S-LINK, a Data Link Interface Specification for the LHC Era

H.C. van der Bij[1], R.A. McLaren[1], O. Boyle[1,2] and G. Rubin[3]

[1]ECP-Division, CERN, CH-1211 Geneva 23

[2]Edinburgh University, Mayfield Road, Edinburgh EH9 3JZ

[3]KFKI-RMKI, PO-Box 49, H-1525 Budapest

ABSTRACT

In the Technical Proposals for the ATLAS, CMS and ALICE experiments at CERN, there is a requirement for several thousand data links. Although there is an obvious need for standardisation, this seems difficult to achieve as the links run at different speeds, over different distances and have various constraints of power consumption, size and radiation hardness. An additional complication is that today we cannot decide which will be the most cost effective technology for the implementation of the final links. Furthermore, we must allow designers of boards at each end of the link, for example read-out electronics and input buffers, to work in parallel with the development of the links.

The S-LINK is a new concept which should provide the benefits of standardisation without these limitations. The S-LINK specification defines, at both ends of the link, a simple FIFO-like user interface which remains independent of the technology used to implement the physical link. The physical link provides transfer of event data and control words, error detection, optional flow control and test facilities. This paper describes the S-LINK specification and gives examples of the use of the S-LINK, the physical links being designed, and the test equipment that is being developed.

I. INTRODUCTION

The ATLAS, CMS and ALICE experiments will need several thousand data links that will transfer on-line physics data. Many of those links have different requirements. For example, some links will be short, such as 25 m, which allows (cheap) electrical transmission, others need to be up to 200 m, which calls for transmission over fibre optic media. Although most of the links need a throughput in the order of 100 MByte/s in which the latency is less important, other links need only a low speed but instead they need transmission with a very low latency. Other constraints that come into play are power consumption, size, radiation hardness and, of course, price. In many cases a cheaper, reduced functionality link (e.g. not rad-hard or lower speed) is sufficient for prototype systems, while for the final system full functionality is needed.

Those different requirements make it very difficult to standardise on a physical link. Also, the components chosen will be obsolete by the time we have to start producing in large quantities for the ultimate construction of the LHC experiments. On the other hand there is a need for standardisation as the detector designers wish to start tests in the relatively near future.

To circumvent these conflicting requirements we have defined the S-LINK specification [1]. The S-LINK is a new concept which should provide the benefits of standardisation without the limitations (Figure 1). The idea is that, unlike other standards that define the physical layer, we define a simple synchronous FIFO-like interface on which the use of the signals remains independent of the technology used to implement the physical link. The mapping of the S-LINK signals to the protocol used on the physical link is left open to the link designers so they can map it in the most suitable way for the technology used. Both Link Source Cards (LSC) and Link Destination Cards (LDC) will be designed by one group. With this approach the link, including both link cards, can be treated as a component in the system with the advantage of having well defined inputs and outputs. This

