Redundancy and Synchronisation in the SCT

R.B.Nickerson, A.R.Weidberg, Oxford University

1. Introduction

There has been considerable discussion within the ATLAS SCT[1)] community of the need for redundant data paths. The motivation for this arises from the known finite lifetime of data transmission links and the fact that repairs will be, at best, very difficult. The desire for redundancy in the data links can be generalised to a desire for 'single-point fail-safety'. In the SCT this can be taken to mean that any single element in the electronics data gathering chain which fails will cause loss of data from at most one front-end chip (128 channels, or 0.002% of strips). The bypass ring architecture[2),3)] of the SCT front-end system is chosen to allow single point fail-safety by providing secondary data paths for use in case of failures.

It should be noted that increasing segmentation does not help reduce the fraction of the system that will fail. Increasing the number of data links reduces the number of strips lost per link failure, but increases the number of failures. The SCT has a statistically large number of links so that the fraction of strips lost is independent of the number of links.

A separate issue within the SCT read-out architecture is that of synchronisation. It is grouped with the redundancy issue in this document because the two are closely related via the interconnections between modules that they mandate. In the current context, synchronisation is a reference to the need to ensure that all the clock and control lines which come individually down to the detector have the correct timing relative to the beam crossing at the front-end chips. This synchronisation can be achieved using beam events, but during the (extended) pre-beam commissioning period there are insufficient cosmic-rays to perform this function.

The purpose of this note is to summarise the arguments for single-point failure protection and beam independent synchronisation and to describe briefly the system under consideration to address these issues[2)]. Emphasis in the latter case is on the connections that are required between chips and groups of chips.

2. SCT Architecture

It is easier to discuss redundancy and synchronisation in the context of a specific system. The SCT binary baseline system is used here. More detail can be found in[2] and a very similar architecture is described in [3] . The system is shown in Figure 1.

