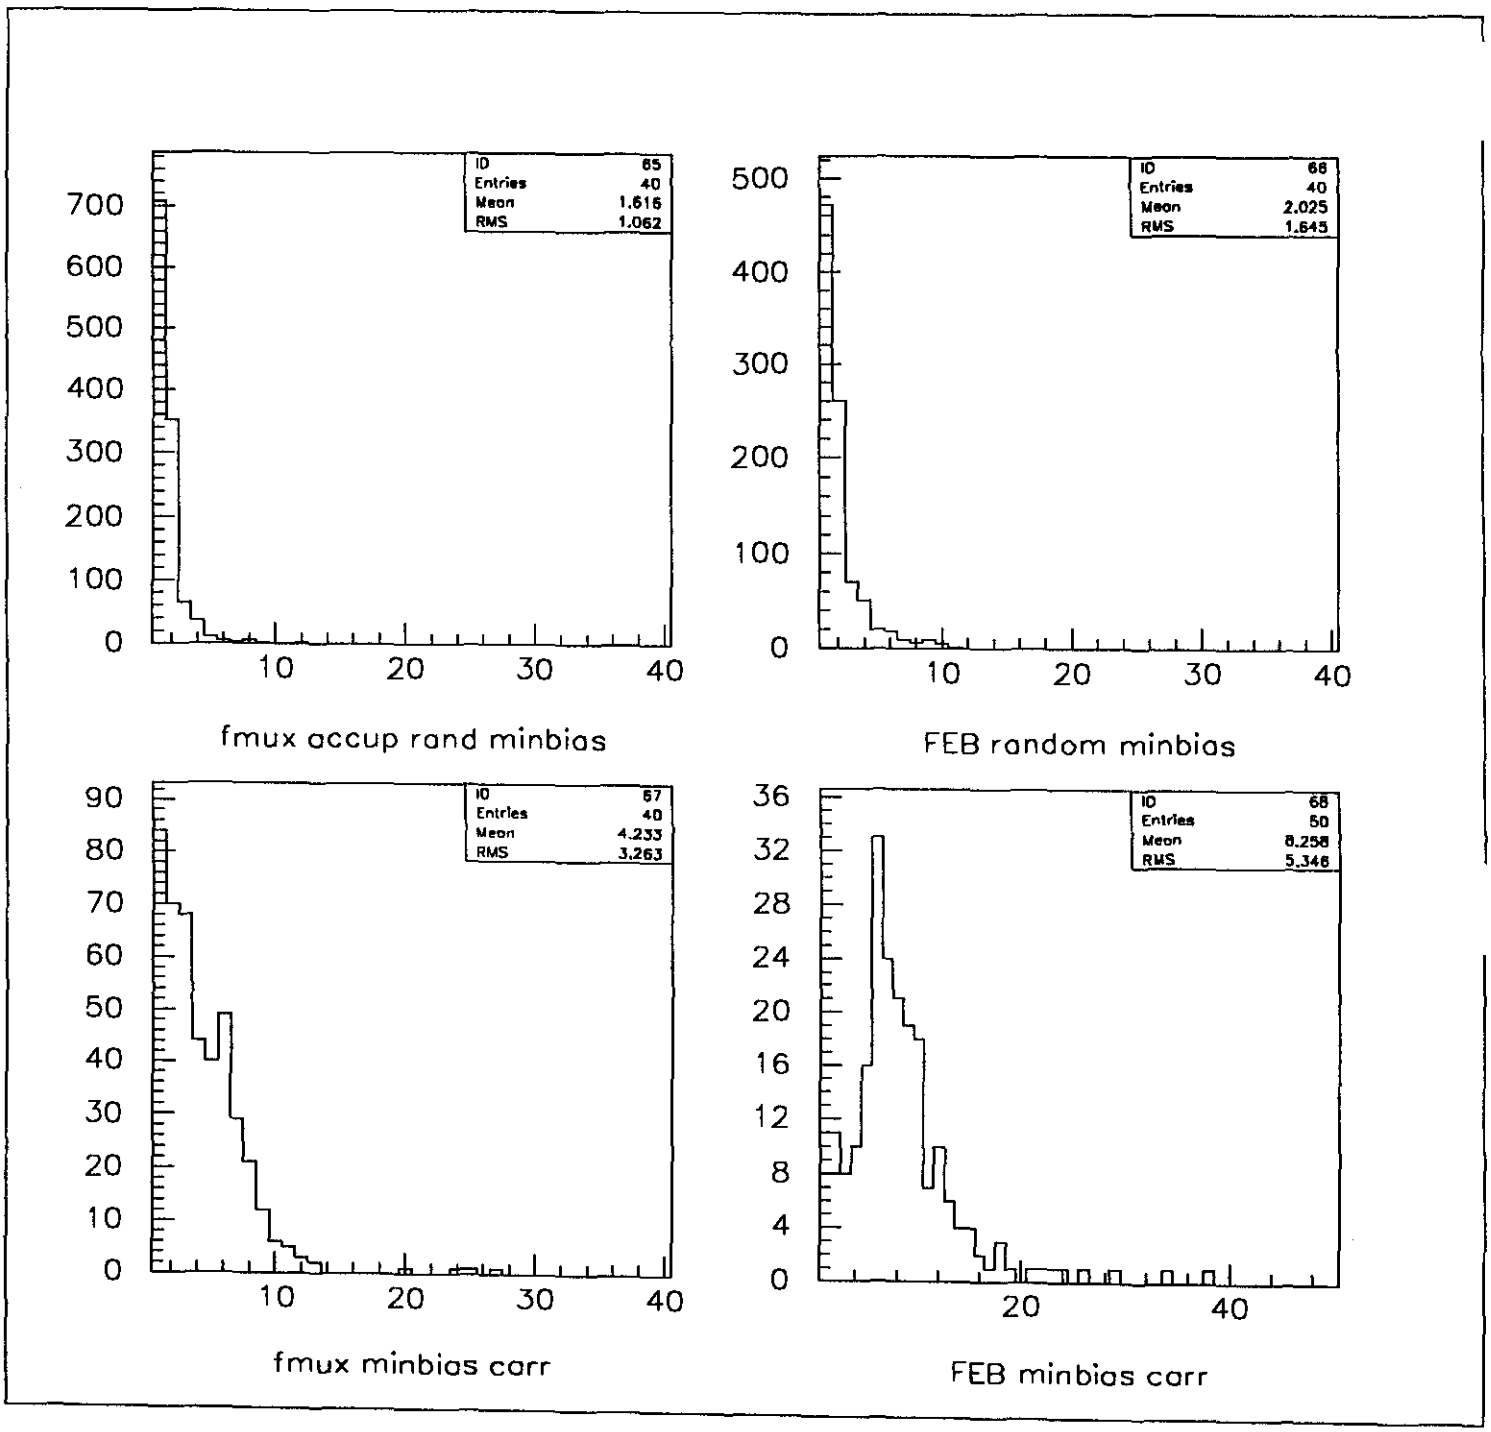

Figure 3: FEB and FMUX occupancies for one minimum bias event. The upper two plots show the modules arranged in ϕ, the lower two like in the real design.