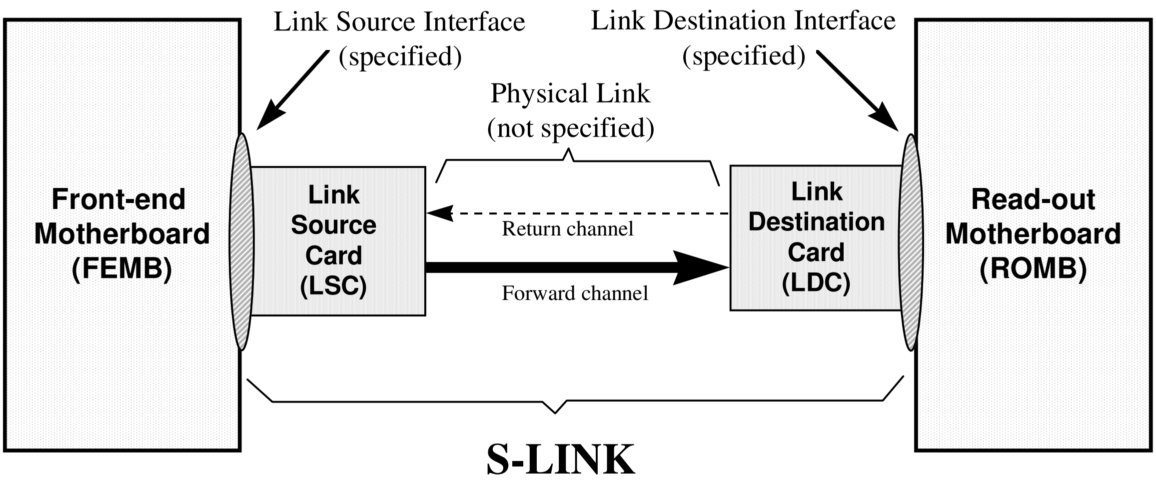

Figure 1: The S-LINK concept

means that various implementations of S-LINK can be developed using new technologies, independently of the design of Front-end Motherboards or Read-out Motherboards. The link users may change from one technology to another without the need for any modifications on the Motherboards.

The S-LINK was designed for moving data from the first level of electronics after the read-out of detectors to the next level. However, the specification is that general that it may be used in many other applications that need an easy-to-use interface to transfer data from one point to another.

II. S-LINK SPECIFICATION

A. *User Requirements*

Based on prior experience with links in several LEP experiments, the Fibre Optic Link in the NA48 experiment [2] and discussions with potential users of links in the LHC experiments, we made a list of user requirements that are common for the different data links.

These requirements are that the link should:

- detect errors introduced in the link
- be able to send data words and control words
- handle any size of event data and event headers
- provide a self-test function
- transfer with a low latency (< 1 μsec)
- provide flow control (not needed for all users)
- provide a return channel for slow control information (not needed for all users)
- be integrated easily in a dataflow system

B. *Hardware Requirements*

The users demanded that the amount of logic needed to interface to the link should be minimal. This means that the complexity of the link should be pushed into the link cards. As in most cases the DAQ electronics has a pipeline based dataflow organisation, a FIFO-like interface seemed the right choice for the interface to the Front-end electronics and Read-out electronics. To reduce cost and power consumption, the link cards should not buffer any significant amount of data; ten words being considered the maximum. The concept that the physical links may easily be changed, calls for an implementation of the S-LINK as a daughter card.

C. *S-LINK Mechanics*

To allow interchangeability of different implementations of S-LINK cards, the specification recommends the use of small daughter boards. A single, high density, 64-pin connector connects the Front-end Motherboard to the Link Source Card and the Link Destination Card to the Read-out Motherboard. Both the format of the daughter board and the type of connector used, correspond to the IEEE Common Mezzanine Card standard [3], which is the same format as used by PCI Mezzanine Cards. This daughter card implementation should be used during the prototyping phases, while for the final systems the link logic may be integrated onto the motherboards.

D. *Link Card Interface*

1) *Link Source Interface*

Figure 2 shows the signals that a Front-end Motherboard should connect to the Link Source Card. It starts with two lines that can be used to select the data width. These lines can be set to 8, 16 or 32 bits and will have a fixed setting for each application. The CNTRL/DATA line allows the user to send control words as opposed to normal data. The DATA lines make up a 32-bit word that must respect the set-up and hold times with respect to the CLOCK signal.