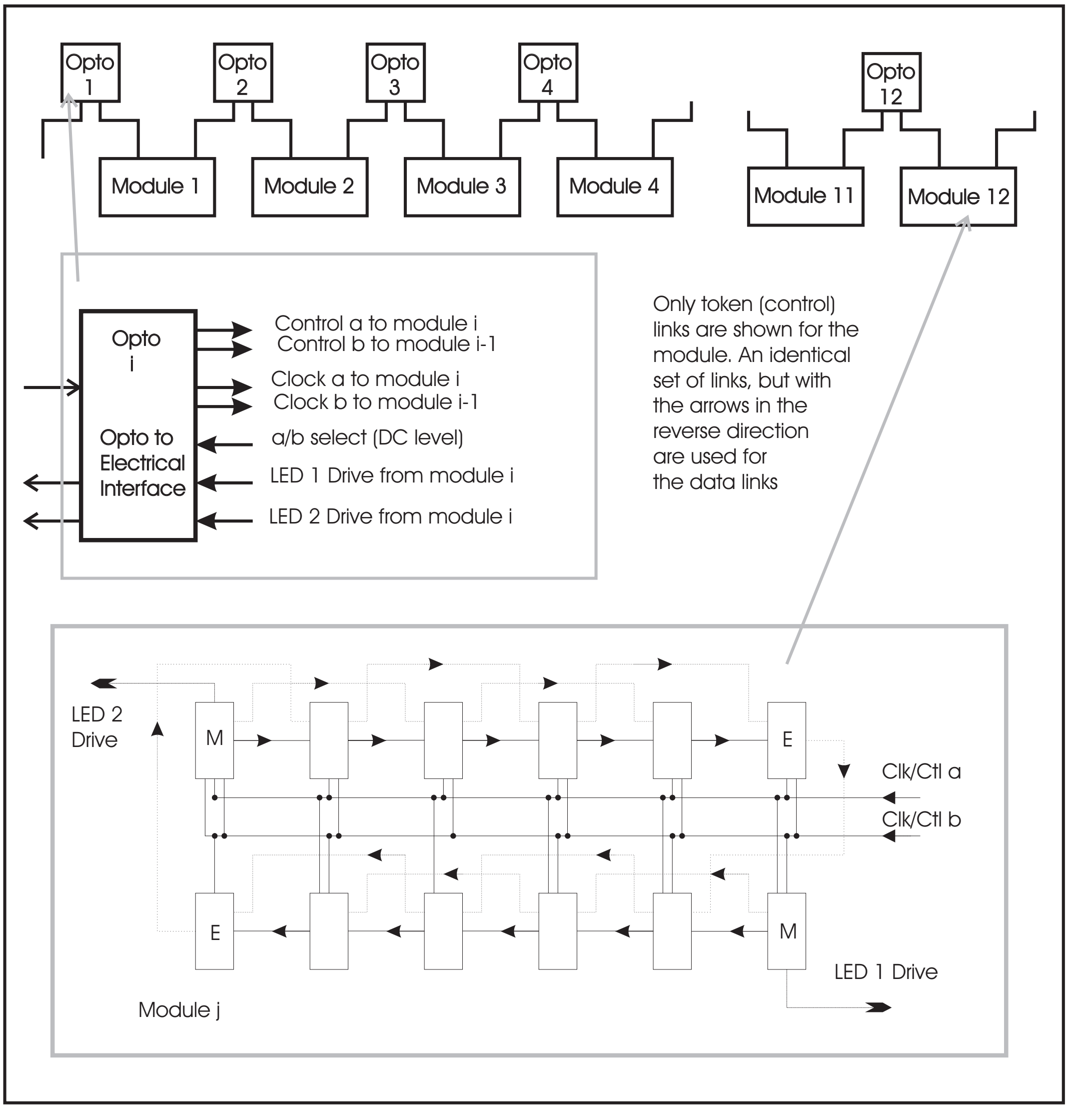

Figure 1 SCT readout architecture

The data is readout from the FE chips via a serial link. There is also a serial path for the data control token ring. After the receipt of a Level 1 Trigger the Master Chip starts the readout by sending the preamble, header data and its own event data to the LED driver on the nearest optoelectronic interface board. A few clock cycles before the last bit of data has been sent it sends a token to the neighbour slave chip which responds by sending its data packet to the Master where it is appended to the data stream. The process continues until the End chip sends its data after which it appends a trailer to the end of the data stream. Single-point fail-safety is implemented by the bypass links (shown dotted). If a slave chip fails then the data and tokens can be re-routed to bypass the dead chip. If a Master chip or an LED data link fails then the data can be re-routed via the LED link for the opposite side of the module.

The clock and control information for a module are sent down a single fibre using a biphase mark encoding[4]. The separate clock and control data are decoded by the DORIC chip[4] on the optoelectronic interface board (as shown in Figure 1). The clock and serial control data are sent to the nearest module. Single-point fail-safety is implemented by the use of a second copy of the clock and control lines to the neighbouring modules. If a line to a module fails for any reason, the module can take its clock and control data from the neighbour optoelectronic board. In order for this scheme to apply to the first module in a group of 12 there would also have to be links from one group to the next.

3. Data Link Redundancy

Figure 2 shows the 50% failure lifetime curves for one candidate radiation hard diode that might be used in the SCT.

