

# Evaluation of twisted pair data transmission for the SCT

**R. Kowalewski and S. Marconi, University of Geneva**  
**A. Beer and P. Lichard, CERN**

## 1 Introduction

A data transmission system for the SCT using individually shielded twisted pair cables within the detector volume was proposed in June 1996 as an alternative to a system using optical fibers. A program of work was undertaken to evaluate the feasibility of such a system and to compare it to the optical system. This note gives the results of that program of work.

The note is structured as follows. A description of the cables will be given, followed by an overview of an SCT data transmission system using twisted pair cables. Some of the important aspects of performance will then be discussed in separate sections. Where possible these projections of performance are based on measurements made on prototype cables and SCT modules. The final sections describe how to optimize the performance versus cable mass and give the material budget and cost estimate for the twisted pair and optical systems. While the primary purpose of this note is to aid in evaluating shielded twisted pair (STP) cables for use in the SCT data transmission system, remarks will be made where relevant concerning the possible application of STP cables for pixel data transmission.

This work was initiated in response to the decision by the TRT community to use twisted pair readout. Their decision was motivated primarily by cost projections indicating savings of a factor of  $\approx 3$  over fiber optic readout.

## 2 Description of the cables

The fine gauge (AWG36) cables studied for use within the detector volume are now available at the CERN stores, item SCEM 04.01.64.036.4. They consist of

- 2 insulated signal wires, one drain wire, Al/polyester foil wrapping
- Signal wire Ag-plated Cu, diameter 0.13 mm; 4 drain wires diameter 0.0635 mm
- 0.39 mm silicone polyetherimide copolymer (HFI150) insulation
- Outer cable diameter 0.9 mm

The cable cross-section is shown in Fig. 1. The cables have been supplied by Habia Cable<sup>1</sup>.

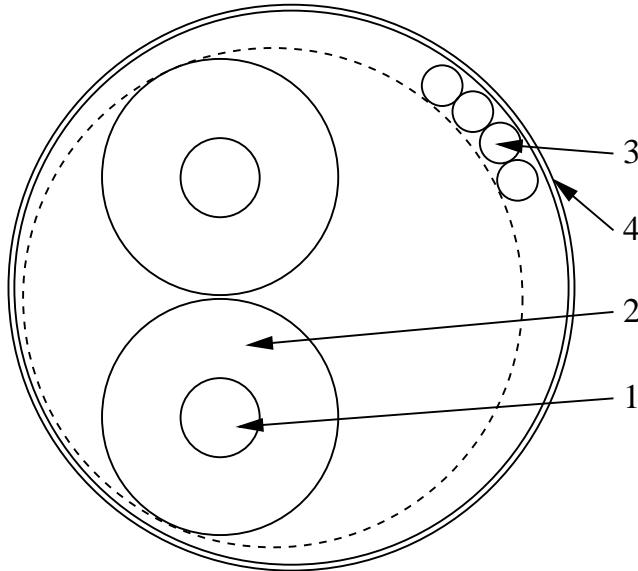


Figure 1: The cross-section of the AWG36 fine twisted pair cables described in the text. The copper conductors (1) are surrounded by a silicone-based insulation (2) and twisted. A 4-stranded drain wire (3) is twisted around the pair in the opposite sense to ensure good contact with the aluminized inside of the plastic outer shield (4).

The TRT group measured the electrical properties for a recent version of the thin STP cables under consideration. The electrical properties as measured using the LANTEK PRO XL cable tester:

Impedance	92.5 $\Omega$
Attenuation (10 m)	6 dB @ 40 MHz, 8.9 dB @ 100 MHz
Tpd	5.3 ns/m
Capacitance	57 pF/m
DC resistivity	14 $\Omega/10m$

These cables have been approved by the CERN Technical Inspection and Safety commission. The cables have been tested successfully to a dose of 1 MGy of ionizing radiation. The anticipated production cost for this cable is approximately CHF 1/m. A p.c. board to facilitate the soldering of the signal cables and drain wires to connectors has been developed by Arno Beer and is available at CERN.

It is anticipated to switch to standard AWG28 multi-conductor twisted pair cables once outside the ATLAS calorimeters. The cables (typically 25 twisted pairs/cable) are surrounded by an overall shield, but are not individually shielded (denoted UTP for unshielded twisted pair). This thicker cable has the following properties:

---

<sup>1</sup>Habia Cable BV, Voorerf 12, 4824 GN Breda

Impedance	100 $\Omega$
Attenuation (100m)	29 dB @ 40 MHz, 59 dB @ 100 MHz
Tpd	6.2 ns/m
Capacitance	61 pF/m
DC resistivity	21 $\Omega$ /100m

The anticipated production cost of this cable (per pair) is CHF 0.1/m. The choice of AWG28 cable is still under investigation in the TRT group.

### 3 System for SCT data transmission

The requirements for the SCT data transmission are to distribute a 40 MHz clock along with trigger and control data to each SCT module, and to provide data transfer from the module to the ROD at an average rate of 35 Mbit/s. The BER should be kept below  $10^{-9}$ . Power consumption, mass, and size of the links should be kept to a minimum. For reasons of redundancy on the hybrid (the ability to bypass a bad FE chip), two independent data links with adequate bandwidth (40 Mbit/s) should be provided to every module. The FE chips output data using the LVDS signalling standard and can drive the signals over 8 m of STP cable<sup>2</sup> to the patch panels in the muon crack. The cable runs within the inner detector volume will need to be as light and small as possible. A switch to larger gauge cables will be made at the patch panels. A repeater will be needed at this point to restore good signal quality. The 70 m from the patch panels to the ROD might require a second repeater or additional compensation circuitry.

The FE chips currently use an NRZ encoding for the output data (motivation: to minimize the duty cycle of the optical components); this encoding could also be used with STP transmission, although a balanced scheme might be more natural.

As redundancy is not considered necessary for passive components, use of STP links would eliminate one of the motivations for distributing 2 independent clocks to each SCT module and therefore for interlinks between SCT modules. (These interlinks would still be needed if one wished to synchronize SCT modules in the absence of beam data.)

#### 3.1 Use of DORIC

The DORIC chip has been designed to serve as a receiver for the signals coming from the PIN diode on the optoboard and to demultiplex the clock and control signals, which are sent from the ROD using a bi-phase mark encoding. A successful development of this chip may allow use of a similar chip (the same decoder with an LVDS receiver) to be used in an STP data transmission system, thereby eliminating one cable per module; the use of multiplexing with an STP system would need to be tested. Even in the absence of multiplexing an on-hybrid repeater may be necessary since the load represented by the hybrid may not be well matched to the cable impedance.

---

<sup>2</sup>The maximum length for the power tapes, which have the same routing as the cables, is quoted as 6 m in the TDR, so 8 m may be an overestimate.

## 3.2 Power budget

According to the ABC designers[1] the LVDS output driving the data is estimated to consume about 64 mW to 100 mW. (These values are larger than the specified and measured values of commercial LVDS drivers (e.g. National Semiconductor DS90C031) which draw 4 mA/driver at 40 Mbit/s.) This output is only enabled on 2 of the 12 chips per SCT module. The power budget for optical links includes this plus the driver power for the LEDs. The latter varies with drive current (and therefore total fluence), and is expected to be[2] 96 mW at 10 mA drive current and 176 mW at 20 mA. The total power budget of the 12 chips per SCT module is listed as 5.8 W (3.8 mW/channel). The power needed to drive the LEDs (at 20 mA) is 3% of this.

## 3.3 Cooling requirements

The STP cables need no cooling. A cooling contact is foreseen for the package holding the optical components to enhance component lifetime.

## 3.4 Repeater and compensation circuits

A repeater using TTL-LVDS driver and receiver chips has been built at the University of Geneva and will be tested shortly. Repeaters in the final system will need to be developed in a radiation tolerant process due to the neutron fluence at the outer patch panel where the repeater will reside. They will also need to be compact.

Compensation circuits could be used to increase the bandwidth available on the AWG28 UTP cables. Simulations indicate compensation circuits should more than double the available bandwidth. Another approach is to use improved UTP cables for the 70 m segment.

## 3.5 Electrical isolation

The SCT has adopted a “groundless” system in an attempt to minimize cross-coupling between disparate elements of the system. Separate power supplies are foreseen for each module. Individual voltage sources are referenced to each other via high-value resistors.

A number of  $\sim$ DC electrical connections exist between the FE hybrids and the off-detector electronics, e.g. power cables, calibration lines, control lines, bias voltages, temperature monitoring. Noise picked up along these conductors is a potential problem even in an optical fiber readout system.