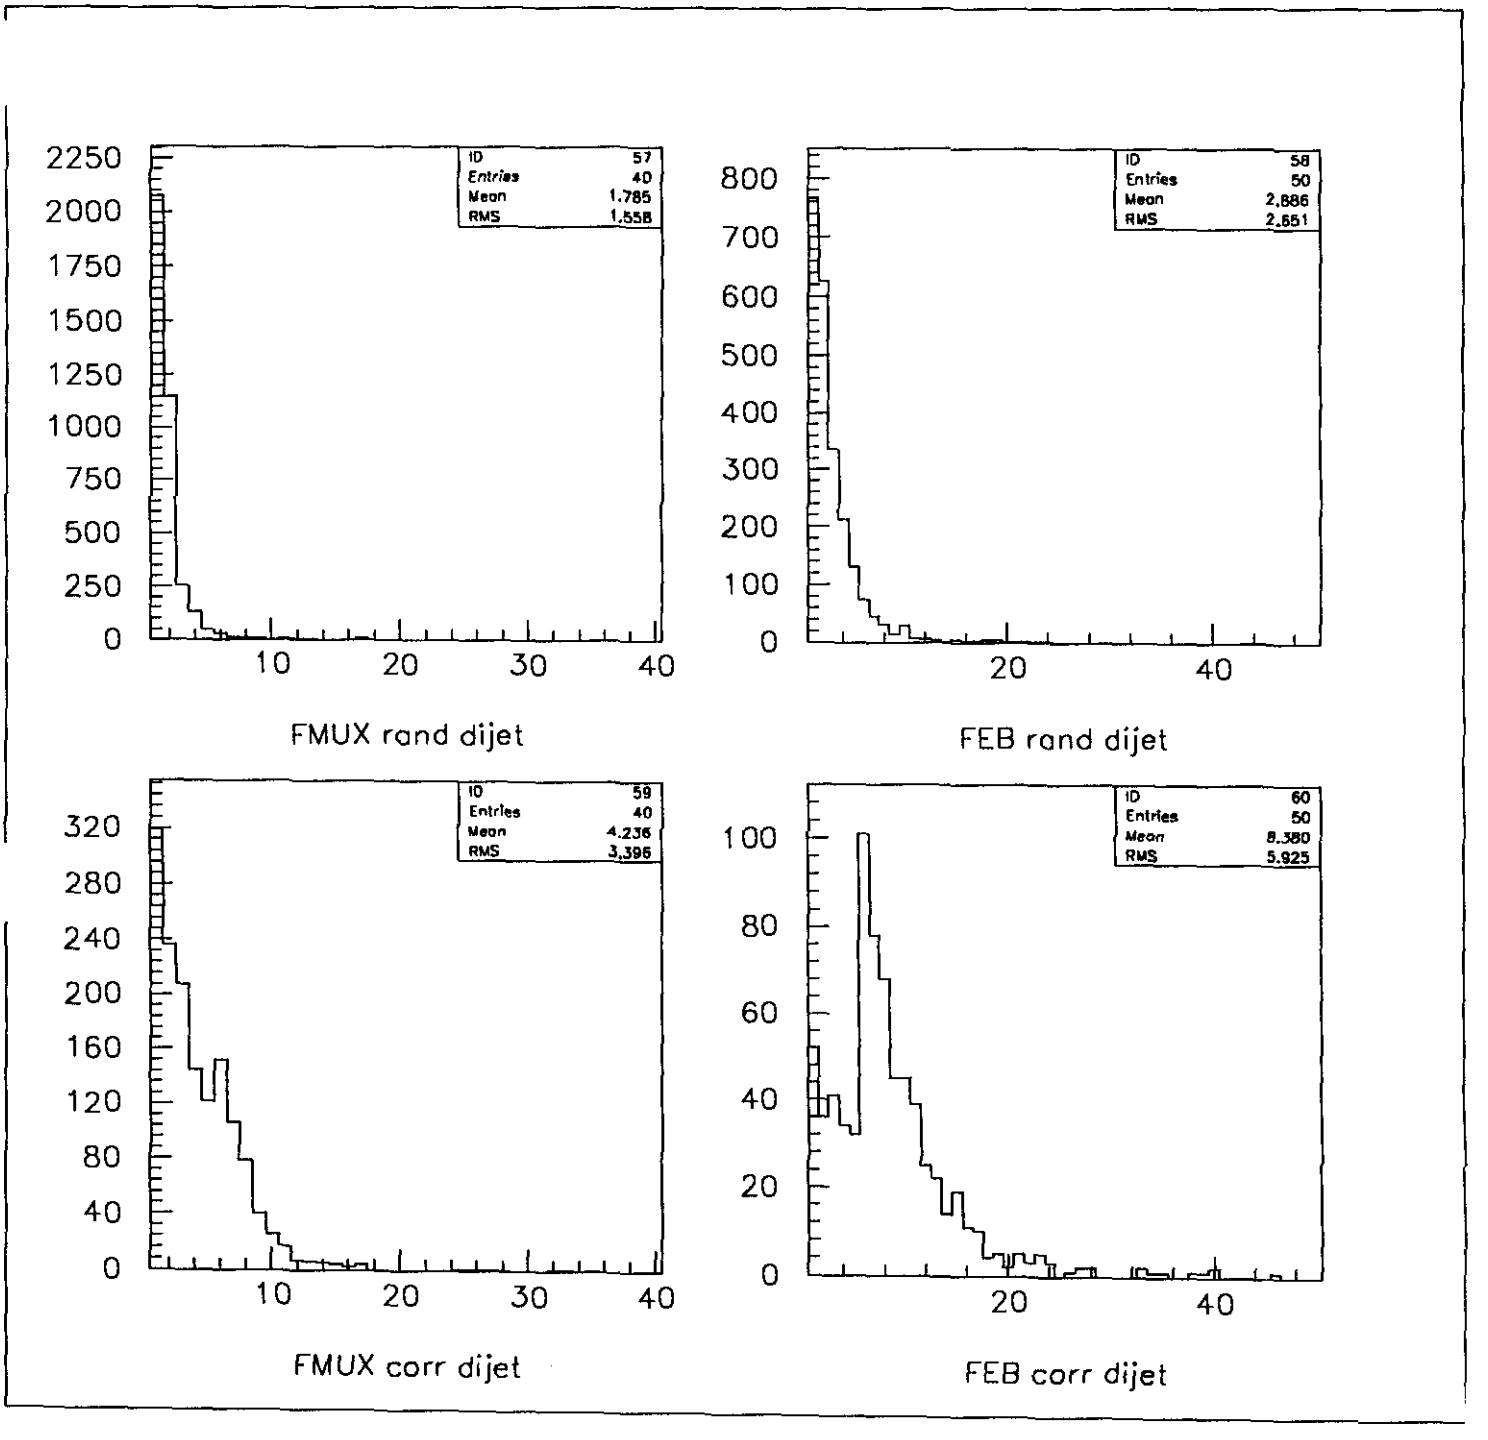

Figure 4: FEB and FMUX occupancies for a single dijet event without any underlying minimum bias events.