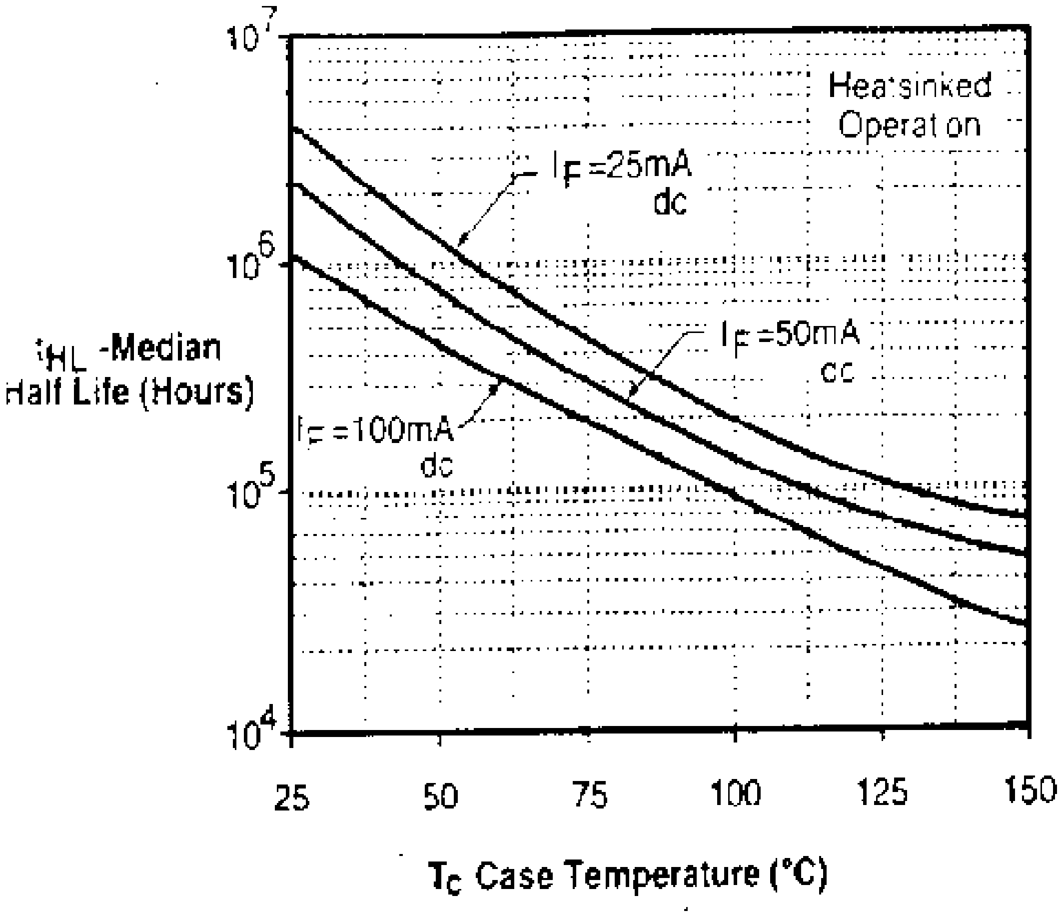

Figure 2 Manufacturer's data for LED Median lifetimes as a function of current and temperature

The expectation is that the diodes would be operated at no more than 20mA, with a