An STP data transmission system would add a high frequency connection between FE and off-detector. In such a system two additional sources of noise must be considered, namely noise radiated from the cables and ground loops via the cable shields. The powering of the repeaters in such a system and the electrical connection of the cable shields will need to be made with care. The two easiest solutions to implement within the inner detector volume (up to the first patch panel) would be to either ground the cable shields at each FE hybrid and keep them floating at the patch panel, or to ground them all at the patch panel and leave them floating at the each FE hybrid. The FE and the repeaters

will in any case need to be referenced to each other. Only a relatively large scale system test can reliably determine which solution is better.

The results of tests of noise induced on a prototype SCT module by the STP cables are discussed below.

### 3.6 Assembly and testing

The cable assembly into bundles (e.g. 24 pairs) is labor intensive; the costing given later includes this labor. Testing of assembled cables can be done very quickly and simply by plugging the fully assembled cables into a BER test board. Cable faults at the connector can be easily repaired; those in the cable itself would require replacing the affected cable.

## 4 Tests using twisted pair cables

Two sets of tests were carried out on the STP cables. Tests of the data transmission in the cables (e.g. bandwidth, bit error rate (BER), jitter, cross talk) were performed using a custom board. Additional tests were performed to study the potential for inducing noise on sensitive FE electronics due to signals in the STP cables or on the cable shields.

### 4.1 Signal transmission on miniature twisted pair cables

A custom test board (see Fig 2) has been designed by Peter Lichard to allow BER tests on 4 bundles of 20 signal pairs. The board consists of separate driver and receiver circuits for sending and receiving LVDS<sup>3</sup> signals. Random or periodic data can be generated at rates up to 80 Mbit/s (some components on the test board stop functioning beyond this rate). Each connector has 20 pairs with identical signals, 4 dummy pairs for studying cross talk, and 3 unattached spares. The signals sent over the 20 pairs are either LVDS levels (one of the 4 bundles is configured this way) or a reduced voltage swing (roughly 1/2 of LVDS) obtained by adding  $100\Omega$  termination resistors at the driver. The TTL signals from the driver side are sent both to the TTL-to-LVDS driver chips<sup>4</sup> and to a comparator on the receiver side of the board via a strip line. Digital and analog delays allow the synchronisation of the data and the adjustment of the strobe at the comparator. Separate counters are kept for single pair errors and simultaneous errors on 2 or more pairs.

The cable shields are all connected to each other and to two pins of a 68-pin connector using a small custom PC board (see Fig. 3). Jumpers on the test board allow the shields to be attached to either the driver ground or receiver ground.

BER tests could be performed to required levels ( $10^{-9}$ ) in a few minutes at 40 Mbit/s. The quality of the data transmission was measured using either the timing window for the strobe over which one had good transmission or using “eye plots” (persistence traces on the scope) of the type shown in Fig. 4. The open area in the eye plot is a measure of the transmission quality.

---

<sup>3</sup>The voltage difference between the signal and its complement is typically 350 mV.

<sup>4</sup>The driver and receiver chips are National Semiconductor DS90C031 and DS90C032, respectively.

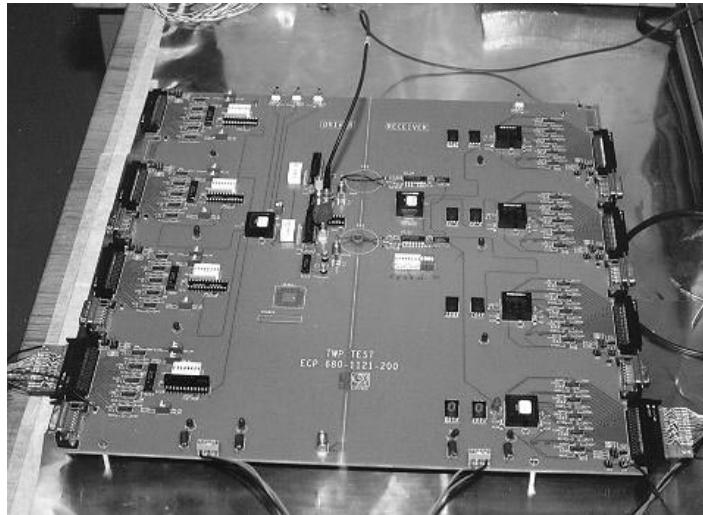


Figure 2: Photograph of the test board for bundles of STP cables.

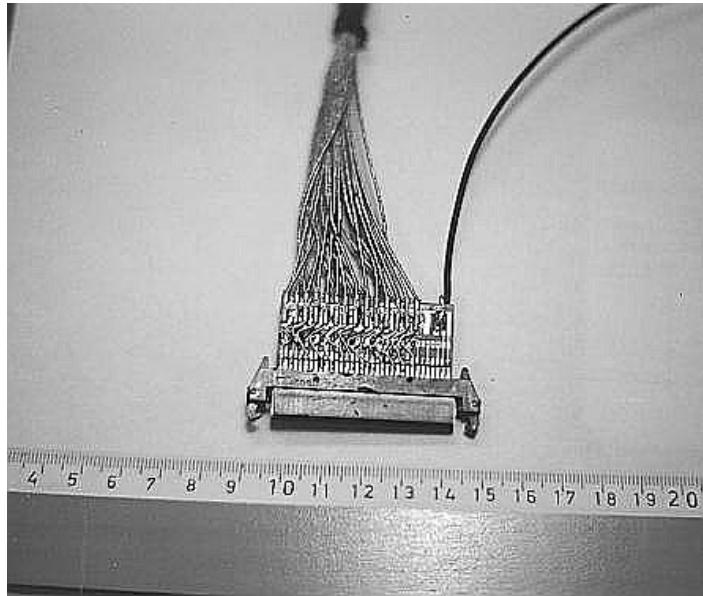


Figure 3: Photograph of a 27-pair connectorized cable for use in the TRT system. The large wire allows easy connection to the cable shields; it would not be present in a real connector. The connector shown also has provision for additional power cables (not shown) which would be removed in an SCT implementation.

The board was operated for a variety of different clock speeds, signal swings, cable lengths and cable types. Table 1 gives a summary of the results. The BER was acceptable in all cases except those marked bad. The quality of the transmission, as determined by the eye plots, is given in the final column. The transmission was similar under all shield connection options (floating, driver side, receiver side, and both sides). AC and DC signals of up to 2V were introduced between the driver and receiver ground planes

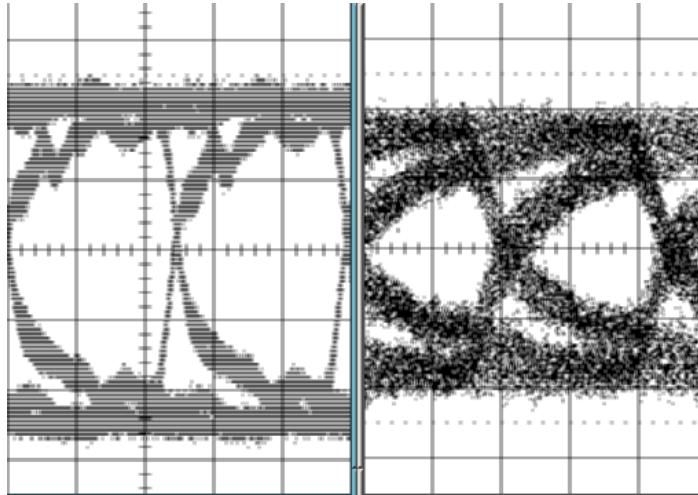


Figure 4: The persistence traces for transmission of random data at 40 Mbit/s over 10 m (left) and 20 m (right) of AWG36 STP cable. Each plot corresponds to about 2500 sweeps.

without introducing errors in the transmission.

The jitter measured on a 40 MHz square wave signal sent over 10 m of AWG36 STP cables was  $\sigma = 170$  ps. Measures of the signals on the dummy lines were used to estimate the cross talk. With LVDS/2 signals on 20 pairs the r.m.s. noise on the adjacent dummy lines was approximately 5 mV.