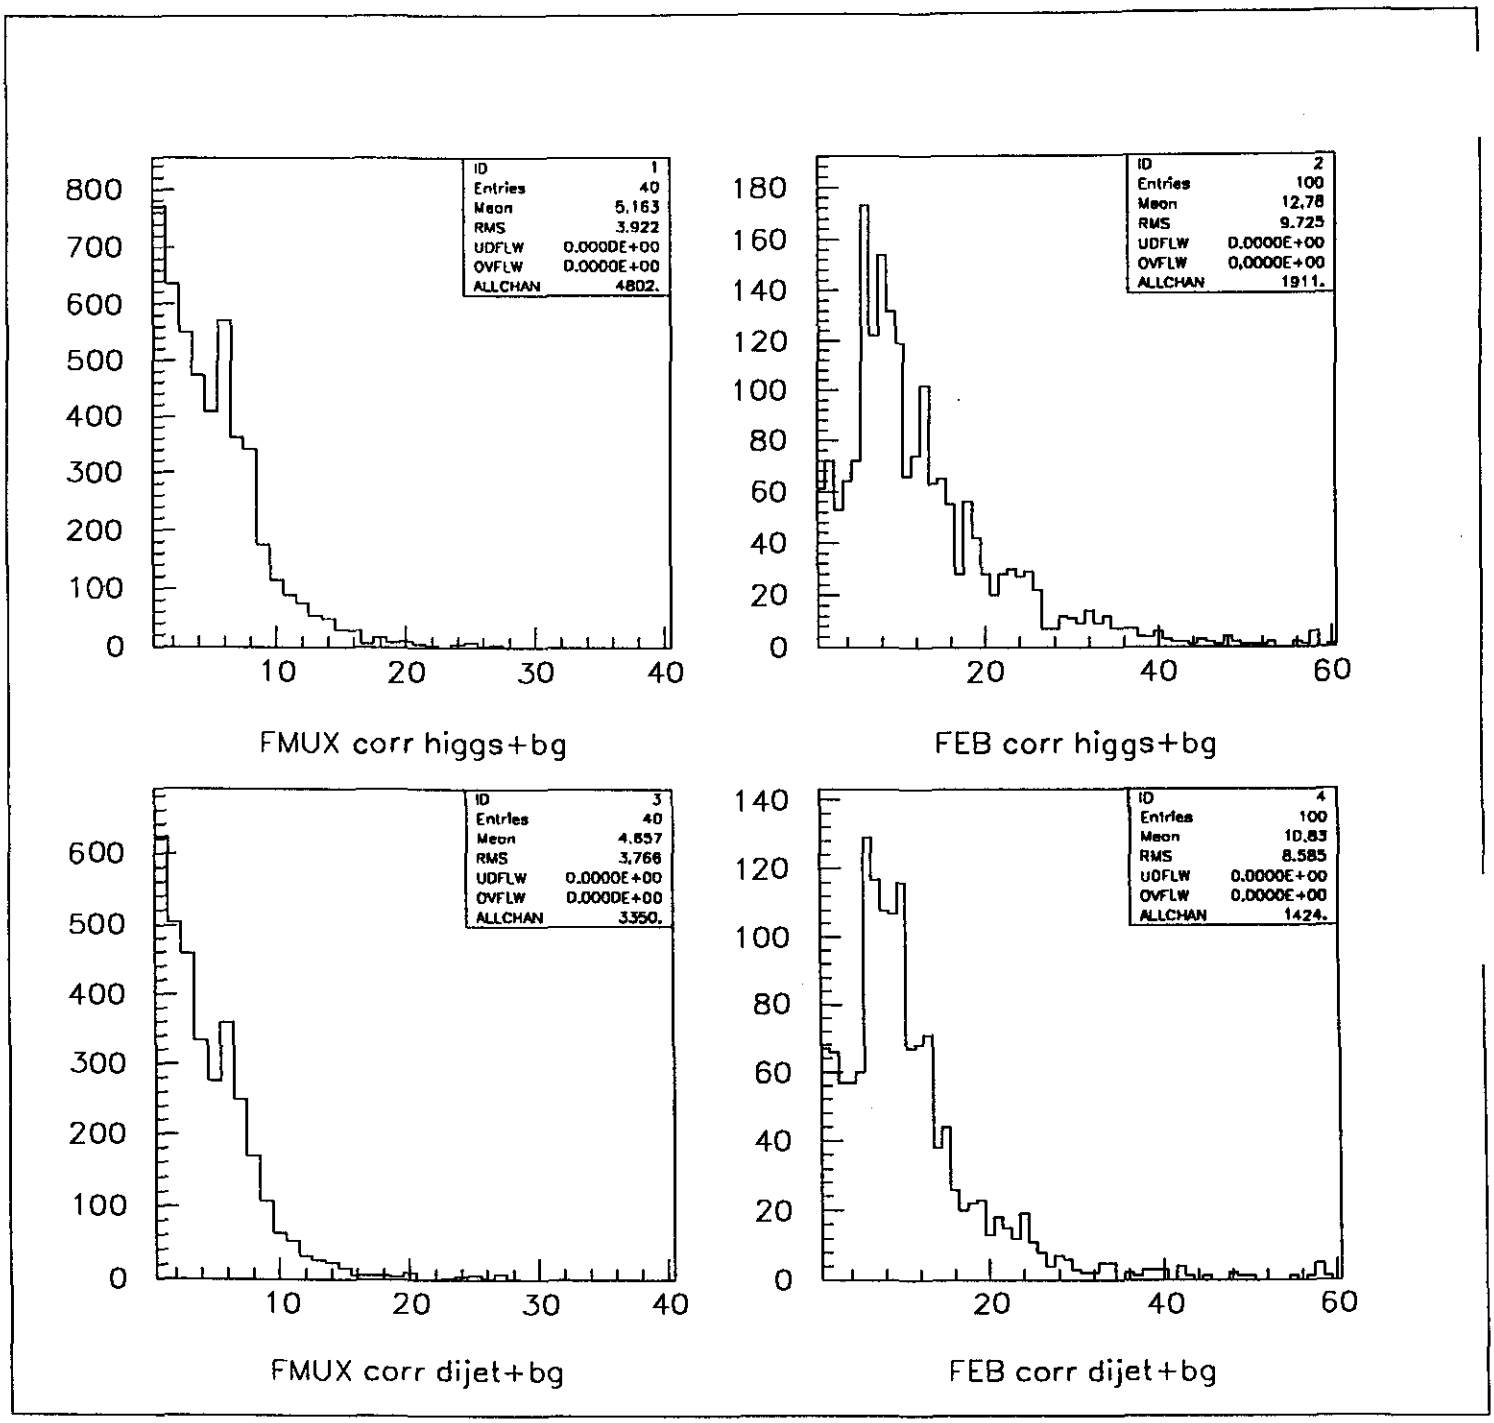

Figure 5: FEB and FMUX occupancies for Higgs events (upper two histograms) and dijet events (lower two histograms). Both events contain 5 underlying minimum bias events.

Front End Buffer Requirements for the SDC Straw Tube Tracker

A. Hölscher, G. Stairs and P. K. Sinervo
University of Toronto
Toronto, Canada

March 16, 1992

Abstract

In this note we present the results of calculations of the buffer occupancies at the front end card for the straw system. For an occupancy per crossing of 5%, the L1 buffer would need 18 storage locations and the L2 buffer 8 storage locations. One can impose a minimal distance requirement between two L2 trigger accepts of up to 50 μs with only little additional burden to the L2 buffer.

1 Introduction

We simulate the Front end TVC/AMU for the straw system to determine the buffer requirements of the L1 and L2 buffer. The system is supposed to meet the following requirements:

- sustain a L2 trigger accept rate of up to 10 kHz, with less than 10 % losses.
- work for occupancies of up to 5% / crossing and integrate the signal over 3 crossings.

Model simulations show that the occupancuies of the inner layers of the straw system are about 10% per trigger without any electronic noise. The above assumption contains some safety margin to incorporate this noise.
The model used is a simplified version of the front end simulation model of Sinervo et al. [2]. It uses the beam clock as the minimal time step and starts generating events for any L1 accept. The L1 triggers are generated with an $L1$ accept rate between 62MHz*(0.001-0.003). For L1 trigger accepts, hits are generated with a probability/crossing, occup=0.05. For each trigger 3 beam crossings are read out. These hits then enter the L2 pipeline. The

L2 delay ($L2lat$) is parametrized by a a flat uniform distribution between 10 and 50 μs (a) or 10 and 100 μs (b). The trigger proposal [1] requires the L2 latency time to be between 10 and 50 μs. Here we will usually allow for a latency of up to 100 μs. L2 trigger accept signals are generated with the probability $L2$. We further require a minimum time $L2min$ between two L2 triggers. This time is the same for L2 accepts or L2 rejects.

2 L1 buffer occupancy

The L1 buffer occupancy depends only on the delay of the L1 trigger decision, which we take to be 240 crossings, and the occupancy of one channel per crossing. It follows a binomial distribution:

$$prob(n) = \frac{240!}{(240-n)! * n!} * occup^n * (1-occup)^{240-n}$$

Simulations werde done with a detailed model of P. Sinervo et al. [2]. Figure 1 shows the L1 buffer occupancy distribution for an occupancy of 2% and 5% per crossing. The maximal buffer length to contain 99 % of the hits, is 10 and 18 for the occupancies of 2% and 5% per crossing, respectively.

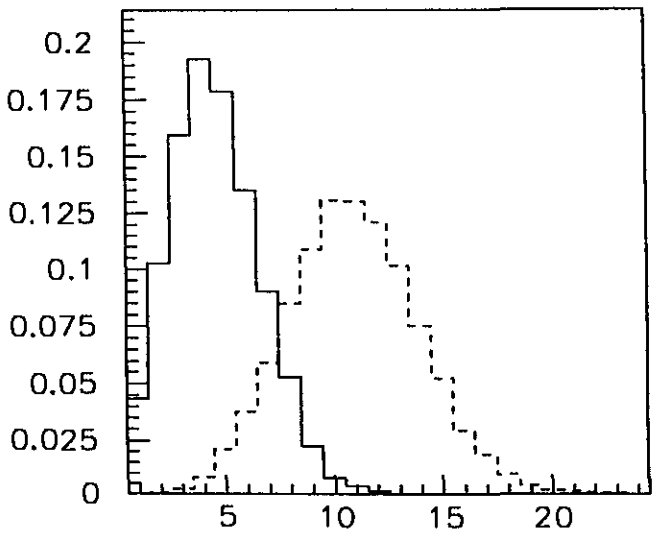

Figure 1: L1 buffer occupancy distribution for an occupancy of 5% per crossing and 2% per crossing.

3 L2 buffer occupancy

We test the required L2 buffer size in a simplified version of the front end simulation model of Sinervo et al. [2]. We note the following parameters for the buffer: the average occupancy n (averaging excludes empty buffer conditions), the variance of the occupancy distribution σ and the maximal occupancy $nmax$ of the L2 buffer. The maximal occupancy is defined as the number of buffers needed to contain more than 99% of the hits. Table 1 gives the results for various parameters.

L1	occup	L2lat	L2min(Xing)	n	σ	nmax
0.001	0.05	a	250	1.3	0.6	4
0.001	0.05	b	250	1.4	0.7	5
0.001	0.1	a	10	1.5	0.8	5
0.001	0.1	b	250	1.9	1.1	6
0.002	0.05	a	250	1.6	0.9	6
0.002	0.05	b	10	1.9	1.1	7
0.002	0.1	a	600	14	7	>24
0.002	0.1	b	10	3.1	1.9	11
0.003	0.05	b	250	2.9	1.6	8
0.003	0.1	b	250	5.0	2.2	13

Table 1: L2 buffer occupancies.

From this table one can infer that the variables n, σ and nmax scale approximately with the square root of L1, occup and L2lat.

$$\mathrm{n}, \sigma, \mathrm{nmax} \propto \sqrt{\mathrm{occup}} * \sqrt{\mathrm{L1}} * \sqrt{< \mathrm{L2lat} >}$$

The introduction of minimal distance between L2 triggers of 250 Xings=4μs does not have a big influence on the performance of the system. This parameter, however, starts to have a large influence as soon as this distance comes close to the mean time between L1 trigger accepts (60 MHz*L1=10-20 μs).
Figure 2 shows the maximal occupancy of the L2 buffer, nmax, as the function of the L1 trigger rate for occupancy/crossing of 0.1, 0.05 and 0.02. For an occupancy of 5% per crossing and a L2 latency time distribution between 10 and 100 μs, a maximal L2 buffer length of 8 hit locations is sufficient.