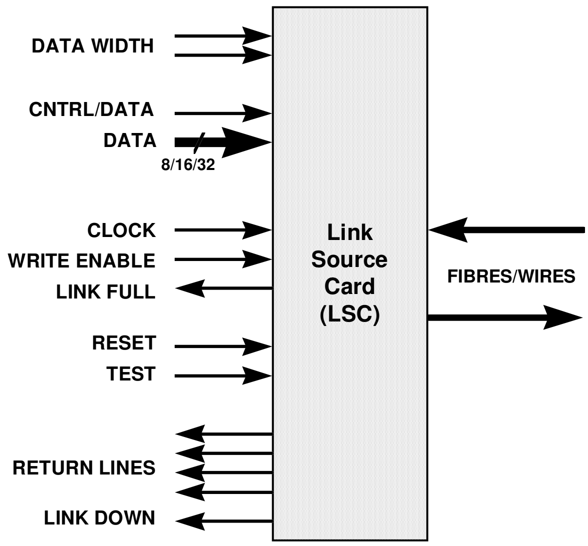

Figure 2: Link Source Interface signals

Data will only be written when the WRITE ENABLE is active. The signal handling corresponds exactly to that of writing to a synchronous FIFO. LINK FULL is generated when the Read-out Motherboard can no longer accept data or when the Front-end Motherboard is sending data faster than the link can handle. This should cause the Front-end Motherboard to suspend the data transfers, but for ease of interfacing the Front-end Motherboard is allowed to send one more word. RESET is used to initialise the link. When TEST is asserted the LSC enters the test mode and generates a predefined data pattern. At the other end of the link the LDC recognises the test mode and begins checking the incoming data words. The RETURN LINES allow the Read-out Motherboard to send a few user-defined control signals back to the Front-end Motherboard. Finally the LINK DOWN signal is asserted when there is a major problem with the link, for example when there is no light in the fibre. The flow control via LINK FULL and the RETURN LINES are not available on the simplex version of the S-LINK, in which the Link Source Card and Link Destination Card are connected with a physical link going in one direction only.

2) *Link Destination Interface*

The signals arriving at the Read-out Motherboard (Figure 3) are almost a copy of those transmitted at the Front-end Motherboard. The data width lines to the LDC should have the same setting as on the LSC. The CNTRL/DATA bit informs the Readout Buffer if normal data is received or that a control word is received. DATA ERROR is asserted if an error has taken place during the transfer, either on a word, or an event; dependent on the link type, the Link Source Card will transfer parity on each word transmitted or will send a CRC calculated over the data encapsulated between control words.

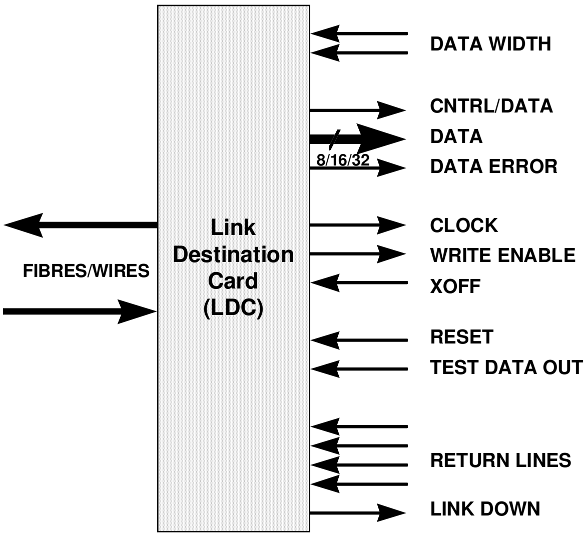

Figure 3: Link Destination Interface signals

The CLOCK and WRITE ENABLE signals provide the timing for the received signals. XOFF is used by the Readout Buffer to indicate that it can no longer accept data. This set of signals is selected so that the LDC-signals DATA, CNTRL/DATA and DATA ERROR can be connected straight into the datalines of a synchronous FIFO. The CLOCK and the WRITE ENABLE lines will connect to the corresponding pins of the FIFO, while the XOFF line may be connected to the half full flag or almost full flag pin. RESET initialises the link and TEST DATA OUT allows the data generated in test mode to be output on the DATA lines.

The RETURN LINES are used to send control bits back to the Readout Driver. LINK DOWN indicates a serious problem on the S-LINK. The flow control via XOFF and the RETURN LINES are only available on the duplex version of the S-LINK.