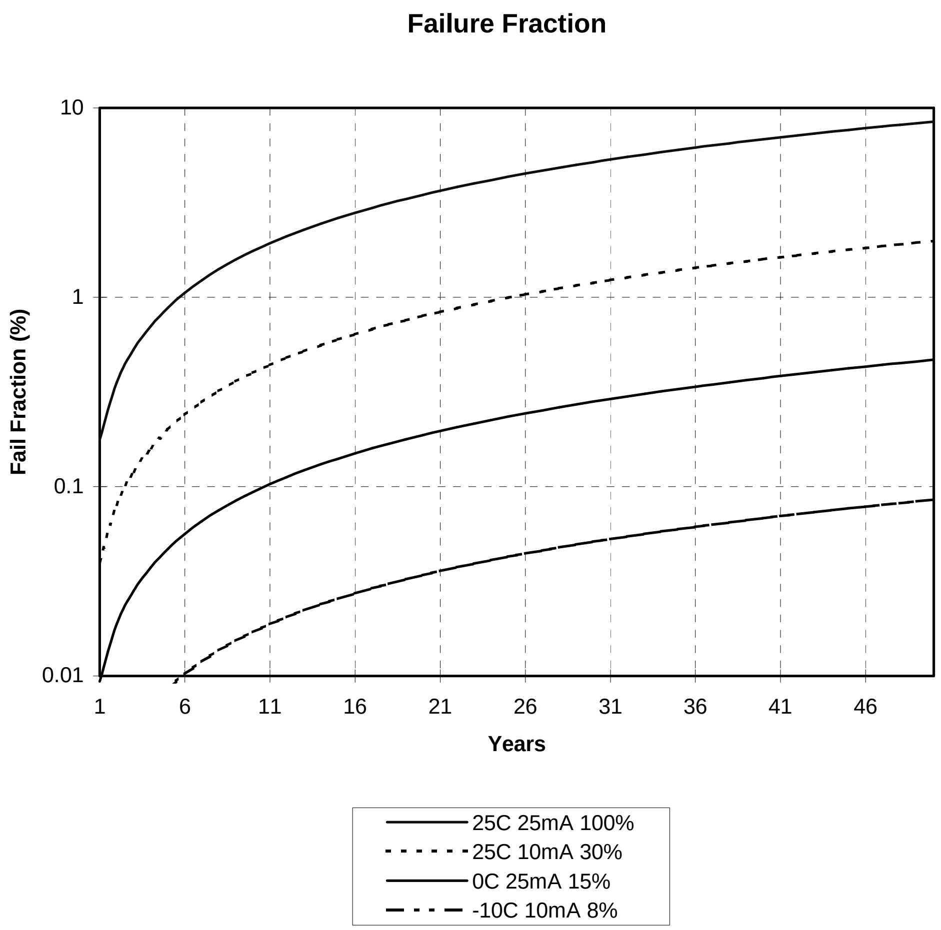

Figure 3 LED failure fraction versus time for different operating conditions.

duty factor <50%, and at a junction temperature that lies between -10C and 0C. Figure 3 shows the fraction of the data links that will fail with a range of assumptions. Taking an LHC year as 10^7 s gives a duty cycle of 30% for DC operated LEDs. The most optimistic case assumes -10C operation, 10mA in the diode, and a 25% duty cycle[1] (this corresponds to the curve labelled 8%). These are not unreasonable assumptions. The worst case corresponds to the case that the temperature is 0C, 20mA operation and a 100% duty factor. This would correspond to the diodes being on all the time, including when ATLAS was not running.

What is clear from Figure 3 is that with the pessimistic assumptions regarding operating conditions, the system is quite marginal in terms of lifetime. All the safety factor resides in assumptions regarding more beneficial operating conditions.

[1] during data taking the link is expected to be used 50% of the time and 50% of the transmitted bits will be “0”s.

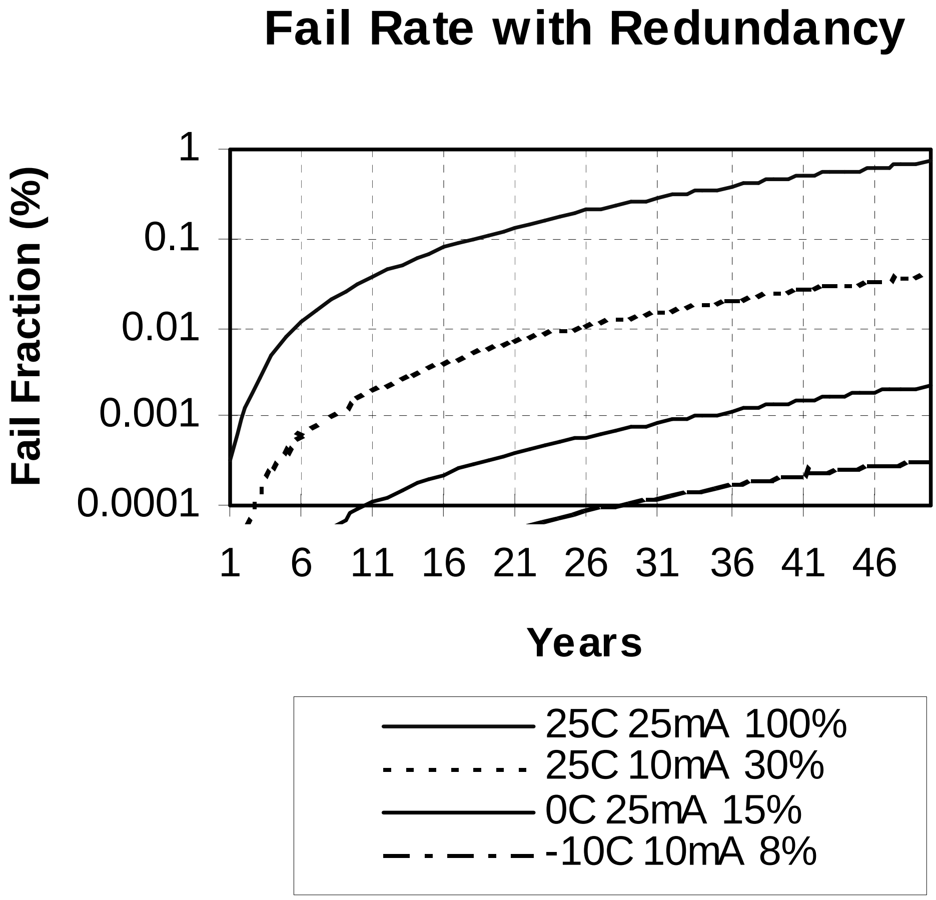

Figure 4 Link failure rates with a redundant system for different operating conditions.