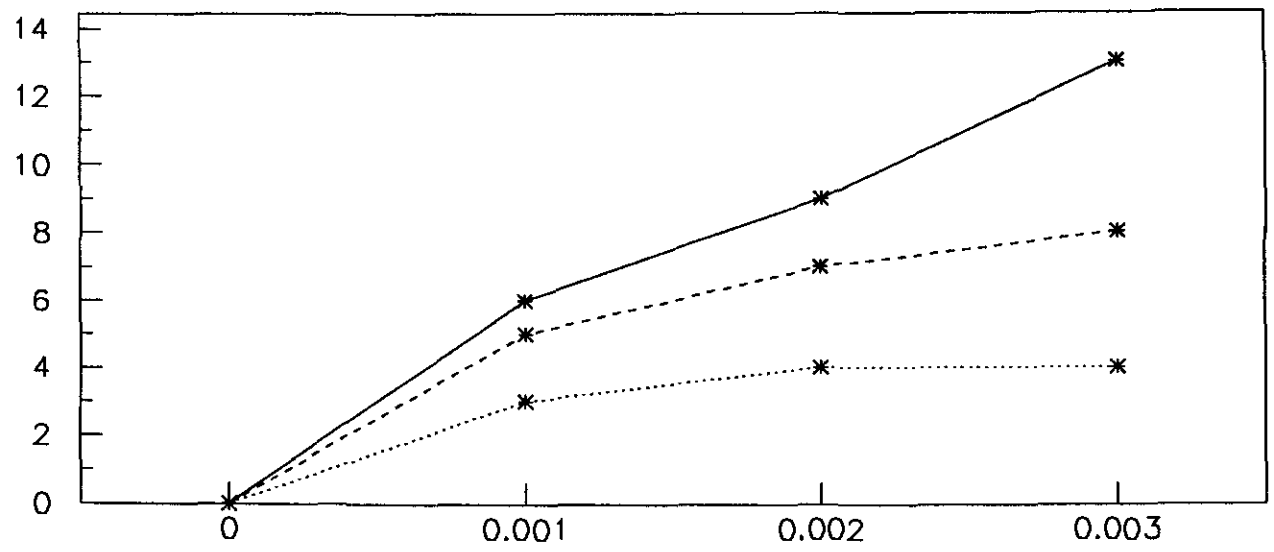

Figure 2: nmax as a function of the L1 trigger accept rate for an occupancy/crossing of 0.1 (upper curve), 0.05 (middle curve) and 0.02 (lower curve).

4 L2 accept separation

For some reasons like event ordering one is interested in the influence of an additional latency between L2 trigger accepts on the performance of the front end buffers. Let τ be the minimal distance between two *L2 trigger accepts*. This is not to be confused with the L2 trigger latency, which is the time the system needs for an L2 trigger decision. This additional latency would require the L2 buffer to take the burden of this additional buffering. For an occupancy of 5% per crossing we show in the following table the L2 buffer occupancies, for different latency times τ.

The last column in the table gives the percentage of L2 trigger accepts that had to be buffered so that the minimal distance criterion was satisfied.
For the L2 latency time we used a uniform distribution between 10 and 100 μs, so that the typical latency was usually much larger, than the additional latency caused by the minimal distance requirement.
The increase in required L2 buffer space on the L2 buffer side seems to be modest, if noticable at all for the assumed parameters. However, as soon as τ approches the L2 accept

L1	L2	L2rate(kHz)	$\tau(\mu s)$	n	σ	nmax	(%)
0.002	0.02	2.4	1	2.1	1.2	5	0
0.002	0.02	2.4	32	2.1	1.2	5	6
0.002	0.02	2.4	64	2.1	1.2	5	11
0.003	0.05	9	1	2.9	1.6	7	0
0.003	0.05	9	32	3.0	1.7	8	25
0.003	0.05	9	64	5.1	2.1	> 8	53

Table 2: L2 buffer occupancies with an minimal L2 trigger accept distance, for an occupancy of 5% per crossing.

rate, the required buffer size increases dramatically. For the last entry, 16 % of the hits are lost for the L2 buffer length of 8.

The small increase in buffer requirement can be understood in the following way: The minimum L2 accept requirement introduces *on the average* a latency of L2 * τ for each event, which is quite small compared to the L2 decision latency, since L2 is small. This effect, however, becomes very important when the probability that during this latency another L2 accept trigger occurs, becomes important. At this point the L2 buffer queue is not emptied any more.

Figure 2 compares the L2 buffer occupancy distribution for (a) L1=0.002 and L2=0.02 and (b) for L1=0.003 and L2 =0.05 for minimal L2 trigger accept separations of 1, 32 and 64 μs. In the case a) the difference in the distributions is nearly imperceptible. In case b) the distributions for $\tau = 1\mu$s and $\tau = 32\mu$s are very similar, while the case of $\tau = 64\mu$s would need more bufferspace. 16 % of the hits are lost for an L2 buffer length of 8.

Figure 3 shows the percentage of L2 accept triggers which had to be delayed, to enforce the L2 trigger accept seperation.

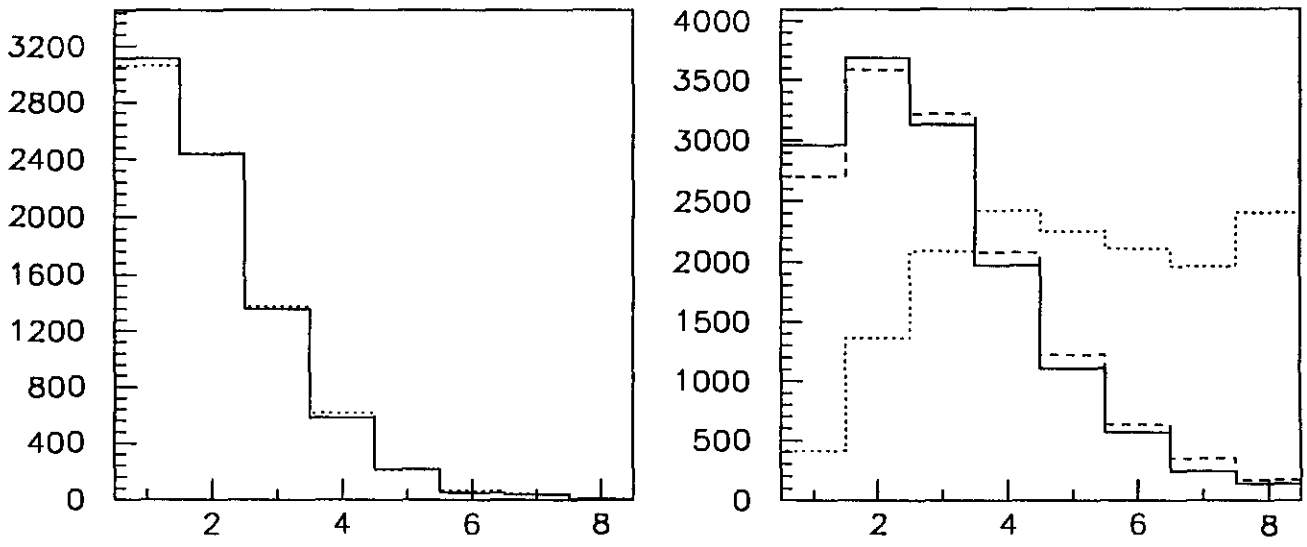

Figure 3: L2 buffer occupancy for a) L1=0.002 and L2 =0.02 and b) L1=0.003 and L2=0.05. The three cases $\tau = 1, 32, 64 \mu s$ are superimposed on each other.

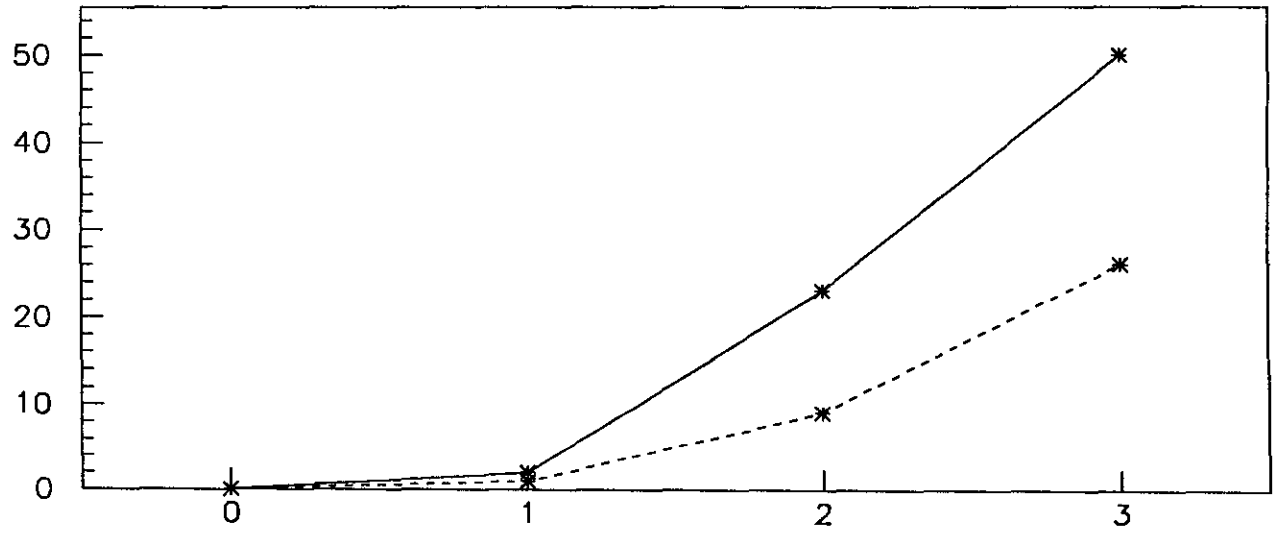

Figure 4: % of L2 trigger accepts which have to wait in order to enforce the minimal distance condition, for separation times of 32 and 64 μs.

5 Conclusions

In this note we have evaluated the buffer requirements for the L1 and L2 buffers, which are located on the TVC/AMU. In order to collect more than 99% of the hits, the L1 buffer needs to have 18/12 storage locations for an occupancy of 5% or 2% per crossing, respectively. The L2 buffer needs 8/4 storage locations for the two different occupancies, assuming the L2 latency to be uniformly distributed between 10 and 100 μs. Without too much additional burden a minimal distance of 30-60 μs between 2 L2 accept triggers can be imposed. The buffer needs are only increased for very large rates (L2=9 kHz) and long minimal distances of 64 μs.