E. Data Format

Each application may define its own data format as the S-LINK specification does not put any constraints on it. A format commonly used is encapsulating event data by two control words. The S-LINK Implementation Guide [4] does suggest to reserve certain bits in the control words to distinct the beginning and end of an event.

III. S-LINK PROJECTS

A. S-LINK Link Cards

Link Source Cards and Link Destination Cards are being developed in research institutes outside CERN, while companies will be approached for the large scale production and support of the cards. Table 1 shows the developments of link cards that are ongoing, each one based on a different technology. All five designs can handle 32-bit data transfers. They are duplex implementations of the S-LINK which means that flow control and return lines are available.

The Fibre Channel implementation [5] uses standard GigaLink Modules which are available from several manufacturers. The MATCH version uses a chip developed by INFN, implemented in GaAs technology which is radiation hard. Optobus is a component available from Motorola, which has integrated optics for driving and receiving a 10-bit parallel ribbon fibre. The G-LINK is a proven, reliable component from Hewlett Packard which can serialise and de-serialise data and transfer over a co-axial cable. The low voltage differential signalling (LVDS) implementation is a low speed but inexpensive solution meant for prototyping systems. It uses a standard SCSI cable.

B. S-LINK Test Equipment and Interfaces

To help Motherboard designers carry out the first tests of their S-LINK interfaces, a dummy destination and a dummy source have been built. The dummy destination, SLIDAD, replaces the Link Source Card using the same connector and form factor (Figure 4). Words written to the link are displayed on LEDs with a single step function. Logic state analyser connections are provided for measurements at full speed. The reaction of the Front-end Motherboard logic to signals from the link, for example LINK FULL, can also be verified. At the Read-out Motherboard a similar module,

Table 1.
Link Card developments

Institute	Medium	Technology	Speed	Status
RMKI, Budapest	Serial Fibre	Fibre Channel [6]	103 MByte/s	Finished
INFN, Rome	Serial Fibre	MATCH [7]	100 MByte/s	Being built
Univ.of Stockholm	Parallel Fibre	OPTOBUS [8]	150 MByte/s	Under test
KEK, Tsukuba, Japan	Serial Electrical	G-LINK [9]	125 MByte/s	Under test
INP, Cracow	Parallel Electrical	LVDS [10]	50 MByte/s	Being built

SLIDAS, can be used to emulate an S-LINK input by generating data with selectable patterns and lengths. Another design, based on the SLIDAS and SLIDAD, is made to provide a stand-alone test bench for Link Cards.

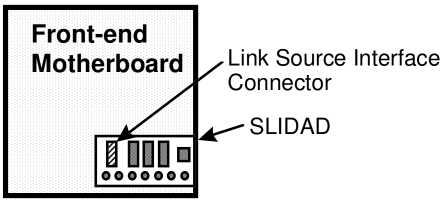

Figure 4: Front-end logic under test with SLIDAD

Once the basic functions of the Front-end Motherboard have been tested, the designer would normally like to read the data via VME-bus. The recommended method is to use a VME-bus processor board with a PCI Mezzanine Card (PMC) site, for example a Motorola MVME1600 or a CES RIO2 8060. An S-LINK to PMC card has been built for this purpose and a LynxOS user-level library and test software has been written (Figure 5). The S-LINK to PMC interface has a maximum transfer rate of 65 MByte/s.

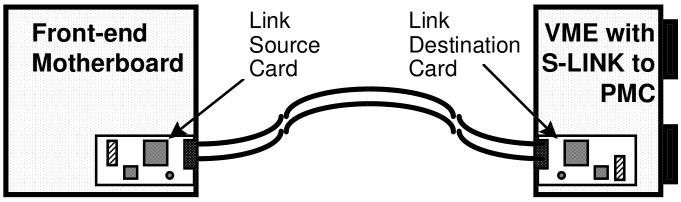

Figure 5: Front-end logic in application

A PMC to S-LINK interface also exists. This one may be used to send data to a data acquisition system. Also this interface has a throughput of 65 MByte/s. To allow the use of standard PCs for testing of links and of Read-out Motherboards, a PCI form-factor of the S-LINK interfaces is being designed as well.