The only way to extend the average lifetime is to utilise a redundant system. As highlighted in the introduction, increased segmentation does not help; alternative data paths for a given region of detector are required. Figure 3 shows the fractional loss of data as a function of time for the same range of assumptions as in figure 2, but with two data links available for each region of the detector. Each link is individually able to transmit all the data, so that two links must fail before a region of strips is lost. Clearly when the lifetime of the system is short in comparison with the requirement, there is still a large probability that two links fail and redundancy is of no value. When the intrinsic lifetime of the non-redundant system is very close to the requirement, the gain is very large, since the probability of both links failing is the square of the probability of a single link failing. A marginal system becomes very safe.

4. Front-end Chip Failure

In the case that one chip fails in a serial ring, the ring is lost. Very little is known about the lifetime of chips after irradiation. Lifetime is also dependent on the details of the design, manufacturing and handling and presumably the environment the chips will operate in. Overall the lifetime of the front-end chips is very poorly understood.

Provided that the cost is low, it becomes prudent to design the system such that a single chip failure does not result in the loss of the data from other chips in the data ring. For a given chip failure rate, this reduces the quantity of lost data by a factor of 3.5 for low chip fail probabilities in the SCT baseline implementation. (6 for a true ring architecture). As with the data links, this is of little value unless the chips are close to having adequate lifetime *a priori*.

5. Synchronisation

For the efficient operation of the FE electronics the BC clock must be synchronised to about 1 ns. (This allows for other timing errors expected in the system.) It would be difficult to achieve this by dead reckoning; the fibres will probably be cut to an accuracy corresponding to about 10 ns. It will therefore be essential to adjust the synchronisation of all the clock and control fibres. This could be done with particles from real events, however it is desirable for debugging purposes to take cosmic ray data when the detector is installed in the pit, prior to LHC operation. To measure detector efficiencies properly requires precise timing. It would therefore be preferable to have a mechanism for synchronising the BC clocks for the different modules independent of particles. It is also desirable to have a system capable of quickly checking the synchronisation for use during data taking since any beam related technique requires analysis of data. If there were no other method the synchronisation could be done using cosmic ray data. To calculate the time required to do the synchronisation with cosmic ray data in a binary system we can use the slope of the efficiency versus time delay measured with the DDR2 module in the H8 test beams shown in Figure 5 overleaf.

From a least squares fit to a straight line in the rising edge part of the curve the slope was[2] 0.11 ns^{-1}. Two estimates of the time required will be made:

1)Optimistic. In this case we assume that the muon system provides accurate track impact predictions in the SCT so that an efficiency per strip can be determined with negligible background from noise hits. If we do a 3 point scan, we can determine the time offset by performing a least squares fit. If the scan is done at the optimal points on the curve then the required 1ns precision can be determined with as few as 10 events per scan point. It might be necessary to do a first coarser scan to determine the optimal points for the fine scan. This would mean that about 60 events would be sufficient for a scan. As shown in Appendix A The Cosmic Ray rate is estimated to be 2.6/hour for a "vertical" module (ϕ=90 degrees). Therefore the scans would take about one day for a "vertical" module and slightly longer for the other modules.

2) Pessimistic. In this case we assume that the trigger is given by scintillator hodoscopes above and below the experiment and that the muon tracking is not available. In this case one would have to look for hits in area which would depend on the granularity of the scintillator hodoscope. For a typical size scintillator this

[2] A. Grewal, analysis of DDR2 H8 test beam data.

