A study of ROD compression schemes for the TRT

Mogens Dam
Niels Bohr Institute
Copenhagen, Denmark

31 August 2001

Abstract

Zero suppression of the TRT data is foreseen already at the ROD level. Valid straws will be identified by applying a gate of tunable position and width to the low threshold data bits. Various schemes have been proposed for the encoding of data from valid straws. The schemes are answering the requirements posed by the different running modes of LHC/ATLAS: low luminosity, low luminosity b-physics programme, and high luminosity. This note presents a Monte Carlo study of the validity gate approach and of the proposed encoding schemes. Results are presented on the obtainable efficiency and background suppression as function of the validity gate position and width. The different encoding schemes are investigated and estimates of the output data volumes are presented for different luminosities.

1 Introduction

The TRT comprises about 420,000 read out channels. For every First Level Trigger Accept (L1A) each of these channels (straws) produces 27 bits of information (1 high threshold bit + 8 low threshold bits for 3 consecutive time-slices). Thus, for an assumed L1A rate of 75 kHz the total data flow out of the TRT is about 106 GByte/sec. From the front-end electronics the data is read out via 256 Read Out Drivers (RODs) and passed along to 256 Read Out Buffers (ROBs), where it is residing during Second Level Trigger processing. To limit the data flow into the ROBs it is foreseen to perform zero suppression already at the ROD level.

In this note, various possible zero suppression schemes are studied based on a detailed simulation of the TRT.

2 Programs, Data, and Methods

The study was performed using the XTRTDIG code run inside the ATLSIM framework. Two main data types were used:

Minimum Bias Events: A set of "old" data files with GEANT hits were read into ATLSIM. Thus, for this data, ATLSIM was only responsible for the digitization part;

Single Muon Events: Generated on the fly by ATLSIM using the single particle generator;

Low threshold discriminator level	200 eV
High threshold discriminator level	5000 eV
Jitter on drift time: box of width	± 5 ns
Minimum discriminator width	6 ns
Discriminator dead time	6 ns
Discriminator settling time	1.5 ns
Fraction of straws with operational electronics	0.97

Table 1: Some important parameter of the TRT simulation.

Because the minimum bias events were produced with an earlier version of the TRT geometry, this geometry was used throughout the study. Some important parameters of the TRT simulation are summarized in Table 1.

A special data file was written from XTRTDIG for this study. Its format was originally defined by Fred Luehring. It was extended, however, with one extra word per straw to store Monte Carlo *truth* information on which bunch-crossing(s) contribute to each straw. The format is summarized in Table 2. The output file contains information on all straws which have at least one low threshold bit set within the first three time-slices after the in-time bunch-crossing (BX0).

For the study of minimum bias events, overlapping events were simulated in four bunch-crossings before and two bunch-crossings after BX0. That is, a total of seven bunch-crossings were simulated. This was done by using the ALTSIM framework to read an average of 23 events per BX from the data file (Poisson distribution). For other luminosities the mean number of events to simulate was scaled correspondingly.

3 Validity gate

The first step of the zero-suppression scheme will apply a gate of tunable position and width to the data received from the front-end electronics. Straws with at least one low threshold bit set within this gate are considered "valid". Our first task here is to optimize this gate. The study was carried out on clean signal-like tracks. For this, the single muon sample of $p_{\mathrm{T}} = 10$ GeV was used.

Figure 1 shows various distributions for the single muon sample. As can be seen, 8.3% of the signals occur as two "trains" so that two rising and two falling edges are observed. This percentage depends critically on the electronics shaping function. In the current version of the simulation, a constant (i.e. independent of the charge on the pre-amplifier input) shaping function with a peaking time of 8 ns and a base width of 17 ns was used as shown in Figure 2. It was found in a dedicated study, that the splitting probability decreased from 8.3% to 1.4% for a shaping function which was "stretched" by a factor 1.5 in time relative to the standard one.

Figure 3 shows the validity gate efficiency as function of the gate starting position for different gate lengths. Also shown is the efficiency as function of gate length for optimally positioned gates. A signal efficiency of 97.2% (95.8%) is observed for a gate of length 4 (3) bits equivalent to 12.5 ns (9.4 ns). Also shown is the efficiency in the case of the "stretched" shaping function. Here the lower signal splitting probability results in somewhat higher