To test a data acquisition system without having front-end electronics or a detector, an S-LINK daughter module for the SLATE [11] is designed. With the SLATE, one can send emulated event data with programmable delays between events. One SLATE can drive three Link Source Cards with a total bandwidth of 100 MByte/s over each link if they all send the same data (Figure 6). The software can handle up to ten SLATEs, which may all be synchronised.

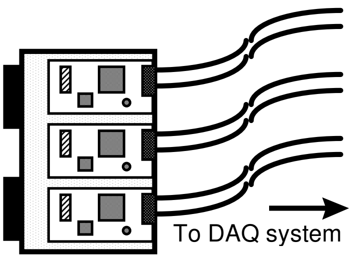

Figure 6: SLATE with Link Source Cards emulating front-end electronics

C. Development Support

To support users and designers of S-LINK cards, help is offered in different ways. An implementation guide [4] is available which clarifies certain issues and design choices that have to be made. For board-level simulations, VHDL models of the different S-LINK components are available. Also personal assistance is given for detailed technical issues. To inform users of the state of the different developments, information is distributed via the World Wide Web [1]. Additionally each S-LINK product has a page on which an overview of the product, the status of the development and all documentation can be found.

IV. APPLICATIONS

ATLAS is the first experiment to adopt the S-LINK (Figure 7). For the next few years it will use the S-LINK as de-facto standard for the link between the Read-Out Driver (ROD) and the Read-Out Buffer (ROB) [12]. The ROD multiplexes the data coming from the links from the detectors and moves it to the ROB, where the data are stored until a level two decision is made.

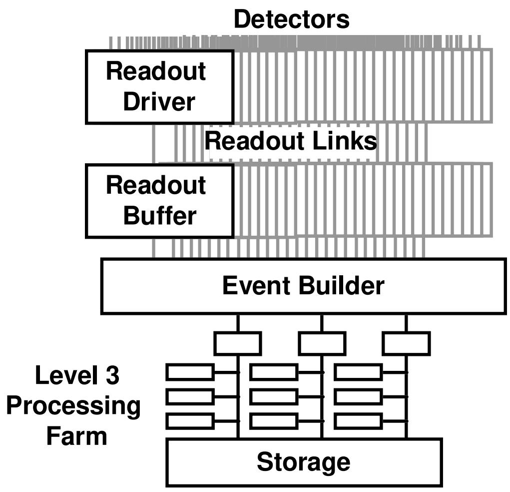

Figure 7: S-LINK used for Readout Links in ATLAS

Originally the S-LINK was developed for the use of read-out links in the DAQ system. However, the low latency aspect of it attracted the ATLAS Trigger Supervisor developers to use it for the transfer of trigger information. This means that both the links from level 1 to the Trigger Supervisor (which transfers the Region of Interest information) and from the Supervisor to the Read-Out Buffer will be implemented in S-LINK.

Also the ATLAS Level 2 Trigger Demonstrator Architecture A and the DAQ Prototype-1 with the Royal Holloway Read-out Buffer will use S-LINK. In addition engineers from the ALICE collaboration are designing an S-LINK implementation using Fibre Channel components and will evaluate if the S-LINK specification meets the ALICE requirements.

V. FUTURE

Currently all efforts in the S-LINK project are focused on finalising and testing the different physical link implementations. All S-LINK developments will be made commercially available to ensure availibility and support. To ease successful integration, help will be given to projects needing S-LINK. Higher speed PMC and PCI to S-LINK interfaces will be developed at CERN while S-LINK products developed outside CERN will be tested for compliance and reliability.

The experience that will be built up in the following two to three years, will be used to verify that the S-LINK specification satisfies the LHC requirements. By evaluating how the S-LINK is used in the different prototype systems, we possibly can restrict the number of options in the S-LINK. If a real need appears for new features, it is likely that those could be incorporated in the S-LINK specification.