Cable type	length	signal size	data rate	quality
AWG36 STP	10 m	LVDS	40 Mbit/s	excellent
AWG36 STP	10 m	LVDS/2	40 Mbit/s	excellent
AWG36 STP	10 m	LVDS	80 Mbit/s	excellent
AWG36 STP	10 m	LVDS/2	80 Mbit/s	excellent
AWG36 STP	20 m	LVDS	40 Mbit/s	good
AWG36 STP	20 m	LVDS/2	40 Mbit/s	good
AWG36 STP	20 m	LVDS	80 Mbit/s	marginal
AWG36 STP	20 m	LVDS/2	80 Mbit/s	marginal
AWG28 UTP	20 m	LVDS	40 Mbit/s	excellent
AWG28 UTP	20 m	LVDS/2	40 Mbit/s	excellent
AWG28 UTP	30 m	LVDS	40 Mbit/s	good
AWG28 UTP	30 m	LVDS/2	40 Mbit/s	good
AWG28 UTP	50 m	LVDS	40 Mbit/s	marginal
AWG28 UTP	50 m	LVDS/2	40 Mbit/s	marginal
AWG28 UTP	80 m	LVDS	40 Mbit/s	bad
AWG28 UTP	80 m	LVDS/2	40 Mbit/s	bad

Table 1: Summary of BER tests using bundle test board.

## 4.2 Tests using standard unshielded twisted pair cables

Tests of LVDS signal propagation on standard UTP cables of gauge AWG 28 were carried out for lengths of 20, 30, 50, and 80 m. The eye plots corresponding to these lengths are shown in Fig. 5 for data rates of 40 Mbit/s. The signal degradation limits the usable length to about 40 m, as indicated in Table 1. Longer cable runs (one expects a cable length of 70 m from the muon crack patch panels to the SCT RODs in ATLAS) will require either a repeater or a compensation network or a better cable.

Transmission of higher baud rates (desirable for the pixels) over 70 m of UTP cable would almost certainly require a second repeater. An alternative would be to use radiation tolerant optical links to go from the muon crack patch panel to the ROD. Such links are being developed for the LAr calorimeter.

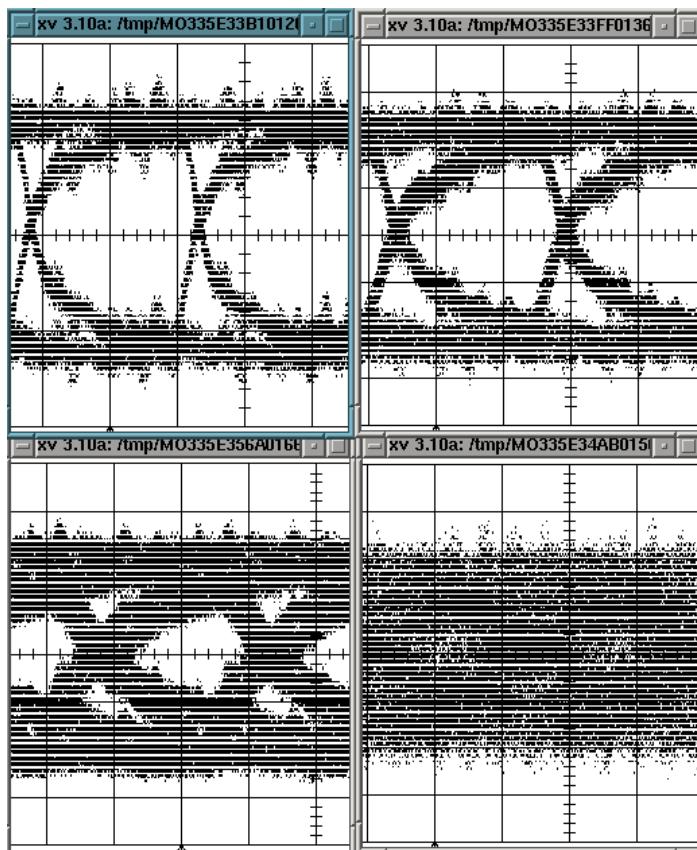


Figure 5: The persistence traces for transmission of random data at 40 Mbit/s over 20 m, 30 m, 50 m, and 80 m of AWG28 UTP cable. Each plot corresponds to about 2500 sweeps.

## 4.3 Noise induced on an SCT prototype module

A second set of tests used the board described above only as a generator of LVDS signals. The STP cables were placed under the backplane of a pair of silicon detectors in the

“CSEM” module<sup>5</sup>, shown in Fig. 6. The distance between the silicon backplane and

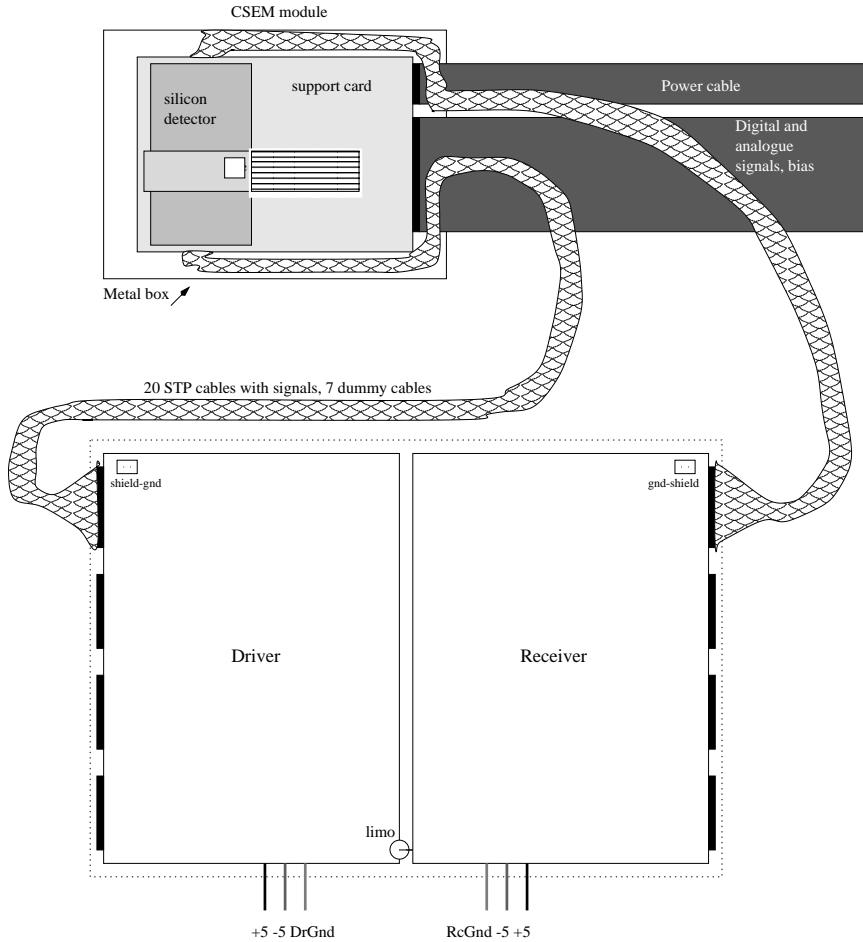


Figure 6: The diagram shows the cable routing for the tests of induced noise.

the bundle was 4 mm. Noise measurements were made on the module by studying the sharpness of the transition from 1 to 0 as a function of the CAFE discriminator threshold. Fig. 7 shows an example of a threshold scan. The noise is derived by fitting an error function

$$F(V, \sigma) = \frac{1}{2} \left( \text{erf} \left( \frac{V - V_{th}}{2\sigma} \right) + 1 \right)$$

to each channel.  $V_{th}$  is the threshold value,  $V$  is the median pulse height, and  $\sigma$  is the noise  $\sigma$ .

The CSEM module suffered two damage incidents, one of known origin, the other of unknown origin. The bonds to one of the CAFE chips were badly damaged when the STP cables were being placed near the module. This caused one of the two chip sets on the module (chip '0' in the jargon) to stop working. All tests were therefore done with the remaining working chip (chip '1'). In an previous damage incident of unknown origin, almost all the short bonds (there were 2 rows of bonds) between the detector

<sup>5</sup>This is a prototype SCT module using detectors from CSEM and binary readout electronics consisting of a separate preamplifier (CAFE) and pipeline chip (CDP).