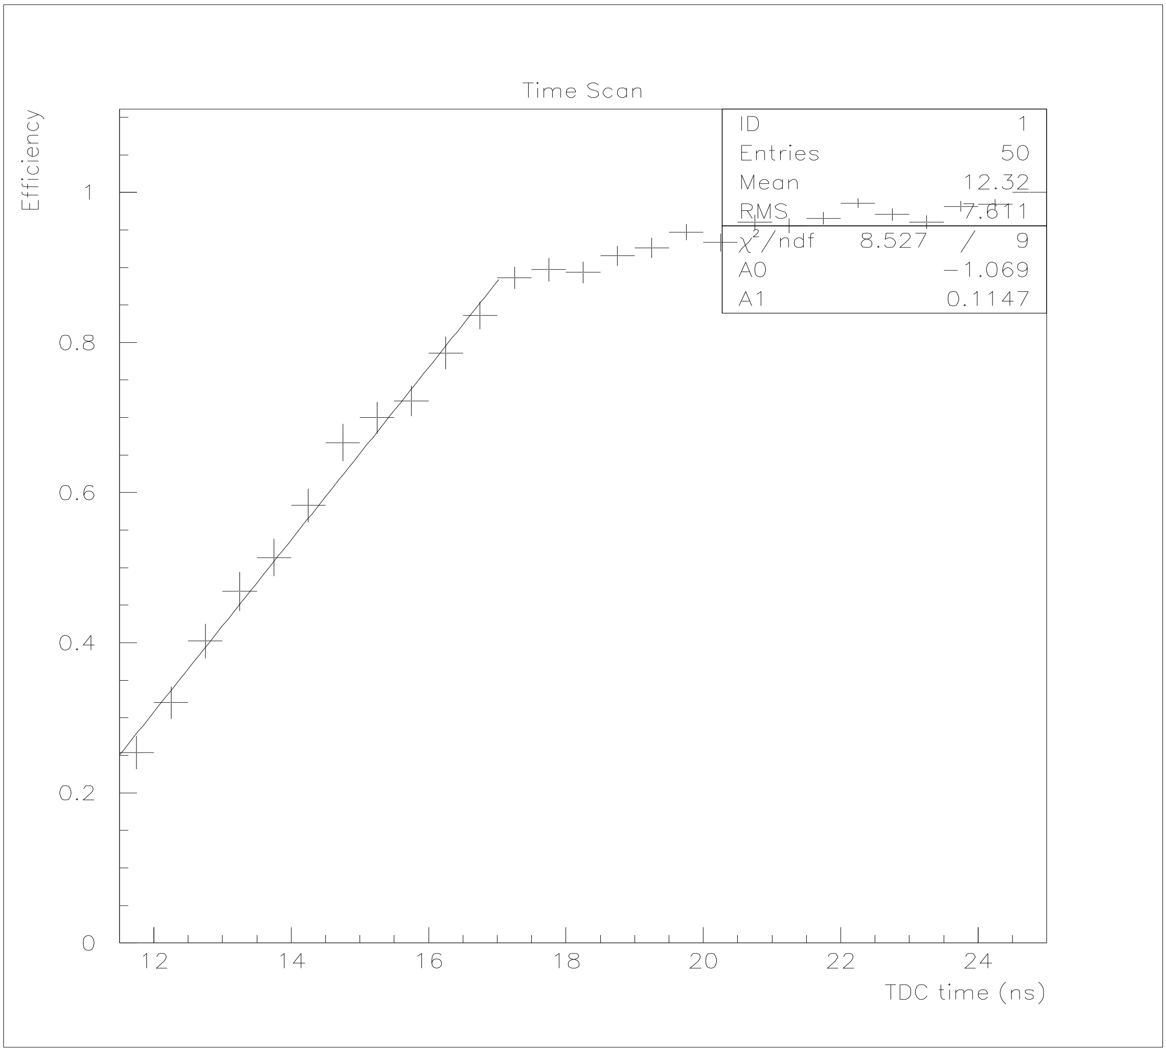

Figure 5 Efficiency versus TDC value for the DDR2 module

would correspond to about 10 modules and therefore in any given module there would only be a real hit in about 10% of the triggers in the corresponding scintillator. The trigger rate in one scintillator hodoscope would, of course, be 10 times the rate in an SCT module. The hit rate would be dominated by background noise hits and so the number of triggers required at each scan point would be dependent on the noise occupancy. The required number of triggers for each scan point is given as a function of noise occupancy in Table 1 below. Hence the total time required for two 3 point scans would be in the range 3 months to 2 years.