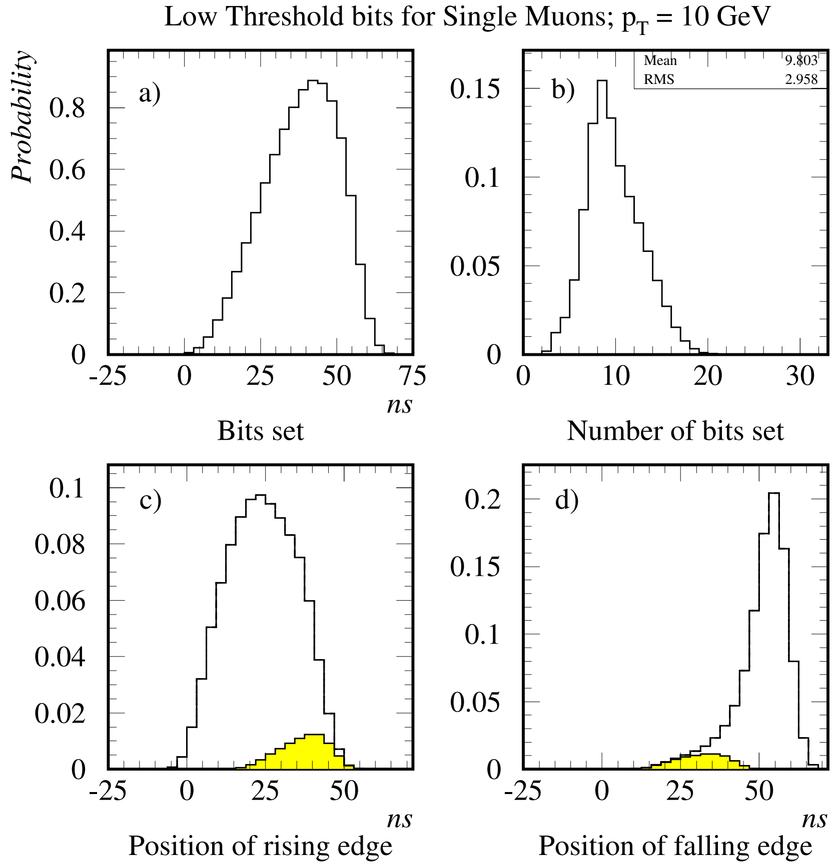

Figure 1: Distributions of low threshold bits for single muons: a) set bits; b) number of bits set; c) rising edge; d) falling edge. For c) and d) the hatched area corresponds to the second (first) rising (falling) edge in the 8.3% of cases where the signal was found to be split into two "trains".

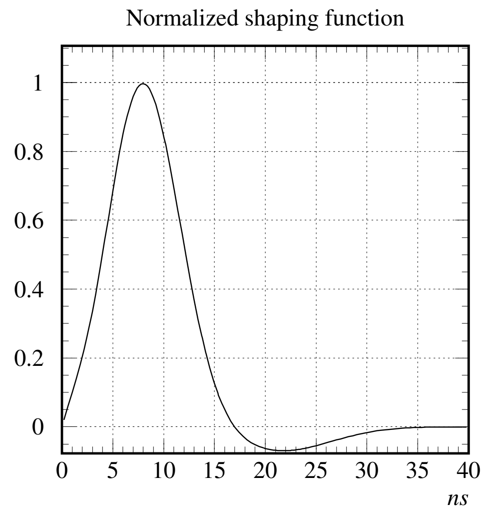

Figure 2: The constant shaping function used in the simulation.

Word	Bits	Description
1	0–9	Layer number: 1–297 ($-z$), 301–597 ($+z$)
		– 1 – 9 short barrel
		– 10 – 75 long barrel (21 and 45 no-existent)
		– 76 – 235 short disks
		– 236 – 299 long disks
	10–19	ϕ zone: 1 to 192
	20–27	Low discriminator threshold for time-slice -1 (-25 to 0 ns)
	28	High discriminator threshold for time-slice -1
2	0–7	Low discriminator threshold for time-slice 0 (0 to 25 ns)
	8	High discriminator threshold for time-slice 0
	10–17	Low discriminator threshold for time-slice 1 (25 to 50 ns)
	18	High discriminator threshold for time-slice 1
	20–27	Low discriminator threshold for time-slice 2 (50 to 75 ns)
	28	High discriminator threshold for time-slice 2
3	0–6	Monte Carlo truth information:
		– bit 0 set: straw hit by track from BX-4
		– bit 1 set: straw hit by track from BX-3
		– bit 2 set: straw hit by track from BX-2
		– bit 3 set: straw hit by track from BX-1
		– bit 4 set: straw hit by track from BX+0
		– bit 5 set: straw hit by track from BX+1
		– bit 6 set: straw hit by track from BX+2

Table 2: The format of the output file from the simulation. Word three was added to the original format specifically for this study. *Actually the format could be optimized to fit into two words. That would require to alter the coding of the two first words*

.

efficiencies of 98.6% (98.3%) for the same gate length.

Figure 4 illustrates how the validity gate works on low luminosity minimum bias events. The gate starting position was chosen according to the above study for an optimal efficiency for single muons. Shown is the efficiency for keeping the signal and the ratio of background-to-signal after applying the gate. Here signal (background) is defined as straws having at least one hit (having no hits) from the in-time bunch-crossing. It is interesting to observe, that if one had chosen to optimize the gate starting position on the minimum bias events themselves, the gate would in general have been positioned one bit (3.1 ns) later with a resulting increase in efficiency for minimum bias tracks of about 3%.

To illustrate the suppression of background, Figure 5 shows the distribution of set low threshold bits for the four most important bunch crossings. Without applying a validity gate (apart from the requirement that there be at least one bit be set in the 0–75 ns interval) the background over signal ratio starts out at about 3.3. By applying a more restrictive gate, this is brought down as illustrated in Figure 4. For a gate of length 4 (3) a background to signal ratio of 0.66 (0.57) is observed.