VI. CONCLUSIONS

A specification for a data link interface has been written in which the link can be considered as a single component in the system. The S-LINK specification describes, at both ends of the link, a set of signals over which an easy-to-use synchronous FIFO-like data movement protocol is defined. A return channel for flow control and four return line signals are included in the duplex version of the S-LINK. In addition to transmission of event data - including optional headers and control words - an S-LINK compliant link also provides error detection and self-test functions. Unlike other standards, the S-LINK specification does not impose the physical layer of the link. It can therefore incorporate new link chips that will appear on the market. While it is foreseen to integrate the link logic on the motherboards for the final designs of the read-out electronics, an IEEE P1386 CMC form-factor board and a connector are suggested for plug-in implementations of the S-LINK. Five different physical links are being developed, while test tools, PMC and PCI interfaces are available. To ease integration, technical support is available for S-LINK users and designers. The ATLAS experiment [13] will use S-LINK in a variety of prototype DAQ systems, while ALICE is building an S-LINK card and is evaluating if S-LINK meets their requirements.

VII. ACKNOWLEDGEMENTS

The authors would like to thank the many members of the ATLAS experiment that have helped shape the S-LINK specification and to thank especially Nick Ellis, Philippe Farthouat, Livio Mapelli and Michal Turala for their support and encouragement.

VIII. REFERENCES

[1] Boyle et al., “The S-LINK Interface Specification”, ECP-Division CERN, Geneva, Switzerland, 29 May 1996, also available at http://www.cern.ch/hsi/s-link

[2] Ph. Brodier-Yourstone et al. “A 10 MByte/s Fibre Optic Link”. IEEE Transactions on Nuclear Science, Vol. 42, No. 4, August 1995. ISSN 0018-9499.

[3] IEEE P1386 Standard Mechanics for a Common Mezzanine Card Family: CMC, Institute of Electrical and Electronics Engineers, Piscataway, NJ 08855-1331, USA

[4] H.C. van der Bij, “S-LINK Implementation Guide”, ECP-Division CERN, Geneva, Switzerland, 19 June 1996, also available at http://www.cern.ch/hsi/s-link/imp_guide

[5] G. Rubin et al., “Development of FCS-Link Interface Cards”, Presented at the second Workshop on Electronics for LHC, Balatonfüred, Hungary, 23-27 September 1996.

[6] Fibre Channel - Physical and Signalling Interface (FC-PH), Rev.4.3. Proposed ANSI standard dpANS X3.230-199x, ANSI, New York, June 1994.

[7] G. Mirabelli, E. Petrolo, E. Valente, R. Cardarelli, R. Santonico and M.L. Ferrer: “A GaAs Transceiver Chip for Optical Data Transmission”, IEEE Trans. on Nuclear Science, Vol.40, No. 4, August 1993.

[8] PC94DL0400: Optobus I Multichannel Optical Data Link, Datasheet, Motorola, Phoenix, Arizona USA, June 1996.

[9] HDMP-1000: Gigabit Rate Transmit Receive Chip Set, Datasheet, Hewlett Packard, Palo Alto, CA USA, January 1993.

[10] Low Voltage Differential Signaling Introduction, Lit# 103288-001, National Semiconductor, Fürstenfeldbruck, Germany, 1995.

[11] E. Denes et al., “SLATE on-line program for trigger and DAQ test systems”, Presented at the second Workshop on Electronics for LHC, Balatonfüred, Hungary, 23-27 September 1996.

[12] R.A. McLaren et al., “An application of S-LINK, a data link interface specification, in the ATLAS readout system”, Presented at the second International Data Acquisition Workshop (DAQ96) on Networked Data Acquisition systems, RCNP Osaka, JAPAN, 13-15 November, 1996.

[13] ATLAS Technical Proposal, CERN/LHCC/94-43, LHCC/P2, CERN, Geneva, Switzerland, 15 December 1994, ISBN 92-9083-067-0