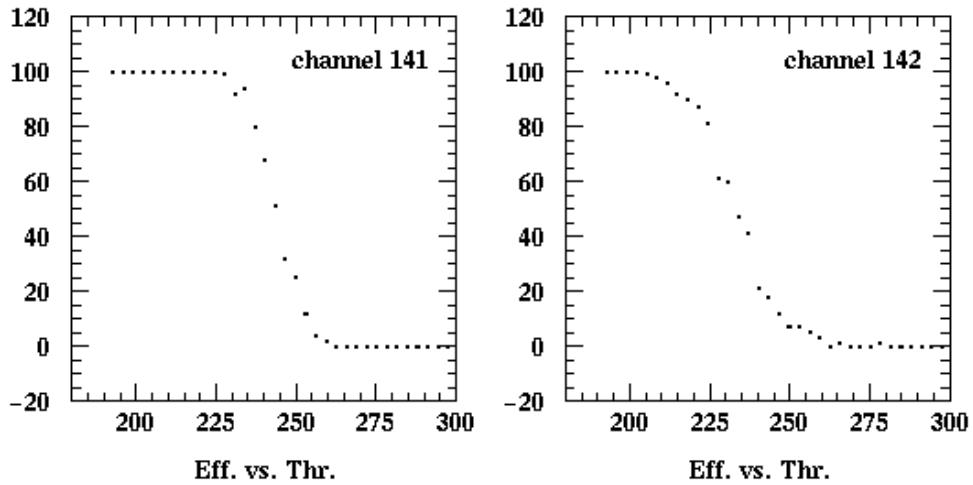


Figure 7: The response to a known calibration charge is shown for two channels. Channel 142 is bonded to a 12 cm long silicon detector, the channel 141 is not.

and chip 1 were broken, while almost all the long bonds remained intact. The noise measurements on the two types of channels (bonded/unbonded) were consistent with pre-damage measurements of the module, indicating that chip 1 was still functioning correctly. The broken bonds between chip and detector enabled the study of whether signals are picked up primarily via the detector or directly on the hybrid. No pickup was observed under any condition on the channels not bonded to the detector. The noise measurements discussed below all refer to channels bonded to the 12 cm long silicon strips.

A series of threshold scans were taken to measure the module noise under various conditions. Initially the bundle test board was placed on an aluminum plate and a copper braid connected this plate to the Faraday shield surrounding the CSEM module. Under these conditions an increase in the module noise was noted whenever the test board was generating signals, *even when there were no STP cables in the module box*. Consequently, we removed the aluminum plate and copper braid and placed the bundle test board further (about 1.5m) from the module. This eliminated the sensitivity to the test board signals when no cables were near the module.

With the cables running inside the module enclosure, parallel to the silicon strips and about 4mm below the detector backplane, a series of runs were taken with and without signals in the cables. The noise measurements, in equivalent electrons at the CAFE input, are shown in Fig. 8. The solid points were taken with the data generation on the board turned off, and therefore with no signals in the STP cables. These are used as a reference, and the line going through these points indicates the average noise value in the absence of data transmission in the STP cables. The open circles are noise measurements taken with signals in the STP cables or on the cable shields. The first 6 open points are with full amplitude LVDS signals (recall there are 20 pairs in the bundle transmitting identical data). The different points have different shield connections, namely

- shield floating,

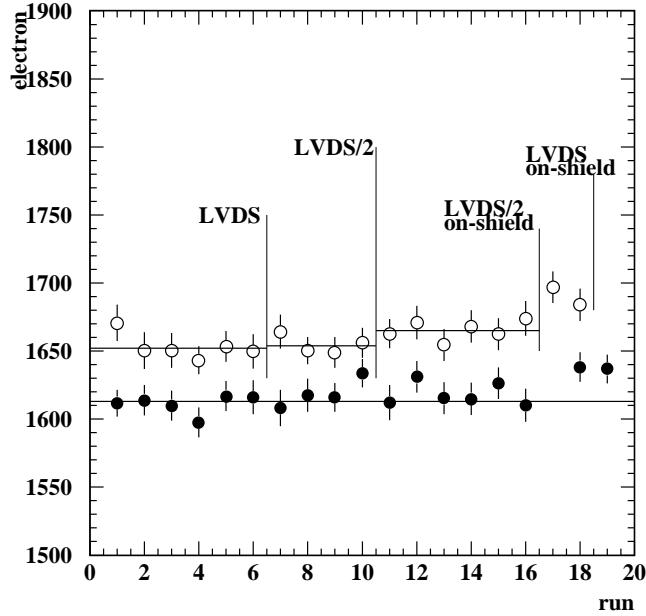


Figure 8: Measurements of the noise at the input to the CAFE preamplifier under a variety of conditions. The solid points are with no signals in the STP cables near the module and the open points are with full or half amplitude LVDS signals. See the text for details.

- shield connected to driver ground plane,
- shield connected to receiver ground plane,
- shield connected to driver and receiver ground planes.

The last two points are measurements repeated as a consistency check. An increase of about 40 electrons is observed with transmission on. Runs 7-10 are the same set of measurements on a bundle with half-amplitude LVDS signals. There is no obvious difference in the noise increase when these reduced amplitude signals are used instead of full LVDS signals. The open points for runs 11-16 (LVDS/2) and 17-18 (LVDS) record the noise observed when one of the wires of a twisted pair was short-circuited to either the shield or one of the dummy wires. In this case a difference is observed between full and half-amplitude LVDS signals.

In addition to the measurements presented in Fig. 8, we observed a significant increase in the noise (from 1600 e<sup>-</sup> to  $\approx 1750$  e<sup>-</sup>) when the cable shield was attached to the Faraday cage surrounding the module.

The observed increases in noise are summarized in Tab. 2. For a module with a nominal noise of 1100 electrons the increase with standard LVDS or half-amplitude signals is about 4.5%. If the observed increases were due to radiation from the pairs, the noise would be expected to scale with signal size; this is not the case. The only difference seen between LVDS and half-amplitude signals was when *non-differential* signals were explicitly sent over the cable shield or through unused lines in the bundle.

Working conditions	$\sigma_{Ref}(e^-)$	$\sigma_{Add}(e^-)$	% inc @1100e $^-$	%inc @1500e $^-$
LVDS signals	$1619 \pm 3$	$331 \pm 25$	$4.2 \pm 0.7$	$2.4 \pm 0.4$
LVDS/2 signals	$1619 \pm 3$	$341 \pm 27$	$4.7 \pm 0.7$	$2.6 \pm 0.4$
LVDS/2 on shield	$1619 \pm 3$	$390 \pm 21$	$6.1 \pm 0.6$	$3.3 \pm 0.4$
LVDS on shield	$1619 \pm 3$	$488 \pm 29$	$9.4 \pm 1.1$	$5.2 \pm 0.6$

Table 2: Summary of the noise measurements using the CSEM module in the presence of STP cables. The reference value is the noise obtained when no signals are present in the cables. The additional noise is quoted as the quadrature difference between the total noise when signals are present and the reference value. The last two columns show the percentage increase for modules with nominal noise values of 1100 e $^-$  and 1500 e $^-$ , respectively.

The source of the observed noise increase is therefore not believed to come from radiation off the twisted pairs. There are other potential sources of noise; any signals coupled into the cable shield can radiate. These spurious signals are unlikely to be a function of the amplitude of the signals transmitted through the pairs. A good candidate for the source of signals picked up on the shields is the large surface of the test board (see Fig. 2), which may couple noise from the TTL signals on the board into the cable shields.

Additional measurements were taken to determine the source of this noise. These measurements were taken with the STP cable bundle in an aluminum box, and with a scope probe making a loop around the cable<sup>6</sup>. The r.m.s. noise measured on the probe is shown in Tab. 3 for LVDS and half-amplitude signals in the cables, and for the signal generation on and off. Values are also given for the case where the signal generation was on but the LVDS driver chips were disabled, i.e. TTL signals running around the board but no data transmission through the cable. The r.m.s. noise measurements have estimated systematic errors of 0.2 mV. The cable shields were floating for the measurements presented in the table, but connecting the shields to the driver or receiver ground planes does not significantly alter the results. These measurements suggest that the bulk of the noise radiated by the cable comes from pickup off the TTL signals on the test board and is independent of the presence or absence (or amplitude) of differential signals on the twisted pairs. Measurements were also taken with one of the signal wires shorted to the shield. In contrast to the results given earlier for the noise induced on the SCT module, short circuiting the wire did not change the noise pickup on the probe.

Translating these results to the SCT is not straightforward. While radiated noise would be expected to scale with the number of cables<sup>7</sup>, noise picked up on the cable shields does not necessarily scale. In addition, pickup on the shields in the SCT would have a different origin than any pickup in our lab tests.

<sup>6</sup>Measurements using the CSEM module could not be taken since the module no longer functioned after having been disconnected to allow another group to use the DAQ system.

<sup>7</sup>There were 20 cables with signals running under the silicon detectors in the tests done here; the number in the SCT would vary from 4(3) to 24(18) depending on the option chosen for clock and control transmission.