References

[1] SDC detector notes, SDC-91-89, SDC Trigger preliminary conceptual design.

[2] P. Sinervo et al., *SSC Front End Simulations* Forth Worth Symposium 1990.

Buffer Occupancies for the SDC Straw Tube DCC System

A. Hölscher, G. Stairs and P. K. Sinervo
University of Toronto
Toronto, Canada

March 23, 1992

Abstract

In this note we present the results of simulations of the buffer occupancies at the front end card and at the crate interface card for the SDC straw tube system.

1 Introduction

The general architecture of this system as it is assumed in this simulation is as follows [1] (see figure 1). On each front end board reside 5 FMUX's, which are connected to 8 four channel TVC/AMU's (or TMC's) each, so that we have 160 straw channels per front end module. The FMUX's have one buffer for each TVC/AMU they are connected to. After each L2 accept trigger the data are first digitized on the TVC/AMU, which might take 1-3 μs, and then transmitted to the FMUX. This link is clocked at about 1 MHz and each clock cycle 4 bits are transferred to the FMUX. One hit from the TVC/AMU generates 3 bytes of data. On the FMUX a 4th byte for the geographical address is added. The 5 FMUX's residing on one front end board are read out one after the other by one FETX that transmits the data to a buffer on the CIC (crate interface card), which resides on a data acquisition crate outside the calorimeter. Each CIC reads 20 front end boards out in this scheme. For the whole straw system we forsee 32 CICs. We are interested in the occupancies of the buffers and bandwidth requirements at various stages.

The system is supposed to meet the following requirements :

- sustain an L2 accept rate of up to 10 kHz, with less than 10 % losses, and
- work for occupancies of up to 15%/trigger/channel.

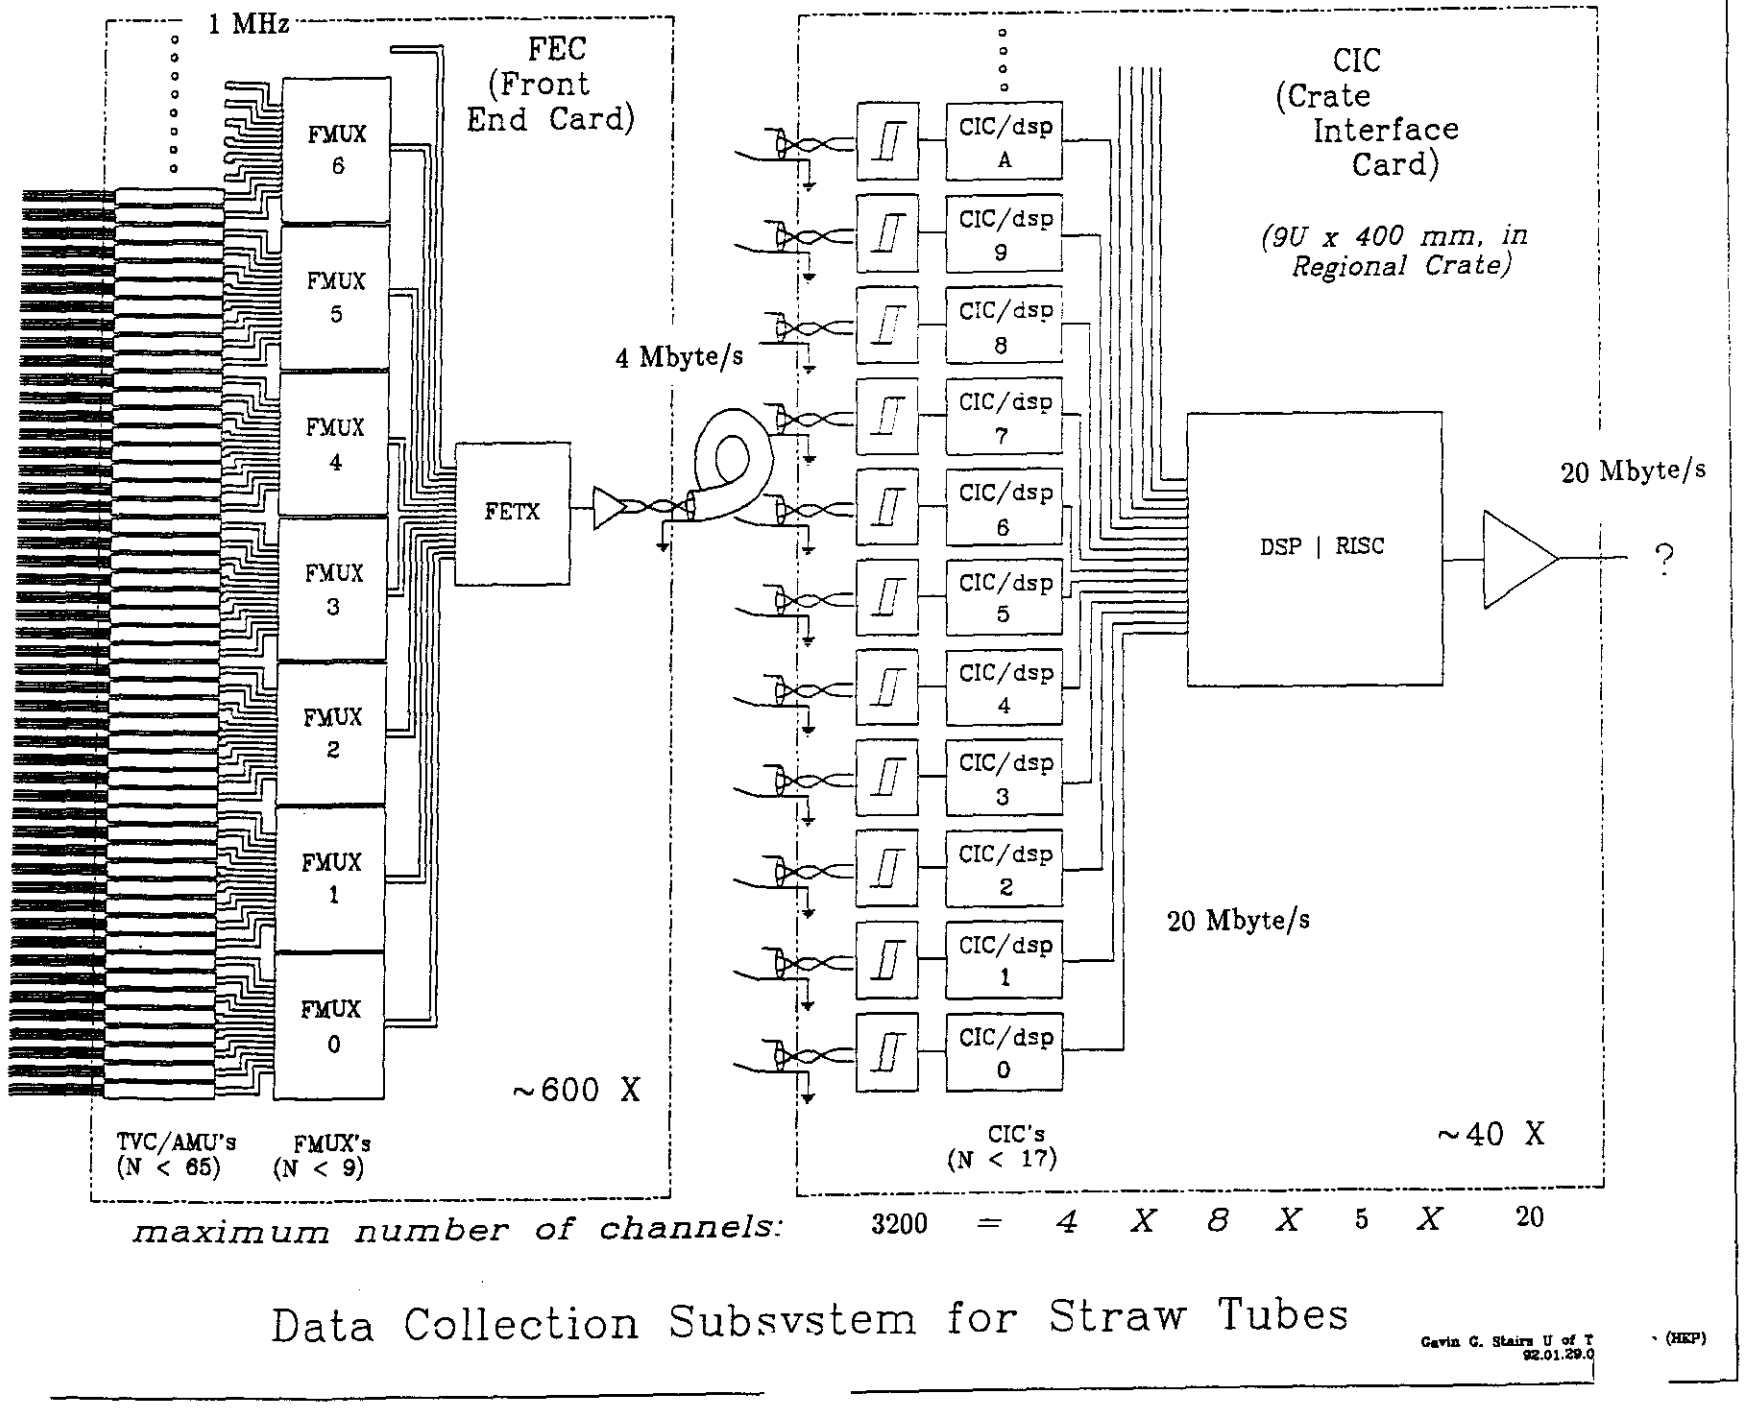

Figure 1: Architecture of the SDC Straw Tube DCC System.