The contributions from different bunch-crossings to the total occupancy averaged over all

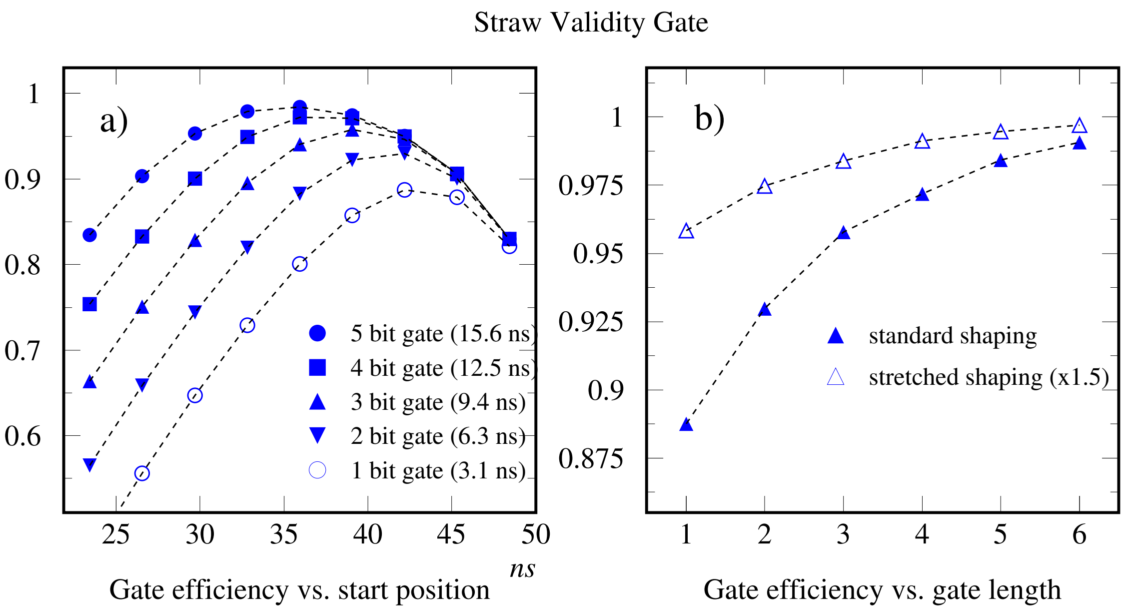

Figure 3: Validity gate efficiencies: a) as function of gate starting position for five different gate lengths; b) as function of gate length – gate placed at the optimal position.

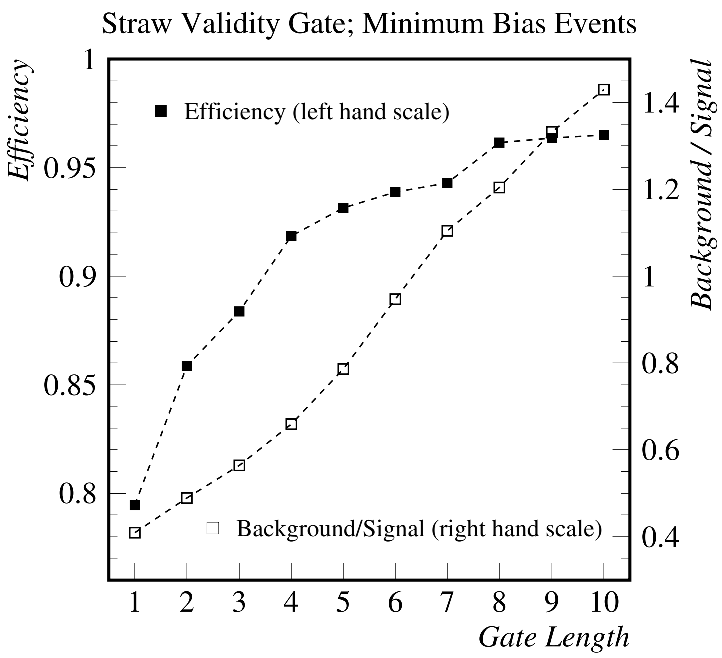

Figure 4: Efficiency and background-to-signal ratio as function of the validity gate length for minimum bias events. Signal (background) is defined as straws with at least one hit (no hits) from the in-time bunch crossing. The gate position has been optimized for single muon tracks.

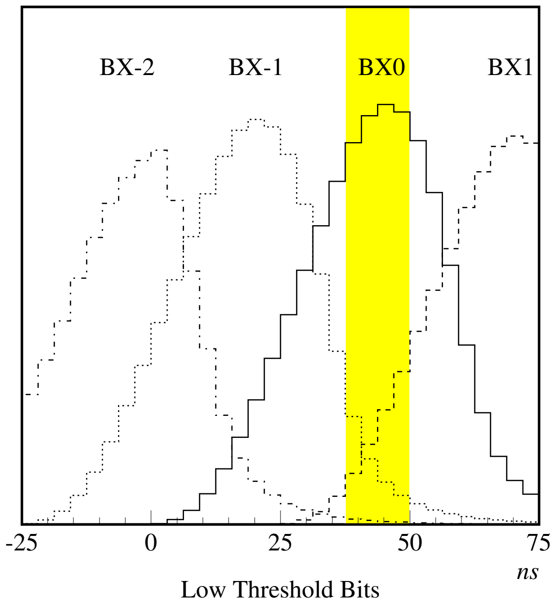

Figure 5: The distribution of set low threshold bits for the four most important bunch crossings. The hatched area indicates the 12.5 ns gate.