Signal source	operating condition	r.m.s. noise (mV)
LVDS	no signals generated	5.8
LVDS	signals generated, drivers disabled	7.9
LVDS	signals generated and transmitted	7.9
LVDS/2	no signals generated	5.7
LVDS/2	signals generated, drivers disabled	7.5
LVDS/2	signals generated and transmitted	7.8
LVDS, 1 wire shorted to shield	no signals generated	5.8
LVDS, 1 wire shorted to shield	signals generated, drivers disabled	7.7
LVDS, 1 wire shorted to shield	signals generated and transmitted	8.0
LVDS/2, 1 wire shorted to shield	no signals generated	5.6
LVDS/2, 1 wire shorted to shield	signals generated, drivers disabled	7.6
LVDS/2, 1 wire shorted to shield	signals generated and transmitted	7.6

Table 3: Measurements of noise picked up on a scope probe near a cable bundle. See text for explanation.

## 5 Bandwidth/mass optimization

There is significant motivation to minimize the mass in the inner detector volume. This can be accomplished by trading bandwidth for mass. Smaller cables have larger attenuation and therefore a more limited bandwidth. The behaviour of cable attenuation versus bandwidth can be found in the literature[3]. A step function input on one end of a twisted pair cable produces a shape at the other end described by

$$f(t) = \operatorname{erfc} \left( \frac{\kappa z}{2\sqrt{t}} \right)$$

where  $z$  is the cable length and the coefficient of attenuation  $\kappa$  is related to the cable attenuation per unit length.<sup>8</sup> In the low frequency limit  $\kappa$  is given by  $\sqrt{R'C'}$ , the product of the resistance per unit length and capacitance per unit length of the cable.

A simple simulation program incorporating the function given above was used to study the data shown in Figs. 4 and 5. A random stream of bits were generated and used to define an NRZ data stream of low and high states. The output signal followed the function given above between changes in the input state. The output signal was sampled at random intervals of average length 2 ns to mimic the persistence plots obtained with a 500 MHz scope. Random noise of 3% of the maximum signal amplitude was added to the generated points. The simulated traces shown in Fig. 9 correspond to the measurements shown in Fig. 5. The data with AWG28 UTP cables of 50 m length was used to determine  $\kappa$ , which was  $0.08 \sqrt{ns/m}$  for this cable. The expected scaling with cable length is observed.

The value of  $\kappa$  appropriate for the AWG36 cables was extracted using the simulation described above. The eye plot showing 40Mbit/s transmission over 20 m of AWG36 was used to determine  $(\kappa z)^2 = 12 \pm 2$  ns or  $\kappa = 0.173 \pm 0.014$  for this cable (compare Fig. 10 with Fig. 4).

<sup>8</sup>The attenuation per unit length is  $\kappa\sqrt{\omega/2}$

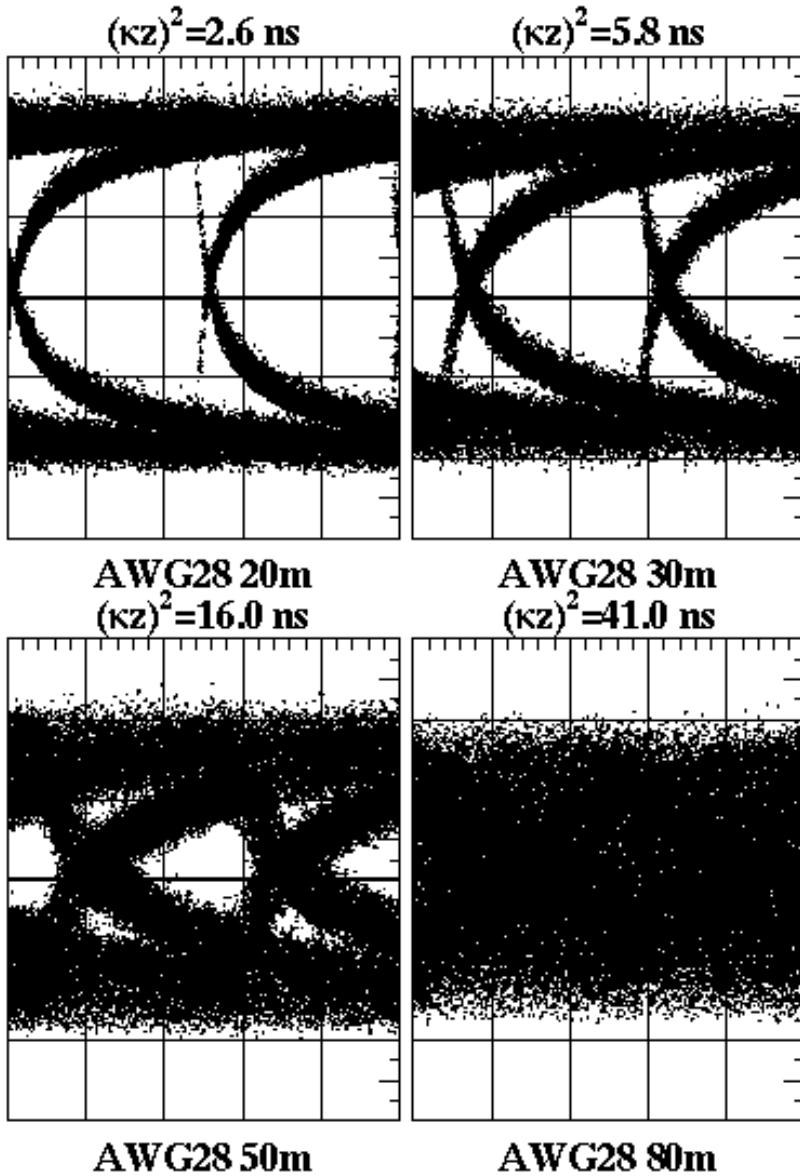


Figure 9: Simulated persistence traces for transmission of random data at 40 Mbit/s over 20 m, 30 m, 50 m, and 80 m of AWG28 UTP cable. The simulation parameters were tuned to reproduce the observed behaviour at 50 m.

A model was developed to allow predictions of  $\kappa$  for other cable sizes. The model was motivated by the low frequency limit,  $\kappa = \sqrt{R'C'}$ , and by the presence of the skin effect at higher frequencies. Firstly, we note that the capacitance per unit length does not change rapidly with cable size, and is taken to be constant;  $\kappa$  is then assumed to vary with the

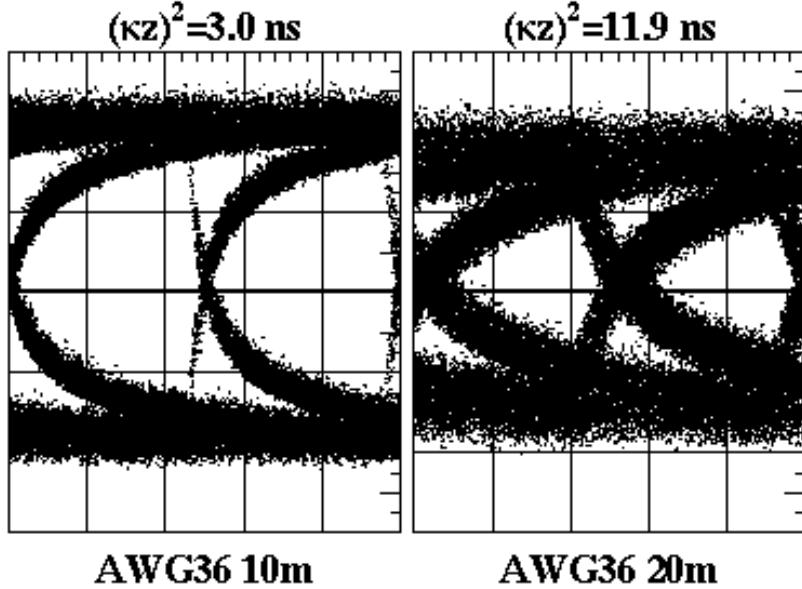


Figure 10: Simulated persistence traces for transmission of random data at 40 Mbit/s over 10 m and 20 m of AWG36 STP cable. The skin depth in the model was tuned to reproduce the plot at 20 m.

effective cross-section of the conductor, taken to be

$$A_{\text{eff}} = 2\pi \int_0^{d/2} r e^{-(d/2-r)/\delta} dr = \pi\delta \left( d - 2\delta(1 - e^{-d/2\delta}) \right)$$

i.e. the area weighted by  $e^{-(d/2-r)/\delta}$ . For cables of diameter  $d_i$  and capacitance per unit length  $C'_i$ , this gives

$$\frac{\kappa_1}{\kappa_0} = \sqrt{\frac{C'_1}{C'_0} \frac{d_0 - 2\delta(1 - e^{-d_0/2\delta})}{d_1 - 2\delta(1 - e^{-d_1/2\delta})}}$$

The skin depth  $\delta$  is treated as an adjustable parameter.