Model simulations [2] show that the occupancies of the inner layers of the straw system are about 10% per trigger without any electronic noise. The above assumption contains some safety margin to incorporate this noise.
In this note the occupancies were generated, unless otherwise stated, with an L2 trigger accept rate of 9 kHz and an occupancy of 5% per crossing with 3 crossings per trigger being read out.

2 Bandwidths

From these requirements one can deduce the following average rates:

$$\text{datarate(TVC/AMU} \rightarrow \text{FMUX)} = 4 * 0.15 * 10\text{kHz} = 0.006 \text{ hits}/\mu\text{s}$$

Our protocol between the FE chip and the FMUX requires the FE to generate an "end-of-event" datum for every L2 trigger (see section 4). Including this additional effect, we obtain an average rate of 0.016 hits / μs.
Assuming the connection TVC/AMU - FMUX is clocked at 1 MHz and that the data are transferred in 1 nibble=4 bit, it will take time = N * 6μs (N number of hits on one TVC/AMU) until the data from one TVC/AMU arrives at the FMUX.
For the datarate from the FMUX to the CIC one gets the following values:

$$\text{datarate(FMUX} \rightarrow \text{CIC)} = 160 * 0.15 * 10\text{kHz} = 0.25 \text{ hits}/\mu\text{s}$$

Assuming one hit is 4 bytes long and allowing for a safety margin of at least a factor 4, the minimal required bandwidth for the bus to the TVC/AMU is :

$$\text{bandwidth(FMUX} \rightarrow \text{CIC)} > 4 \text{ Mbyte/s}$$

For each front end board there is a separate input buffer on the CIC. The bus which reads out these input buffers needs to carry the following rate:

$$\text{datarate(CIC} \rightarrow \text{CIC)} = 20 * 160 * 0.15 * 10\text{kHz} = 5.0 \text{ hits}/\mu\text{s}$$

Assuming again that one hit contains 4 byte and requiring a safety factor of at least 2 times the required bandwidth one obtains:

$$\text{bandwidth(CIC} \rightarrow \text{CIC)} > 40 \text{ Mbyte/s}$$

This is also the rate with which the data are transmitted to the DAQ.
The further one is away from the front end module, the smaller the safety factor needs

to be, since the fluctuations become smaller and smaller. This will be confirmed by the subsequent simulations. Generally, as long as one stays some reasonably safe factor away from the limiting bandwidth, the system performs very well.

3 Correlations

For each L2 trigger accept signal, we generate hits for each channel according to an average occupancy. In a previous note [2] we have shown, however, that the front end boards typically contain either no hit or at least 6 hits, due to the fact that a transversing particle sets a hit in all layers of one superlayer. To simulate this type of correlations, we assign to each event a different average occupancy. The distribution of this average occupancy is assumed to be exponential

$$f(occup) \propto exp(-occup/ < occup >)$$

with a mean occupancy of 5% per crossing. Three crossings are read out for each trigger.

4 FMUX occupancies

The FMUX's reside on the front end card and read out 8 TVC/AMU's (or TMC's) each. On each front end card there are 5 FMUX's, which are read out by one FETX. The FETX reads the FMUX's one after the other and transmits the data to a buffer on the CIC. SDC mandates that all data in the DAQ system to be event and channel ordered. The two issues concerned with this ordering are:

- data ordering by the geographical channel address, and
- event ordering by time and disentangling data from different events.

There are two possibilities to order the data from different events:

- Introduction of a latency between L2 trigger accepts, so that one event can be read out completely from the FMUX before the next one enters it.
- Let the TVC/AMU send an empty hit to indicate the end of an event for one trigger and the beginning of the next trigger. The FMUX would then continued to be read out until the empty hit for each L2 trigger.

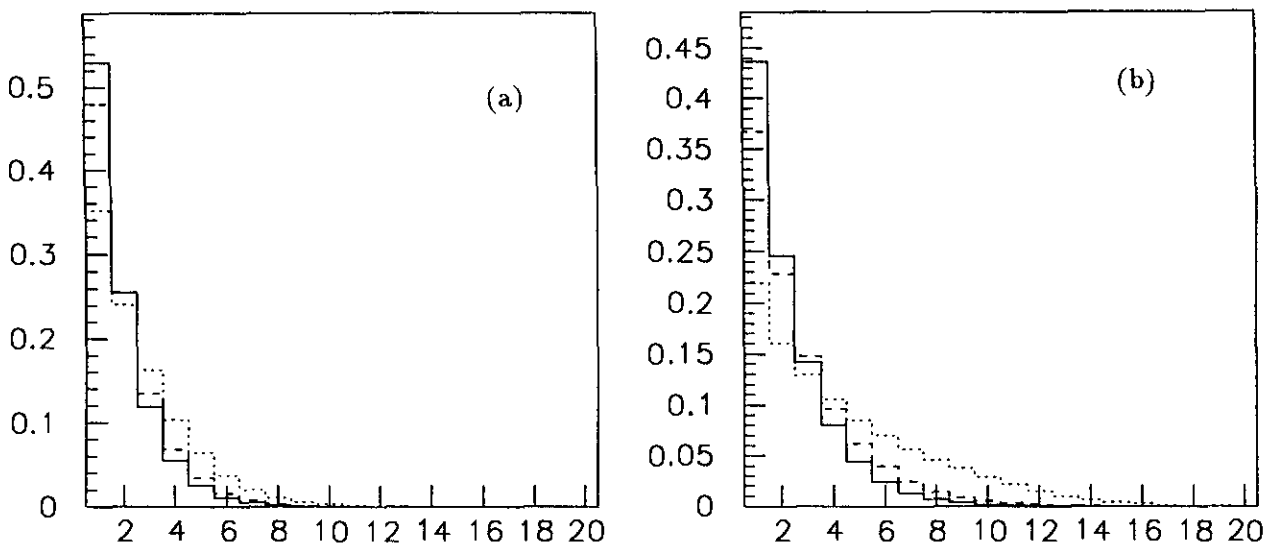

Figure 2: FMUX occupancy (in terms of hits) for one four-channel TVC/AMU without (a) or with (b) one empty hit for each event. The occupancy is shown for three different FMUX-FETX-CIC bandwidths of 2 (dotted), 4 (dashed) and 8 Mbyte/s (solid histogram).

The introduction of a minimal distance between L2 trigger accepts of up to 50 μs introduces a sustainable load on the the L2 buffer as was shown previously [3]. However, problems might occur for events that have a high occupancy in one FMUXR. This scheme also requires some additional logic in order to enforce this minimal L2 trigger accept distance.

In the second solution, the TVC/AMU adds an empty hit to the end of the data for each event. The FMUX is then read out until the occurence of an empty hit and one would naturally get the hits ordered after their event number. The disadvantages of this method are an increased TVC/AMU - FMUX bandwidth and an increased buffer requirement on the FMUX. Both disadvantages don't seem to be severe, however, and we use this method for the subsequent simulations.

The ordering by geographical address can be done by reading out the FMUX's in the same order for each trigger. This reduces the effective bandwidth somewhat, since one might have to wait for the data to arrive at a particular FMUX before the readout can continue. This degradation, however, is acceptable. Figure 2 shows the occupancy of an FMUX buffer for one four-channel TVC/AMU for three different FMUX-FETX-CIC bandwidths of 2, 4 and 8 Mbyte/s and for the case where one empty hit at the end of each event is not included (a) or is included (b). Figure 3 shows the required FMUX buffer length in order to contain 99.9% of the hits for the various cases. For a FMUX-CIC bandwidth of 4 Mbyte/s, the hits

are well contained in a buffer of length 15.

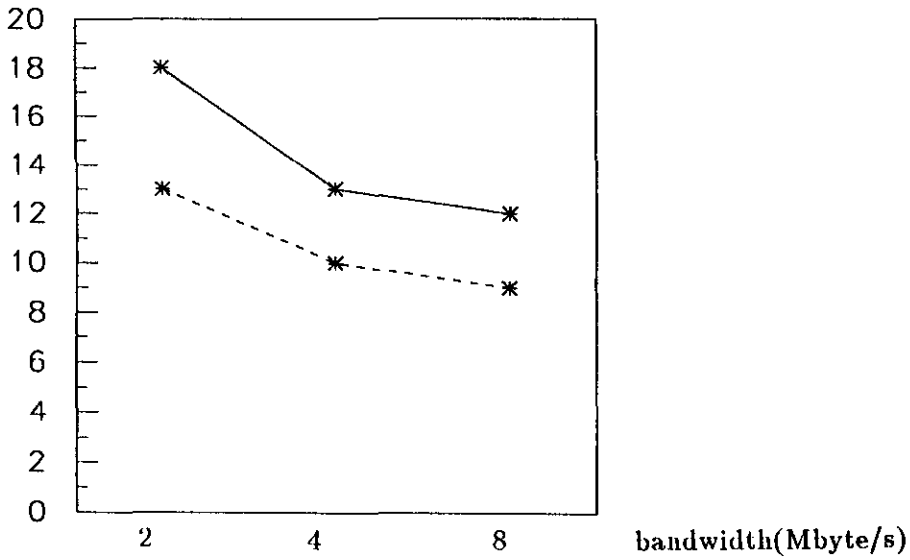

Figure 3: max FMUX buffer occupancy (99.9%) as a function of the FMUX-CIC bandwidth.