Beam-crossing	-4	-3	-2	-1	0	$+1$	$+2$	All
No gate	0.1	0.8	15.7	19.0	19.0	19.0	6.5	53.9
4 bit gate (12.5 ns)	0	0.03	0.3	3.9	17.5	6.3	0.01	25.2
3 bit gate (9.4 ns)	0	0.03	0.3	3.9	16.8	4.7	0	23.4

Table 3: Straw hit occupancy (in percent) from different beam-crossings at the design luminosity of 10^{34} cm^{-2}s^{-1} for different validity gates. The straw hit occupancies from each beam-crossing are estimated as independent contributions, whereas the total occupancy includes the correlations between the different beam-crossings.

straws are shown in Table 3. The numbers for the 12.5 ns gate should be directly comparable with the numbers in Table 3-16 of the TDR [1]. Whereas the individual contributions are roughly of the same size in the two studies, this study finds a total occupancy about 10% above that of the TDR study. Between the two studies there seems to be some discrepancy in the amount of correlation between different bunch contributions. The contributions of all of the bunches and of the just in-time bunch are shown in Figure 6 as function of the layer number.

4 Multiple edges

At high luminosity a compression of the TRT data is necessary already at the ROD level. In the proposed scheme, it is essentially suggested to pass information on the position of the leading edge(s) of the low threshold discriminator output, instead of passing the full

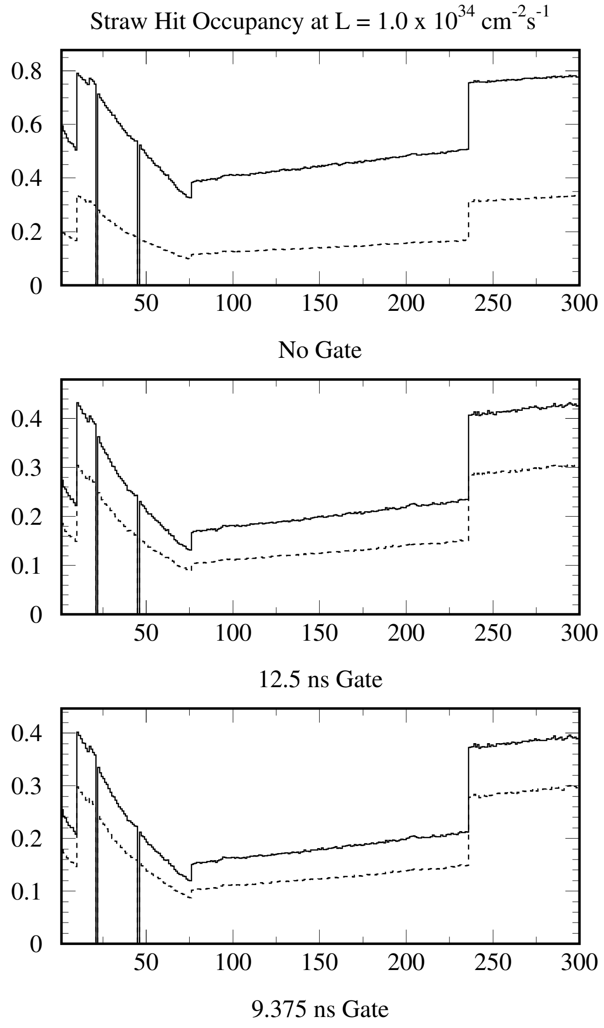

Figure 6: Straw hit occupancy vs. layer number at the design luminosity of 10^{34} cm^{-2}s^{-1}. Full line: data from BX$-$4 to BX+2; dashed line: data from BX0 only.

bit stream itself. The position of the leading edge provides drift time information and is therefore important for pattern recognition purposes. As we have already seen, there is a finite probability of observing more then one leading edge. It is therefore useful to consider the potential gain from the recording of multiple edges. Due to the possibility of obtaining particle identification through the so-called Time-over-Threshold (ToT) measurement, it is also of interest to record the position of the trailing edge(s) – at least at low luminosity.

Table 4 summarizes the ROD compression scheme for high luminosity conditions proposed at the recent Istanbul meeting [2]. It is foreseen to provide for two leading edges – one in each of the two first time-slices. In case of two edges within the same time-slice, the first will be dropped and only the last will be encoded. Table 5 summarizes the number of leading edges observed during the first two time-slices at design luminosity. Only a marginal difference is observed between, on the one hand, the situation where all edges are counted, and on the other hand, the situation where only one edge per time-slice is counted. Since the latter

Straw empty	`0_0`	
Straw not empty, not valid	`0_1_H`	
Straw valid	`1_0_H_T_E1_E2_{DTM}_{DTM}`	
	- `H`	high threshold flag
	- `T`	trailing edge flag
	- `E1` (`E2`)	edge flag for first (second) time-slice
	- `DTM`	3 bit edge position (optional; if `En` flag)

Table 4: "Istanbul" proposal for high luminosity ROD zero suppression scheme.