The skin depth  $\delta$  in copper at 20 MHz is about 40  $\mu\text{m}$  and scales as  $\delta \sim 1/\sqrt{\omega}$ . The values of  $\kappa$  cited above for the AWG28 and AWG36 cables give  $\delta = 0.10_{-0.05}^{+0.20}$  mm. While this method is necessarily approximate it allows some quantification of the dependence of the coefficient of attenuation on the cable size.

The eye plots shown correspond to about 5 000 bits of transmitted data to facilitate comparison with the persistence plots recorded with the oscilloscope. Extending these results to BER of interest is discussed in Appendix A.

The experience with the BER tester suggests that signal transmission is more than adequate (including a safety margin) for  $(\kappa z)^2$  up to 10 ns (see Tab. 1 and Figs. 4,5).

The model can be used to predict the signals expected with other cable sizes<sup>9</sup>. Using the parameters fixed from the data shown above yields the eye plots shown in Fig. 11 for 8 m of AWG38 (0.1 mm conductor diameter) and AWG40 (0.08 mm conductor diameter) cables.

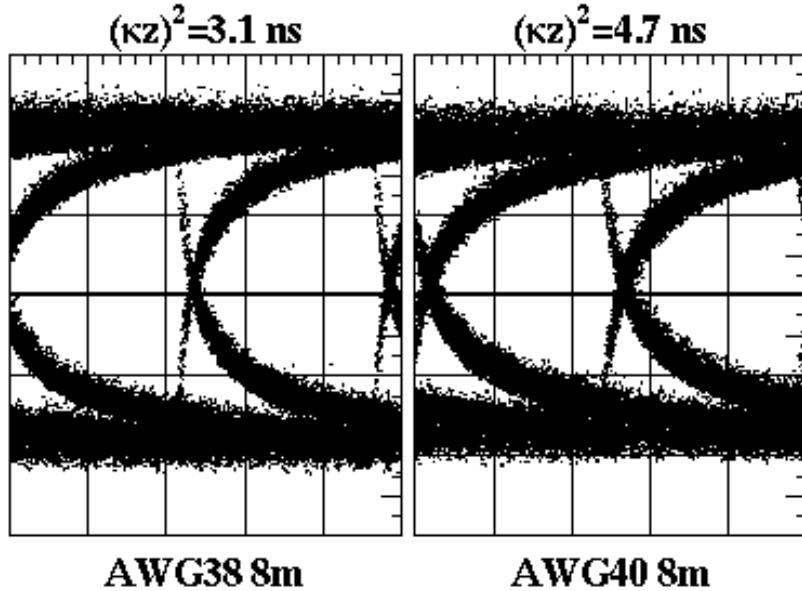


Figure 11: Simulated persistence traces for transmission of random data at 40 Mbit/s over 8 m of AWG38 and AWG40 STP cables.

The calculated dependence of the coefficient of attenuation on bandwidth and cable length is shown in Fig. 12. The uncertainties in the tuning of the model to the AWG36 cables are propagated into the error bars shown. The 8 m of AWG36 and AWG38 cables in the inner detector volume can transmit considerably higher bandwidths than necessary. While the SCT bandwidth per link will remain 40 Mbit/s due to other considerations, higher bandwidths might be of interest for pixel data transmission.

While it appears that cables of size AWG38 or AWG40 would be adequate for the required bandwidth, the feasibility of using these thinner cables must be demonstrated. Habia Cable has given us a drawing for AWG38 cables and will make a test run at our request. The local expert on cable mechanics does not expect problems with AWG38 cables. The use of AWG40 cables is one step further away from current experience and would require some development.

<sup>9</sup>We recently received an estimate from Habia Cable concerning the attenuation of AWG38 STP cables. They estimate the ratio of attenuation of AWG38 to AWG36 to be 1.3; the model used here gives  $1.27 \pm 0.04$ . The Habia number was calculated conservatively as the ratio of conductor cross-sections.

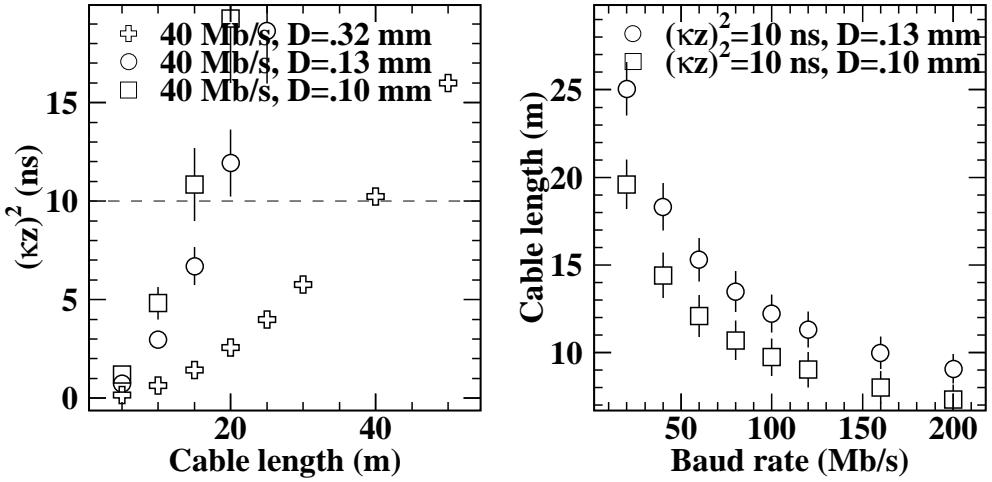


Figure 12:  $(\kappa z)^2$  versus cable length for several conductor diameters. The dashed line indicates where signal reception without compensation becomes marginal. The second plot shows the cable length versus bandwidth with  $(\kappa z)^2$  held fixed at 10 ns; good transmission is possible in the region below and to the left of each curve. The error bars are dominated by the uncertainty in the parameter tuning (see text).

## 6 Required space in cable runs

The number of modules per half-barrel is 1056, and the cables will be brought out in cable trays along the cryostat at each end. The outer diameter of the AWG36 cables is 0.9 mm. A bundle of 27 of these cables was measured to have an outer diameter of 5.5 mm, thus an area of  $23.8 \text{ mm}^2$ ; one would have expected  $18.9 \text{ mm}^2$  for perfect packing (packing factor 0.91). Thus the total cross sectional area required on each side of the barrel will be  $1056 \cdot (23.8/27) \text{ mm}^2 \cdot n = 929 \text{ mm}^2 \cdot n$  where  $n$  is either 3 or 4 depending on whether or not the clock and control links are multiplexed. This is larger than the area currently reserved for the data transmission system. If AWG38 cables (outer diameter 0.8 mm) are used the area required at each end of the barrel will be  $734 \text{ mm}^2 \cdot n$ . The area needed for the forward cables on each end will be  $(988/1056)$  times this number.

Taking the connectors used until now as a model, a 24-pair connector would have dimensions 32 mm long  $\times$  32 mm wide  $\times$  5 mm, or about  $210 \text{ mm}^3$  per pair. Further miniaturization would be desirable for connectors used within the detector volume.

The space required for fibers depends primarily on the fiber protection used.

## 7 Material budget

The material budget for the STP cable system has been calculated for AWG36 and AWG38 cables. The number of radiation lengths is given in Tab. 4 for AWG36 cables and in Tab. 5 for AWG38 cables. In these tables the column labelled A/X0 is divided by the nominal module width of 64 mm to obtain the X/X0 given in the next column. The last column must be multiplied by the number of cables at a given position (there will be 3 or 4 cables per module, and 6 modules per half-barrel).

item	diameter (mm)	material	X0 (mm)	A (mm <sup>2</sup> )	A/X0 (10 <sup>-4</sup> mm)	X/X0/pair /module (10 <sup>-4</sup> )
signal wires (2)	o.d. 0.13	Cu	14.3	0.026546	18.56	0.290
drain wires (4)	o.d. 0.065	Cu	14.3	0.013273	9.28	0.145
dielectric (2)	i.d. 0.13	HFI 150 (polyamide)	440	0.200277	4.55	0.071
inner shield	i.d. 0.38	Al	89	0.052778	5.93	0.093
outer shield	i.d. 0.82					
	o.d. 0.86					
	i.d. 0.86	Polyester	400	0.055292	1.38	0.022
	o.d. 0.90					
total				0.348166	39.70	0.620

Table 4: Calculation of material in AWG36 cables.

item	diameter (mm)	material	X0 (mm)	A (mm <sup>2</sup> )	A/X0 (10 <sup>-4</sup> mm)	X/X0/pair /module (10 <sup>-4</sup> )
signal wires (2)	o.d. 0.10	Cu	14.3	0.015708	10.98	0.172
drain wires (2)	o.d. 0.065	Cu	14.3	0.006637	4.64	0.073
dielectric (2)	i.d. 0.10	HFI 150 (polyamide)	440	0.145142	3.30	0.052
inner shield	i.d. 0.32					
	i.d. 0.73	Al	89	0.047124	5.29	0.083
outer shield	i.d. 0.77					
	o.d. 0.77					
	i.d. 0.77	Polyester	400	0.049637	1.24	0.019
	o.d. 0.81					
total				0.264248	25.45	0.399

Table 5: Calculation of material in AWG38 cables.