4.1 Recovery fom buffer overflow

Special care has to be taken when the FMUX buffer overflows, in order not to loose the control information provided by the empty hit. When the FMUX buffer is full, we propose to stop the data transmission from the front end chip to the FMUX. This will cause first the L2 buffer and then the L1 buffer on the chip to be filled up and eventually cause the chip to disable its input when no space is left for hits in the L1 buffer. As soon as the FMUX has space again, the data transmission from the TVC/AMU (or TMC) to the FMUX buffer can resume.

5 CIC occupancies

The data are transmitted from the front end boards to buffers on the crate interface card (CIC), which is located outside the calorimeter. Each CIC has one data input buffer for each of the 20 front end board it is connected to. It then puts the data belonging to one trigger together in one output buffer, from where the data are transmitted to the DAQ. Apart from collecting the data from the front end board, the CIC card has some control

and monitor functions. The input buffer needs on the CIC are of course dependent on the bandwidth with which theses buffers are emptied. Figure 4 shows the occupancy distribution of the CIC input buffers for a CIC bandwidth of 24 Mbyte/s (again assuming that each hit consists of 4 bytes). In this figure, we compare the case, where (a) the event data are not geographically ordered, but the input buffers on the CIC are read out in the order that they arrive in the CIC, (b) second where we assume a fixed occupancy for each event, neglecting correlations as they were assumed in section 3 and (c) assume the type of correlations presented in section 3 and reading out the CIC front end buffers always in the same order. The percentage of hits lost with a buffer length of 1000 hits per front end card is 0.6%, 0.03% and 1.6%, for the three cases, respectively. The occupancy distributions don't show big differences between the three cases. Reading the front end boards out in the order of their availability (a) gives a slightly larger throughput than always reading them out in the same order, whereas neglecting correlations within one event leads to smaller losses (b) compared to the nominal case where we get ordered event and were we put in some correlations. Although these differences are expected the difference between the distributions is much smaller than one probably would have anticipated.

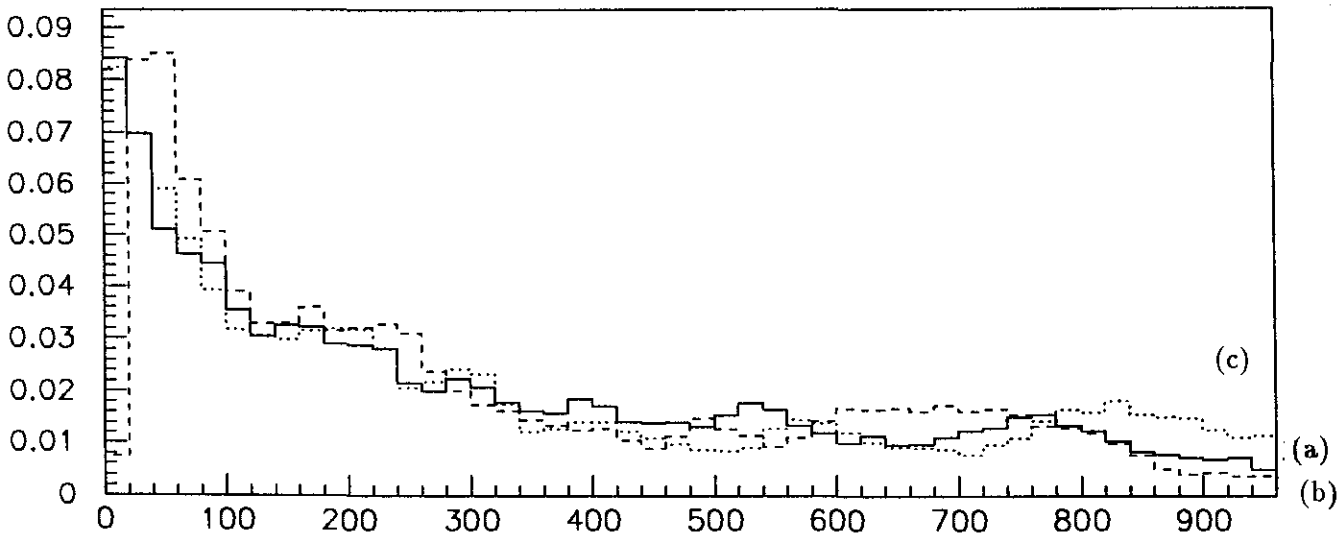

Figure 4: CIC input buffer occupancy (in terms of hits) for one front end board. Case (a) solid, (b) dashed and (c) dotted histogram).

Figure 5 shows the CIC buffer occupancy distribution dependence on the CIC bandwidth. The occupancies are shown for a bandwidth of 24, 36 and 48 Mbyte/s. Also shown are the maximal occupancy for different bandwidths. For a CIC bandwidth exceeding 40 Mbyte/s

the hits are well contained in a buffer length of 400 hits.

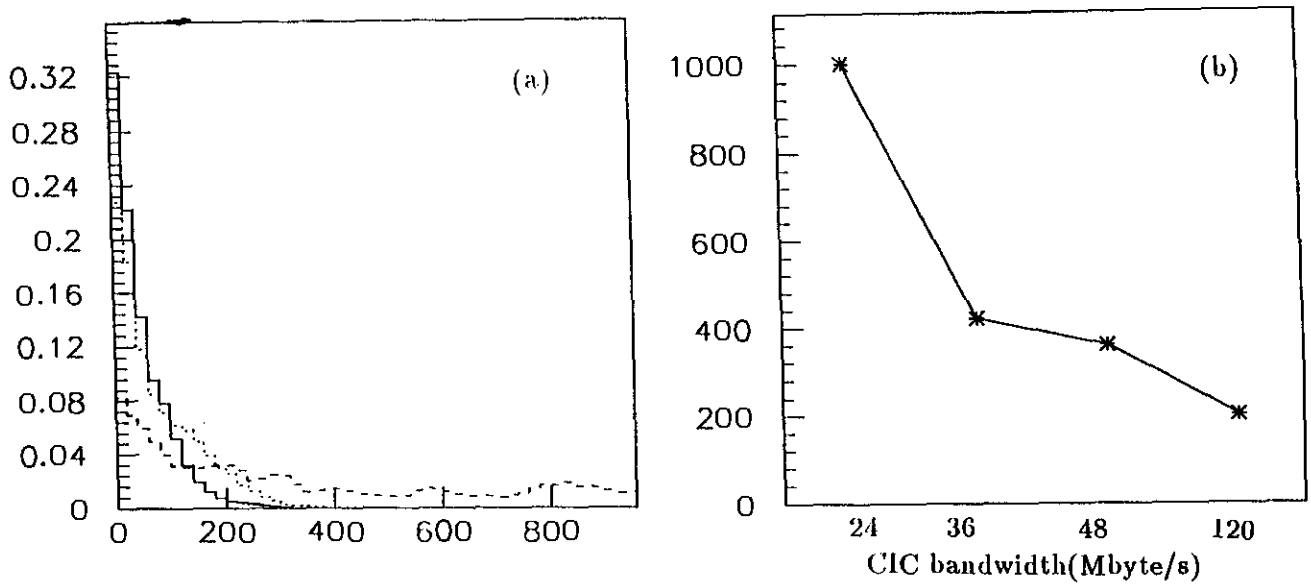

Figure 5: CIC input buffer occupancy for one front end board, for three different bandwidths on the CIC: 24 (dashed), 36 (dotted) and 48 Mbyte/s (solid histogram). The histogram in (b) shows the maximal (99.9%) occupancy.

5.1 Influence of CIC delay times

Since the CIC has also monitor and control functions, we are also interested in the buffer requirements assuming that the hits have to stay a minimal time τ in the input buffer. This would give the a DSP (see figure 1) some time for monitor and control functions. Figure 6 shows the CIC input buffer occupancy distribution for an CIC bandwidth of 48 Mbyte/s for 3 different minimal times $\tau = 0$, 500 and 1000 μs in the CIC buffer. Even for a delay time of 1000 μs, the additional burden on the CIC buffers is small. In this case one would like to have a somewhat larger storage of 500 hits.