	All edges				Max 1 edge per TS		
Number of edges	0	1	2	> 2	0	1	2
All BXs	5.9	85.7	9.3	0.1	5.9	85.4	8.7
BX0 only	0.1	93.0	6.8	0.1	0.1	93.6	6.3

Table 5: Number of leading edges observed during the first two time slices at design luminosity. At the left, all edges are counted. At the right, a maximum of one leading edge per time-slice is counted. All straws satisfy the 4 bit (12.5 ns) validity gate.

scenario allows for a denser packing of the data, it is therefore assumed in the following that leading edges are counted as a maximum of one per time-slice.

Like the result just presented, the rest of this section is based on a study of minimum bias events, where hits from the in-time bunch-crossing are defined as signal and those from out-of-time crossings as background. For this study it would probably have been more to the point to consider signal like tracks (single muons, say) in the dense environment of the underlying minimum bias events. This is technically slightly more complicated, and so the "easy" solution was chosen. Figure 7 shows how various low threshold distribution differ between the single muon tracks and tracks from minimum bias events. Generally, the single muon tracks are about one bit earlier (both leading and trailing edges) than the minimum bias tracks. This will cause the single muon tracks to be more severely shadowed by the preceding bunch-crossing and less by the following crossing than what is the case for minimum bias tracks. Thus, for high momentum signal tracks the probability of loosing the leading (trailing) edge will be somewhat larger (smaller) than what is indicated in the current study.

Table 6 presents the frequency by which 0, 1, and 2 leading edges occur for signal straws; i.e. straws having at least one hit from the in-time bunch-crossing. The bottom row of the table shows, how the fraction of straws with exactly one leading edge goes down relatively moderately from 92.6% at low luminosity to 87.9% at high luminosity. This decrease is compensated by an increase in the occurrence of 0 (2) edges from 0.4% (7.0%) to 2.9% (9.2%). The effect can be understood from the breakdown in the rows above, where the four possible bunch-crossing situations are considered:

a) The top row shows the most favourable situation, where there are no hits from neither the previous (-1) nor the next ($+1$) crossings. Fortunately this is the most common situation observed in 94.1% (67.6%) of all signal straws at low (high) luminosity. In this

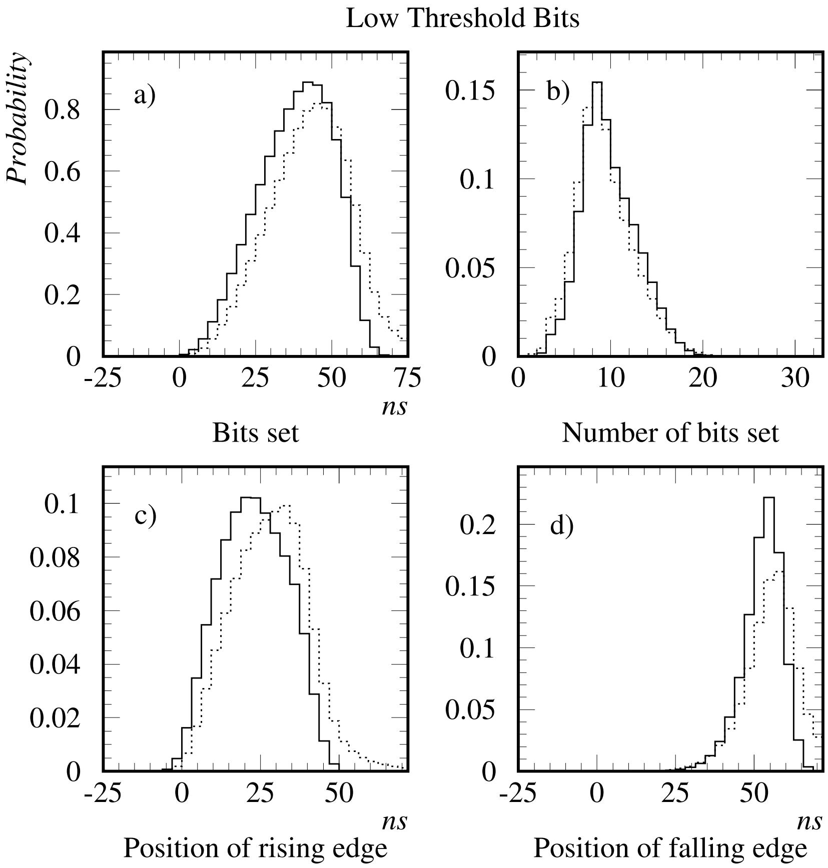

Figure 7: Distributions of low threshold bits for single muons (full line) and for tracks from minimum bias events (dotted line): a) set bits; b) number of bits set; c) (first) leading edge; d) (last) trailing edge.