The assumed cable routing is as shown in Fig. 13. The number of radiation lengths in the cables as seen by a particle coming from the interaction region (including the effect of the longitudinal beam spread) through to the last TRT hit is shown in Fig. 14 for a system with 4 AWG36 STP cables per module<sup>10</sup>. Fig. 15 shows how much of this material is in front of the last SCT hit. A comparison of different STP link options (4 AWG36 cables, 3 AWG36 cables + DORIC, and 3 AWG38 cables + DORIC per module) is shown in Fig. 16.

<sup>10</sup>The plot actually shows the material for an earlier cable type which had a slightly larger drain wire and thus 8% more material than is given in Tab. 4

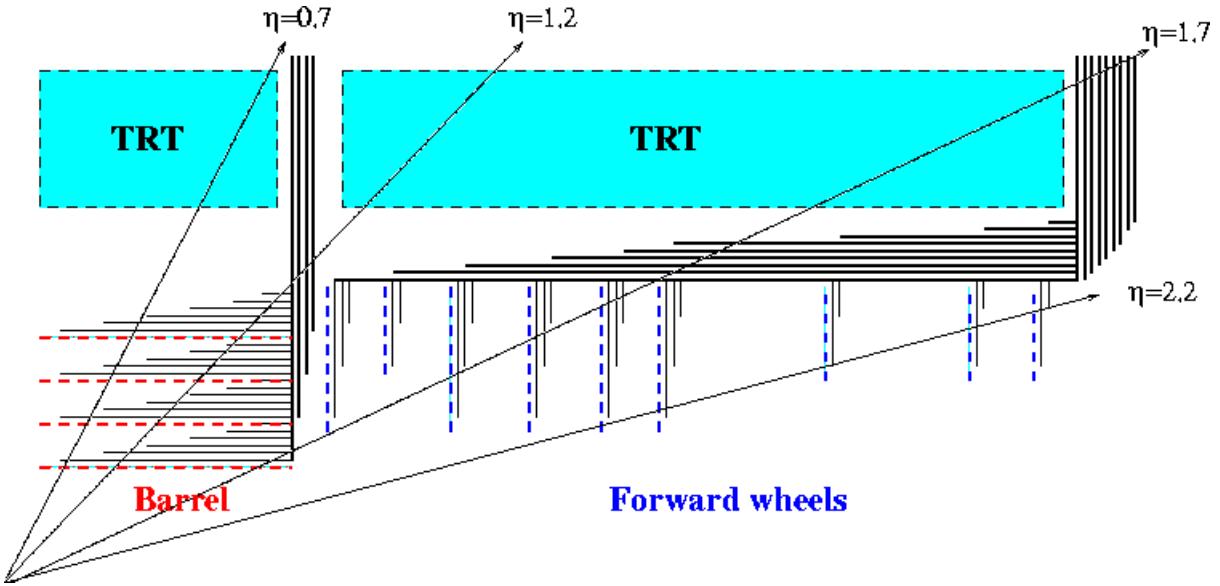


Figure 13: The routing of SCT data and C & C cables assumed for the calculations of the associated material budget is shown.

The effect of adding protection to the optical fibers is shown in Fig. 17 for two different assumptions on fiber protection, namely a minimal mass fiber ribbon and a fairly heavy (1 mm outer diameter) individual fiber sheath. In each case the material of the additional control lines (PIN bias and control line for LDC driver current) is included<sup>11</sup>, but LED driver power, which is derived from the module (176 mW compared to 5.8 W/module), is not. A comparison between STP and optical links is shown in Fig. 18.

Fig. 19 shows the material budget given in the TDR; the material up to the last TRT hit is included. The STP cables in the worst region,  $\eta = 2.1$ , contribute  $X/X_0 = 1.3\%$  to 2.5% depending on the success of the material reduction effort, to be compared with an existing 16% in other SCT related material. A more typical figure ( $\eta = 1.7$ ) has contributions from SCT of  $X/X_0 = 0.8\%$  to 1.5% to be compared with the overall SCT contribution of 20%.

The material budget for SCT cable use by the pixels is subject to considerable uncertainty due to the as yet not well defined pixel data transmission system. After consultation with our pixel colleagues[5] we estimate the material in the pixel data transmission system to be as shown in Fig. 20 if AWG36 cables are used at 40 Mbit/s. Using AWG38 cables and 80 Mbit/s could lower reduce the material by a factor of 2.

---

<sup>11</sup>We estimate 4% of the flat cable described in the TDR could be saved by eliminating these two lines.

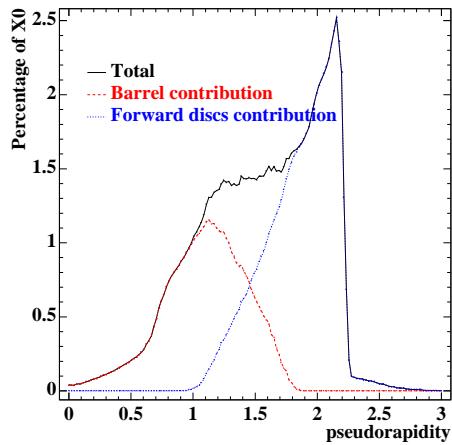


Figure 14: Links system material in front of the last TRT hit for 4 AWG36 cables per module. The contributions of barrel and forward cables are indicated separately.

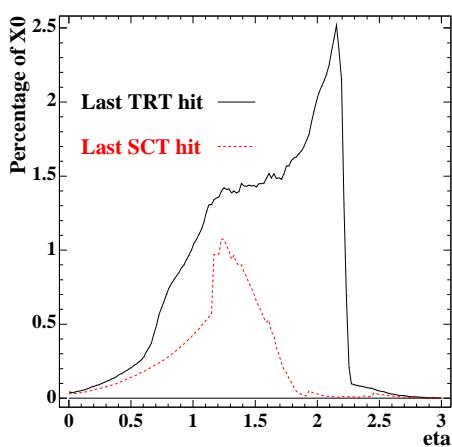


Figure 15: Links system material in front of the last SCT and TRT hits for 4 AWG36 cables per module.

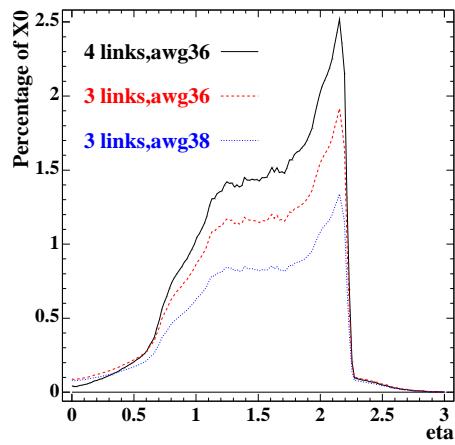


Figure 16: Links system material in front of the last TRT hit for different STP options.

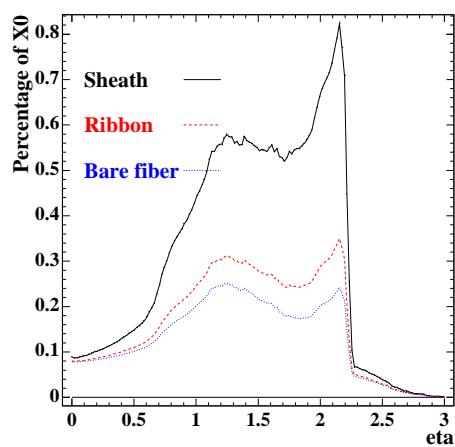


Figure 17: Links system material in front of the last TRT hit for different fiber protection options.