5.2 Load balanced Front End Readout

Simulations of the occupancies of front end modules [2] show that the occupancies have a strong dependence on the superlayer number in the sense that the occupancies of the outer layers are much smaller than the inner layer. Since the CICs are located just outside of the calorimeter in different ϕ positions, geometrical arguments and arguments of load balance suggest, that each CIC reads out front end modules of different layers. This lowers the

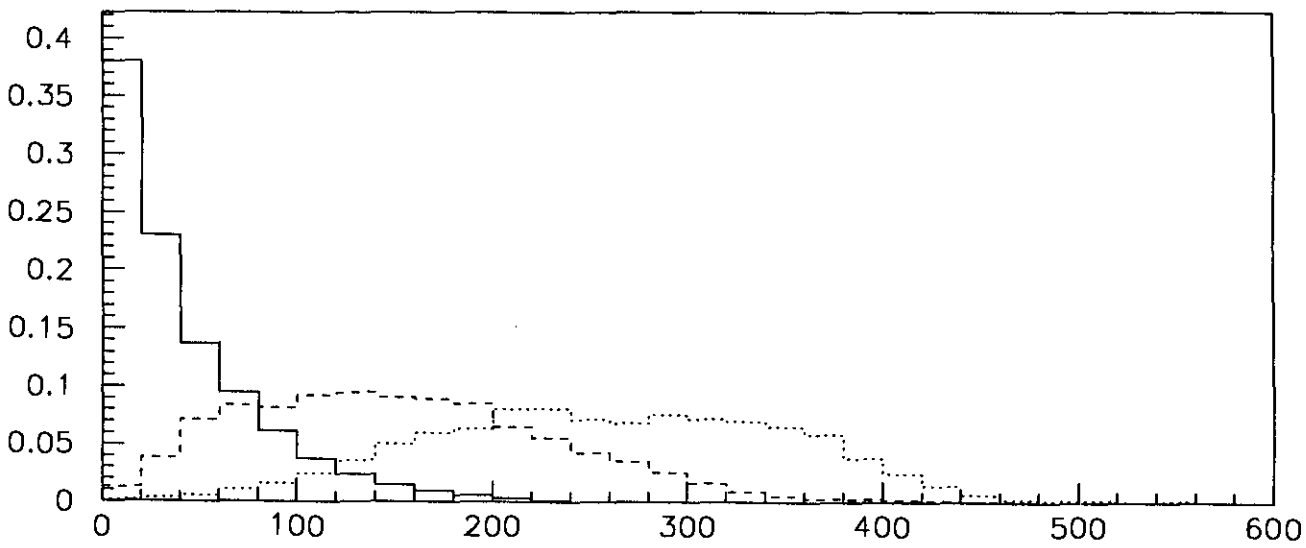

Figure 6: CIC buffer occupancy for one front end board for a CIC bandwidth of 48 Mbyte/s for three different CIC minimal stay times of 0 (solid), 500 (dashed) and 1000 μs (dotted histogram).

bandwidth requirements of one CIC, since the occupancies of the outer layers are much smaller. We assume that the 20 front end boards of one CIC contain 4 front end module of each superlayer. We assume the following average occupancies per trigger for the 5 superlayers of 15%, 9%, 6%, 4.5% and 4.5%, which is still well above the results of the simulation of 10%, 4.7%, 3.3%, 2.5% and 2.5% [2]. Figure 7 shows the CIC buffer ocupancy distributions for different CIC bandwidths of 24, 16 and 12 Mbyte/s. Here no difference is made between CIC reading out superlayers with higher or lower occupancies. For this load balanced front end readout, the required CIC bandwidth is only 20 Mbyte/s.

5.3 CIC output buffer and bandwidth to DAQ

The data from different front end boards are now assembled and added to one output buffer, where they are to be shipped off to the data aquisition system. The buffer size needed here, of course, depends on the specification of this connection, which at this stage is not defined. Assuming the load balanced case (section 5.2), a CIC - DAQ bandwidth of 24 Mbyte/s and no additional delays, one gets the output buffer distribution of figure 8, showing that in this case the buffer can contain as many as 1000 hits. This requirement will certainly grow.

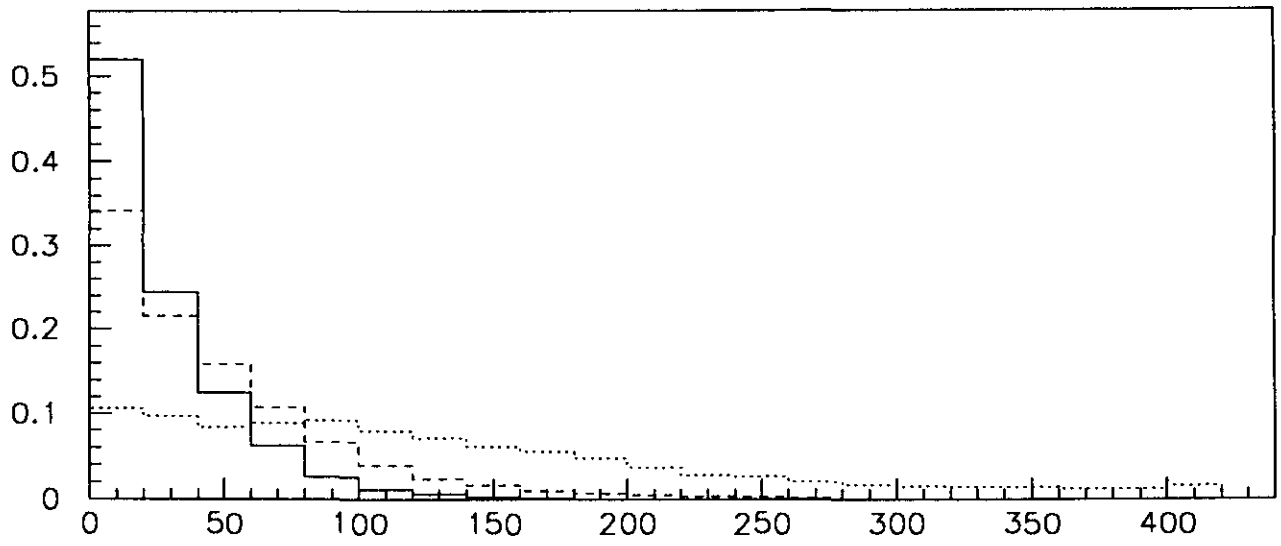

Figure 7: CIC buffer occupancy for one front end board for a CIC bandwidth of 24 (solid), 16 (dashed) and 12 Mbyte/s (dotted histogram), with the CIC reading out 4 front end boards of each superlayer.

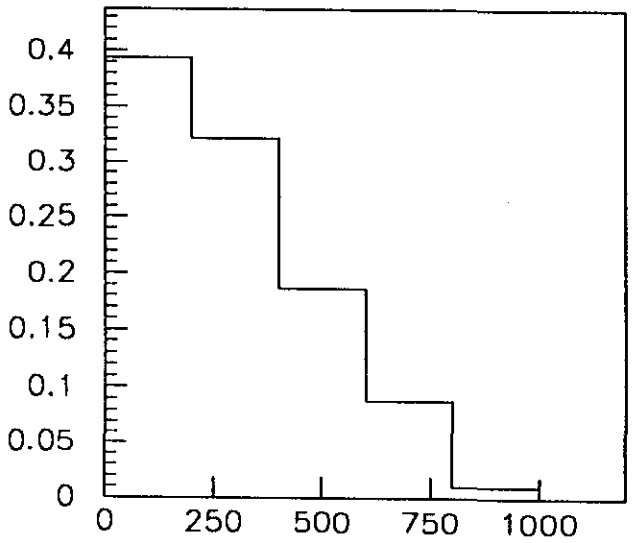

Figure 8: CIC output buffer occupancy for the load balanced front end readout (section 5.2) and a CIC-DAQ bandwidth of 24 Mbyte/s and no additional delays.

6 Conclusion

In this note we discussed the buffer needs and bandwidth requirements of the front end boards and of the crate interface card. Adding the conclusion from a previous study [3] we arrive at a set of minimal parameters for a front end system, assuming the requirements that:

- the system has to sustain an L1 accept rate of 100 kHz
- the system has to sustain an L2 accept rate of 10 kHz
- allow the straw tubes to have an occupancy of 15 % per trigger
- and allow for an L2 latency time of 10 - 100 μs uniformly distributed

The system would have the following parameters:

- L1 buffer length >= 18 hits
- L2 buffer length >= 8 hits
- TVC/AMU (or TMC) - FMUX link clocked by at least 1 MHz with 4 bits transferred/clock
- FMUX buffer length per 4 channel TVC/AMU of at least 15 hits. Assuming each hit contains 4 bytes, one needs 8*4*15 =480 bytes storage on each FMUX.
- the bandwidth FMUX - FETX - CIC exceeding 4 Mbyte/s.
- CIC input buffer length per front end board exceeding 1.6 kbyte. Assuming that each CIC reads out 20 front end boards one would need at least 32 kbyte on the CIC as an input buffer.
- requiring the data to stay some minimal time of 500 μs to 1000 μs in these CIC input buffers to allow for some data monitoring operations, imposes only a small additional burden on the buffer.
- bandwidth on CIC with which the 20 front end board buffers are read out exceeding 40 Mbytes. If one CIC reads out front end boards roughly equally distributed over the superlayers (load balanced), a lower bandwidth exceeding 20 Mbyte/s is sufficient.

- the CIC output buffer for the load balanced case with a CIC-DAQ bandwidth of 24 Mbyte/s is required to be at least 1000 hits. This requirement will certainly grow, depending on the specifications of the CIC-DAQ link.

References

[1] G.Stairs et al., *Front end Collector Bus Design Description*, SDC-92-242.

[2] A. Hölscher, P. K. Sinervo and G.Stairs, Univ. of Toronto,
Simulation of Front End Board Occupancies for the SDC Straw Tube Tracker

[3] A. Hölscher, P. K. Sinervo and G.Stairs, Univ. of Toronto,
Front End Buffer Requirements for the SDC Straw Tube Tracker