Leading Edges		10^{33} cm^{-2}s^{-1}				10^{34} cm^{-2}s^{-1}			
Case	BX contrib.	fraction	0	1	2	fraction	0	1	2
a)	0	94.1	0.1	93.4	6.5	67.6	0.1	93.6	6.3
b)	-1, 0	2.3	12.5	64.0	23.6	14.2	15.7	63.5	20.8
c)	0, $+1$	3.5	0.0	91.9	8.0	14.4	0.1	92.2	7.7
d)	-1, 0, $+1$	0.1	13.7	62.0	24.3	3.8	15.6	61.9	22.5
Total signal		100.0	0.4	92.6	7.0	100.0	2.9	87.9	9.2

Table 6: Number of leading edges within the first two time-slices for four different bunch-crossing situations: a) in-time crossing only; b) in-time and previous crossings; c) in-time and following crossings; and d) all three crossings. The last row, which is the sum of the four contributions, gives the total for all signal straws; i.e. straws with at least one hit from the in-time crossing. All straws satisfy the 4 bit (12.5 ns) validity gate.

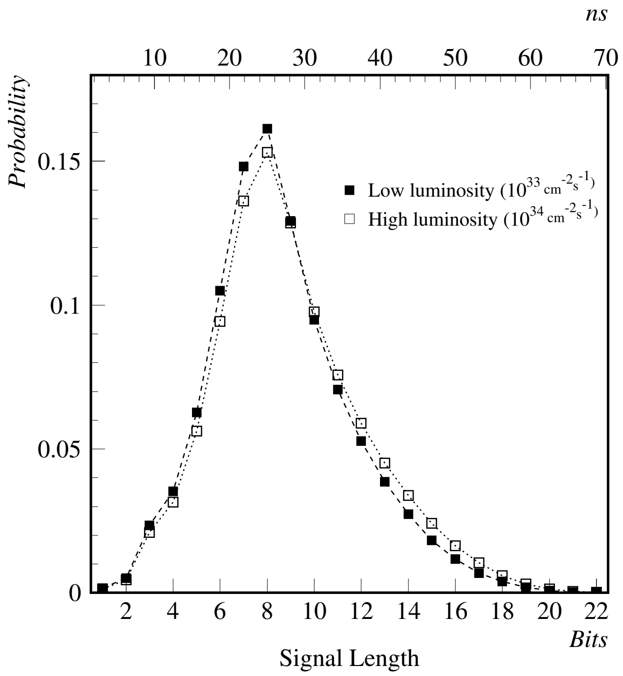

Figure 8: Signal length distribution for straws with hits from the in-time bunch-crossing only.

case, 6.5% of straws have two leading edges. This corresponds roughly to the situation for single muons (where it was 8.3%) as discussed in connection with Figure 1. We therefore conclude that generally the second edge is an artifact caused by the short shaping time, and that the first edge is the one of interest for pattern recognition.

b) Here hits are observed also from the previous crossing. This naturally leads to an increase both in the occurrence of zero edges (the two signals merge to one, which starts before t_0) and two edges (the two non-merged signals both start after t_0). In this case, the second edge generally carries the correct drift time information.

c) Here hits are observed also from the next crossing. A small increase in the two edge occurrence is observed. The drift time information is in the first edge.

d) Here there are hits from both neighbour crossings. The situation looks similar to b) and in general the second edge carries the correct drift time information.

In general only small differences are observed in the individual bunch-crossing situations between the low and high lumonisity cases. This is also illustrated in Figure 8, which compares the signal length for straws with hits from in-time bunch-crossings only for low and high luminosity. The two distributions are rather similar; compared to an average signal length of about 9 bits, the increase going from low to high luminosity is only about 0.5. The major difference between the two luminosities is the different weighing of the four bunch-crossing situations. Following the argument from above, according to which the first edge is expected to carry the good the drift time information in bunch-crossing situations a) and c), we can now conclude that the loss of correct drift time information would be marginal if we would

Luminosity	Two edges	First edge
10^{33}	99.6%	99.4%
10^{34}	97.1%	96.2%

Table 7: Estimated fraction of straws where the correct leading edge is recorded for two situations: i) two leading edges are recorded, ii) only the first leading edge is recorded. Numbers have beed derived by considering the information in Table 6 as described in the text.

Trailing Edges		10^{33} cm^{-2}s^{-1}				10^{34} cm^{-2}s^{-1}			
Case	BX contrib.	fraction	0	1	2	fraction	0	1	2
a)	0	94.1	2.5	90.8	6.7	67.6	2.8	90.8	6.4
b)	−1, 0	2.3	1.4	70.5	28.1	14.2	1.9	71.6	26.5
c)	0, +1	3.5	33.7	60.7	5.6	14.4	37.7	57.3	4.9
d)	−1, 0, +1	0.1	23.0	60.1	16.9	3.8	26.0	58.2	15.8
Total signal		100.0	2.9	87.9	9.2	100.0	8.6	82.0	9.4

Table 8: Number of trailing edges within the first two time-slices for four different bunch-crossing situations: a) in-time crossing only; b) in-time and previous crossings; c) in-time and following crossings; and d) all three crossings. The last row, which is the sum of the four contributions, gives the total for all signal straws; i.e. straws with at least one hit from the in-time crossing. All straws satisfy the 4 bit (12.5 ns) validity gate.

choose instead of recording two edges to record only the first observed edge. An estimate of the size of the effect is presented in Table 7.