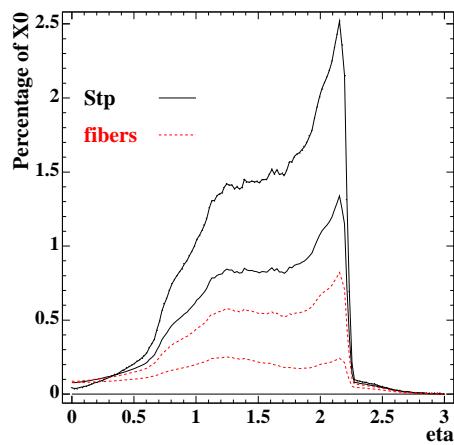


Figure 18: Links system material in front of the last TRT hit for two STP options and two fiber options. The options shown indicate roughly the range of material budget depending on the detailed choice of STP cable or fiber protection.

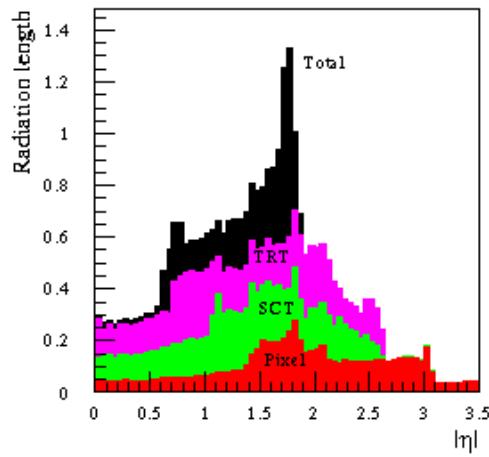


Figure 19: Total inner detector material within the sensitive tracking volume (to the last TRT hit). Taken from the TDR.

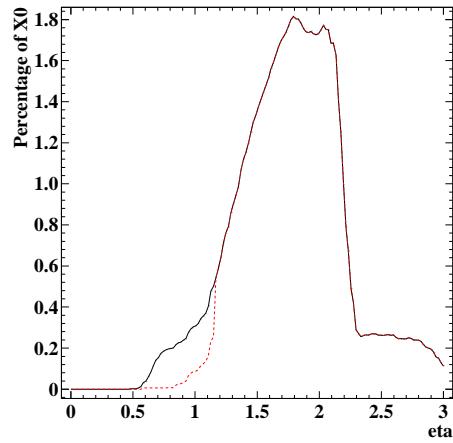


Figure 20: Material in the sensitive volume in a pixel data transmission system using AWG36 cables at 40 Mbit/s. The solid curve shows the material before the last TRT hit, the dashed curve shows the material before the last SCT hit. Using finer cables or higher transmission speeds would save material.

## 8 Cost

A preliminary costing is given in Tab. 6. The unit costs for the STP system come essentially from the TRT group. The cost of a second repeater is included; the price of the radiation hard repeater is used for both repeaters to be conservative. The costs of the opto-electronics associated with the fiber optic system are rather uncertain. The numbers shown are probably underestimates. The total cost difference between the two systems (for 5000 modules) at least 650 KCHF and could exceed 1 MCHF.

Item	(LED-PIN)	(CHF)	(STP cable)	(CHF)
low-mass STP cable			4*8m	48
AWG36 connector+mounting			1	8
AWG28 connector+mounting			4	14
radiation hard fiber	8m	3		
fiber connectors		28		
connection to ROD (80m) including protection, fiber bundles	standard fiber	58	standard STP	36
sub-total		89		106
drivers	Oxford	20??	2 repeaters	< 37
emitters	GEC LEDs	80	(1 rad-hard)	
PIN diodes	GEC PIN diodes	35		
DORIC	DORIC	10??		
Opto-boards + mounting		10??		
sub-total		155??		< 37
Off-detector driver/receiver/PIN		30??		
TOTAL		> 274??		< 143

Table 6: Costing for optical fiber and STP systems. These prices are per module and do not include contingency. Costs are not included for patch panels. Non-recoverable engineering costs are not included.

Notes for the LED-PIN option:

- Price for radiation hard fiber = 0.13 CHF/m (quote); fiber protection for 80 m runs estimated to cost 0.11 CHF/m.
- Connectors: 4 MT-12 plus 2 MPO-12 connectors per 12 fibers. This includes all connections and terminations from opto-hybrid to ROD. Labor and assembly connector assembly costs are included.
- Price estimates for ASICs and optoboards are very uncertain.

Notes for the STP option: Assumes 4 pairs/module in bundles of 24 pairs. Assumes 5 50-pin STP connectors (1 for AWG36 and 4 for AWG28), 2 repeaters.

- CHF 30 labor to mount 24-pair AWG36 cables into connector,
- CHF 10 for AWG36 connector PC board,
- CHF 9.85 per 50 pin connector,

- CHF 5 to mount 24-pair AWG28 cables,
- CHF 16 per 4-channel LVDS repeater chip,
- CHF 15 for PC board to mount repeaters for one 24-pair cable,
- CHF 2.7/m for 24-pair AWG28 cables,
- CHF 1.5/m per AWG36 cable. If AWG38 were used the price might be slightly higher (e.g. CHF 2.1/m, but this is not a firm quote).

## 9 Conclusions

The main issues concerning the use of fine STP cables for the SCT data transmission system have been studied. The main findings are as follows:

- The bandwidth available on AWG36 cables is more than adequate over the distance required (8 m). The prospects for using thinner cables (AWG38) are good. Further studies are needed to determine the feasibility of multiplexing the clock and trigger/control signals.
- Noise radiated from LVDS signals sent through the cables is not a problem. Noise pickup on the cable shields must be controlled. Noise due to eventual ground loops in the full SCT implementation cannot reliably be studied with a small bench-top system.
- The material budget of the STP cables is significantly (probably a factor of 5 or more) above that of an optical fiber system. In absolute terms the number of radiation lengths added by the STP cable system is typically 1.5% for  $1 < \eta < 2$ . There is potential for reducing this to 0.8% by using thinner cables and multiplexing.
- The space required in the cable runs is greater for STP cables than for fiber ribbons. The cables from each SCT module would require a cross-sectional area of 3.5(2.6) mm<sup>2</sup> assuming 4(3) cables per module. These figures go down by about 12% if AWG38 cables are used.
- The STP system has a significant cost advantage over optical fiber systems primarily because no radiation hard opto-electronics are required. The cost savings on the full SCT system would be upwards of 650 KCHF.

The constraints the pixel community must satisfy with their data transmission system are even more stringent than those of the SCT. The STP cables described here could be used to ensure a radiation hard system and keep costs down; the potential for additional noise problems due to pickup and ground loops, however, are the same as for the SCT. For the pixel application there is significant motivation to increase the baud rate beyond 40 Mbit/s, and the studies done here indicate that 80 Mbit/s appears to be feasible on both AWG36 and AWG38 cables. Using a system optimized for minimal mass, the STP contribution to the material budget could be kept below 1%.

## 10 Acknowledgements

The work presented here is based on the already well advanced work on TRT data transmission undertaken by Philippe Farthouat, Peter Lichard and Arno Beer. Many of the tools used in this study were provided by them along with valuable advice. The electronics group at the University of Geneva has provided excellent support. The authors appreciate the work of Christian Couyoumtzelis in teaching us how to use the DAQ setup for the noise tests on the CSEM module.

## 11 Appendix A

Estimating the BER achievable with a given attenuation requires the simulation of a large number ( $10^9$  or more) of transmitted bits. It becomes rather tedious and time consuming to produce each point as was done in producing the plots shown in this note. An alternate method for estimating the behaviour for large numbers of transmitted bits is to consider the longest series of consecutive '0' bits which is statistically likely within the total number of transmitted bits, since the first '1' bit after a long string of zeros is the hardest to correctly identify. This limiting curve is obtained by finding the smallest  $n$  such that the probability of  $n$  consecutive bits of one type is less than the desired BER. This probability is given by

$$P(n) = 2^{-n+1}$$

or

$$n = \text{nint}(1 - \ln(\text{BER})/\ln(2))$$

We find that the difference between the limiting curves for BERs of  $10^{-4}$  and  $10^{-9}$  is very small. This validates the experience obtained with the BER tester, namely that the transmission either fails on a very short time scale (fractions of a second) or remains good for days.

## 12 References

### References

- [1] David Campbell, private communication.
- [2] David White, private communication.
- [3] "Lossy Transmission Lines" by Fred E. Gardiol, published by Artech House, Inc. (1987).
- [4] Arno Beer, private communication.
- [5] Martin Holder and Kevin Einsweiler, private communication.