Noise Occupancy	Number of Triggers
10^{-3}	8000
5×10^{-3}	40000
10^{-2}	80000

Table 1 Number of triggers required per scan point as a function of noise occupancy.

In conclusion, if we can not rely on precise muon tracking being available before the start of the LHC and we want to study detector efficiencies using Cosmic Ray data

then it is essential to have a scheme for synchronisation of modules which is independent of beam or Cosmic Rays.
As a subsidiary motivation for a hardware synchronisation scheme, it might be noted that, with a highly parallel system of clock distribution, it is desirable to have a quick method of checking the timing at the start of each data taking period. A technique that relies on interpreting efficiency curves for the detectors is less likely to be a fully automatic process.

A suitable synchronisation scheme will be implemented in the SCT binary system. The master ABC chip will have a mode in which it sends synchronisation data back to the counting room. In this mode the clock is used to toggle a T type flip-flop[3] The output of the flip flop is thus a 20 MHz clock signal, the phase of which is fixed by the phase of the clock. Note that normal data will be sent at 40 Mbits/s NRZ[4] ,which requires the same bandwidth transmission system as a clock of 20 MHz. At the receiving end of the data links, the phase of the 20MHz data relative to a local copy of the BC clock, t_1, will be measured by the data receipt electronics. The front-end module will then be instructed to take its clock from the neighbour module and the new phase of the returned 20MHz data relative to the same local BC clock, t_2, will be measured. The measured phase difference between the two clocks is then $\Delta T = t_2 - t_1$
The phase difference can be adjusted away in the counting room clock & control distribution module by delaying one or other of the clocks compared.

6. Cost Implications

The costs that should be associated with this scheme depend on the rules used for evaluation. In the case of the SCT, a factor of two safety is designed into the bandwidth of the data links to allow for higher than anticipated data rates. No additional links are added to incorporate the redundancy scheme, so that the costs associated with redundancy are small. The factor of two safety is assumed to be enough to allow for all sources of bandwidth increase. The number of clock/control lines are not changed, but they must be bussed further. The cost increases are therefore those associated with interconnects and small changes to front-end chip areas.

7. Conclusions

This note has reviewed the arguments in favour of redundant systems in the SCT readout. The main conclusions are:

- Data link redundancy should be a requirement in the SCT architecture.
- Extending this to single point failure protection seems prudent, given the lack of knowledge of irradiated chip lifetimes.
- The ability to synchronise without beam should be included in the system.
- The cost of these features is expected to be small.

References

[3] A T type flip-flop behaves like a JK flip-flop with the J and K inputs held at logic 1.
[4] NRZ = non return to zero.

1 SCT Backup document is available on W3 at URL: http://atlasinfo.cern.ch:80/Atlas/Welcome.html

2 M. French et. Al, ”The SDC Silicon Tracker DAQ”, Nucl. Instr. Meth. A342 (1994) 214.

3 ABC Specification document is available on W3 at URL: http://hepunx.rl.ac.uk/atlas_binary/documents.htm

4 D.J. White, DORIC chip specification, RAL.

Appendix A Cosmic Ray Rates.

The total flux of vertical cosmic muons at sea level is

$I = 80\ m^{-2}\ s^{-1}\ sterad^{-1}$

The stopping power of the calorimeter plus concrete shielding is approximately 4 GeV (for muons that pass through the entire detector). The fraction of cosmic muons with p>4 GeV is about 6%. The dominant contribution to the cosmic rate will come from muons that come down the pit without penetrating additional earth. The pit has a diameter of 16 m at a depth of 100m. Therefore the solid angle subtended is given by

$\Delta\Omega = 2\ \pi\ (1 - \cos\theta\ ) = 0.02$ sterad

The area of a SCT module is approximately 0.06 by 0.12 m^2 . This gives an estimated (lower bound) rate of 2.6 per hour for a module at ϕ=90 degrees.