Table 8 lists the occurrence of trailing edges in a way similar to what was done above for leading edges. At low (high) luminosity 97.1% (92.4%) of straws have one or more trailing edges. Of the 9.2% (9.4%) of straws with two trailing edges most are coming from cases a) and b), where the last edge is expected to be the good one. By choosing always the last edge we would therefore expect to have the correct falling edge information in at least 94% (80%) of all straws – probably quite a bit more.

As seen from Figure 7, the trailing edge distribution spans time slice 2 and 3 and has a total width of 10–12 bits. Four bits would therefore be needed to record the position of the trailing edge.

The studies presented in this section suggest the following conclusions:

- For the drift time measurement it seems sufficient to record only the first leading edge after BX0. The potential gain in the number of straws with correct drift time information from the currently foreseen scheme, where two edges are recorded, seems to be less than 1%.

  4 bits would be needed to record the postion of a single leading edge.

- The Time-over-Threshold measurement could be accomplished from the recording of the position of the trailing edge. In case of multiple edges, the last edge should be chosen as the good one.

4 bits would be needed to record the postion of a single trailing edge.

5 Data Compressions Schemes

The choice of zero suppression scheme naturally depends on the detector occupancy. At low luminosity, where the occupancy is moderate, it is possible to retain the full stream of low and high threshold bits for all hit straws. This would allow for a straightforward extraction of the ToT information, which is mostly of interest for the low luminosity b-physics programme. However, specifically for the b-physics programme it is necessary to perform a full scan of the TRT already at the Second Level Trigger level. In order to minimize the bandwidth requirements on the switching network connecting the ROBs with the trigger processor farm a further data compression might be of interest. At high luminosity, the large detector occupancy will necessitate an aggressive data compression scheme.

5.1 High Luminosity Scheme

The proposed high luminosity scheme is summarized in Table 4.

5.2 Low Luminosity Scheme

In the "standard" low luminosity scheme, the full stream of 24 low threshold bits and 3 high threshold bits will be recorded for all non-empty straws. The detector data will then be accompagnied by the straw address. In the recent Istanbul meeting, it was suggested that the data be partitioned in the following way:

1. the TRT is read out through 256 RODs;
2. each ROD is divided into 4 groups (corresponding to 4 zero suppression engines);
3. each group is divided into 16 blocks;
4. each block contains data from a maximum of 26 straws.

In the data stream, each of the $256 \times 4 \times 16 = 16384$ blocks would then be preceded by a 32 bit header word which includes the group address and the block number. Each straw in a block would occur as a 32 bit word where the data occupies 27 bits and the straw identifier the remainder 5 bits. The block would be terminated by a 32 bit trailer word.

With 26 straws per block, we assume that all blocks would have to be read out for every L1A even at low luminosity. Therefore the data volume corresponding to the block wrapping (header + trailer = 8 bytes) is 512 bytes for the 64 blocks of one ROD.

5.3 Alternate Scheme

If one were to apply the high luminosity compression scheme also at low luminosity (low occupancy), the data volume would be dominated by the empty straws. This observation has led Nick Ellis to propose an alternate scheme [3], which is based on the high luminosity scheme, but with the following modifications:

- Discard the empty straw codes;

Luminosity	1×10^{33}	2×10^{33}	1×10^{34}	2×10^{34}
Non-empty straws	8.01	15.01	53.98	77.16
Valid straws (12.5 ns gate)	2.93	5.87	25.16	43.32
- no leading edge	0.06	0.13	1.48	5.59
- 1 leading edge	2.73	5.41	21.49	33.67
- 2 leading edges	0.14	0.33	2.19	4.07

Table 9: Details on average straw occupancy at various luminosities. Numbers are in percent.

Luminosity	1×10^{33}	2×10^{33}	1×10^{34}	2×10^{34}
Headers/trailers	512	512	512	512
Straw data + address	532	1000	3592	5136
Total	1044	1512	4404	5648

Table 10: Output data volume (bytes) per ROD per L1A for the "full stream" low luminosity compression scheme.

- Send all other codes preceded by an n-bit address offset;
- In case of more than $2^n - 2$ consecutive empty straws, send the empty straw code preceded by the address offset ($= 2^n - 1$) after each group of $2^n - 1$ consecutive empty straws.

For this scheme, clearly the optimum value of n depends on the detector occupancy.

6 Estimated Data Volumes

Table 9 summarizes the average straw occupancy at four different luminosities. Details are given on the occurrence of *non-empty* straws and *valid* straws, and for the valid straws on the occurrence of 0, 1, and 2 leading edges. This information can be used to estimate the data volume for the three compression schemes described above.

Tables 10 and 11 presents the average output data volume per ROD per L1A for the low and high luminosity compression schemes, respectively. Assuming an L1A rate of 75 kHz, we notice that 1000 bytes per ROD per L1A correspond to 75 Mbytes/sec on the ROD to ROB link. This should be compared to the tolerable bandwidth which has an absolute maximum limit of 160 Mbytes/sec. Already at a luminosity of 2×10^{33} cm^{-2}s^{-1} we are approaching this limit with an estimated traffic of 115 Mbytes/sec. On the other hand, the high luminosity compression scheme is comfortably within the bandwidth limit even for twice the design luminosity. Were one to use this scheme at low luminosity, however, about three quarters of the data flow would be caused by empty straws.

Table 12 shows results on the data volume from the alternate compression scheme. Clearly the data from the non-empty straws are the same as in Table 11. The only difference between the two schemes is the way the empty straws are represented. For the alternate scheme, the corresponding data volume has been estimated with a small toy Monte Carlo, where non-empty straws were placed randomly according to their apropriate occupancy in the readout

	Bits/straw	1×10^{33}	2×10^{33}	1×10^{34}	2×10^{34}
Empty straws	2	383	354	192	95
Non-empty, non-valid	3	32	57	180	211
Valid: - no leading edge	6	1	2	18	70
- 1 leading edge	9	51	101	402	630
- 2 leading edges	12	3	8	55	101
Total, valid straws		55	111	475	801
Total, non-empty straws		87	168	655	1012
Grand total		470	522	847	1107

Table 11: Output data volume (bytes) per ROD per L1A for the high luminosity compression scheme.

	$\mathcal{L} = 1 \times 10^{33}$		$\mathcal{L} = 2 \times 10^{33}$	
Condition	All non-empty	Valid only	All non-empty	Valid only
$n = 3$	232	192	325	254
$n = 4$	192	139	**310**	207
$n = 5$	**182**	113	329	**189**
$n = 6$	192	**103**	359	190

Table 12: Output data volume (bytes) per ROD per L1A for the alternate low luminosity suppression scheme proposed by Nick Ellis.

sequence from *groups* of $16 \times 26 = 416$ straws (c.f. Section 5.2). Here each group corresponds to a separate zero supression engine. To flag the start of a group in the data stream, it has been assumed to place a header record of length $n + 2$.[1] No trailer has been assumed. The output data volume depends on n, the number of bits used for the address offset. By adjusting n to its optimal value, 5 (4), this scheme leads to a reduction in the data volume by a factor 0.39 (0.60) relative to the high luminosity scheme at a luminosity of 1 (2) $\times 10^{33}$ cm^{-2}s^{-1}. Under this condition the data rate is 14 (23) Mbytes/sec. A further reduction in the data rate could be obtained by dropping non-valid straws. In that case the data rate would be 8 (14) Mbytes/sec.

Notice that in a recent study of the LVL2 Trigger system an average raw data fragment size of 300 bytes was assumed from the TRT at low luminosity after pre-processing [4]. Under this condition the TRT data would constitute about 40% of the total data volume in the TRT full scan scenario. The assumed TRT data fragment size compares nicely to the numbers in Table 12, where 182 (310) bytes are expected at 1 (2) $\times 10^{33}$ cm^{-2}s^{-1}.

[1]The 2 bit field could contain the sequence $\langle 11 \rangle$ which is unused in the scheme of Table 4.

7 A revised scheme

Based on the conclusions from Section 4 we can consider a revised scheme, which allows for a 4+4 bit encoding of one leading and one trailing edges. For this scheme we will assume to give up the edge flags since the most economical solution would be to reserve a pre-defined bit pattern to flag if no edge is found. Valid straws would then be encoded as `1_0_H_LE_TE`, where `LE` and `TE` are the two 4 bit edge position fields. The encoding of empty and non-valid straws would be unaltered. This solution would lead to a small increase in data volume compared to Table 12: the raw data fragment size would increase from 182 (310) to 194 (333) at a luminosity of 1 (2) $\times 10^{33}$ cm^{-2}s^{-1}.

8 Conclusions

- An optimized 12.5 ns validity gate provides a signal efficiency of 97%. It reduces the amount of straws giving signal due to out-of-time bunch crossings by about a factor 5, and thus the total data volume by a factor 2.5;
- The gain in the drift time measurement from the recording of two leading edges looks to be marginal. For less than 1% of all signal straws, the correct drift time information resides in the second edge;
- The foreseen high luminosity ROD compression scheme is within the bandwidth constraints also at twice design luminosity;
- The "full stream" low luminosity scheme is close to saturating the ROD to ROB link bandwidth;
- An alternate low luminosity scheme, which is identical to the high luminosity scheme apart from the way empty straws are encoded, gives a much reduced data volume. Previous estimates of the data volume could be confirmed. These estimates are being used by the LVL2 trigger group for the design of the trigger network;
- A revised scheme, where one leading and one trailing edge are encoded in two 4 bit fields gives an increase in the data volume of about 7% relative to the alternate low luminosity scheme. This scheme would allow for the Time-over-Threshold measurement.

References

[1] Inner Detector Technical Design Report, CERN/LHCC 97-17, 30 April 1997.

[2] P. Lichard, Transparencies from the Istanbul meeting.

[3] N. Ellis, personal communication.

[4] *Paper modelling of the ATLAS LVL2 trigger system*, J. Bystricky and J.C. Vermeulen, ATL-DAQ-2000-030, 4 April 2000.