

High Level Design of Gigabit Ethernet S-LINK LSC

Authors : J. W. Dawson, D. Francis, S. Haas, J. Schlereth
Keywords :

Abstract

NoteNumber :

Version : 1.3

Date : 27.10.2001

Reference :

Document Change Record

Table 1. Document Change Record.

1. Document Title: High Level Design of Gigabit Ethernet S-LINK LSC			
2. Document Reference Number			
4. Issue	5. Revision	6. Date	6. Reason for change
<i>1</i>	<i>0</i>	<i>9 June. '01</i>	<i>Birth.</i>
<i>1</i>	<i>1</i>	<i>19 Oct. 01</i>	<i>Additional text, comments & corrections from John</i>
<i>1</i>	<i>2</i>	<i>25 Oct. '01</i>	<i>Corrected Spelling mistakes & cleaned up references</i>
<i>1</i>	<i>3</i>	<i>27 Oct. '01</i>	<i>Cost table added and comments from Stefan</i>

1 Introduction

This document presents the high level design of a prototype Gigabit Ethernet S-LINK Link Source Card (LSC).

2 Background

The Trigger and DAQ community aims to maximise the use of Commercial Off The Shelf (COTS) products to minimise the overall purchasing, maintenance and upgrade cost of the Trigger and DAQ system implementation.

Networking “trends”, particularly with respect to multi-gigabit Ethernet, indicate that dedicated processing power is being deployed on network interface cards to enable high speed links to be driven at line speed (greater than one gigabit per second bandwidths). These network interface cards could be deployed at the receiver end of a Read-Out Link (ROL), i.e. the ROS, thus increasing the use of COTS components within the ATLAS Trigger and DAQ.

To date the RODs within ATLAS are being designed to provide data to a ROL according to the S-LINK interface specification [1], therefore for COTS network interfaces to be deployed within the ROS a Gigabit Ethernet implementation of the S-LINK LSC must be designed and implemented. The design, implementation and testing of a prototype Gigabit Ethernet S-LINK LSC is the aim of this task.

The design presented is not novel, it aims to employ hardware and software techniques which have already been used at CERN [2], with the aim of demonstrating proof of principle on a minimum time scale, i.e. rapid prototyping.

3 High Level Design

3.1 Required functionality

The card must provide the following functionality:

- receive a stream of words in S-LINK format, i.e. 32 bit data and control words, via the standard connector.
- output the data words as one or more Ethernet frames, i.e. the control words are removed.
- map between the Ethernet flow control signals and the S-LINK flow control signals
- map between the Ethernet “link down” signal and the S-LINK link down signal.

3.2 Block diagram

A block diagram of the Gigabit Ethernet LSC is shown in Figure 1:. The input conforms to the S-LINK protocol with regard to reset, link down, and flow control. The data stream, including the control bits, are written to a FIFO. The FPGA reads words from the FIFO and the Beginning Of Fragment (BOF) control word is used to provide the Transmit Start of Frame (TXSOF). The End Of Frame (EOF) control word is used to provide the Transmit End of

Frame (TXEOF). The FPGA writes the 32 bit data words, together with TXSOF and TXEOF, to the Gigabit Ethernet MAC. The S-LINK control bits are lost in this process.

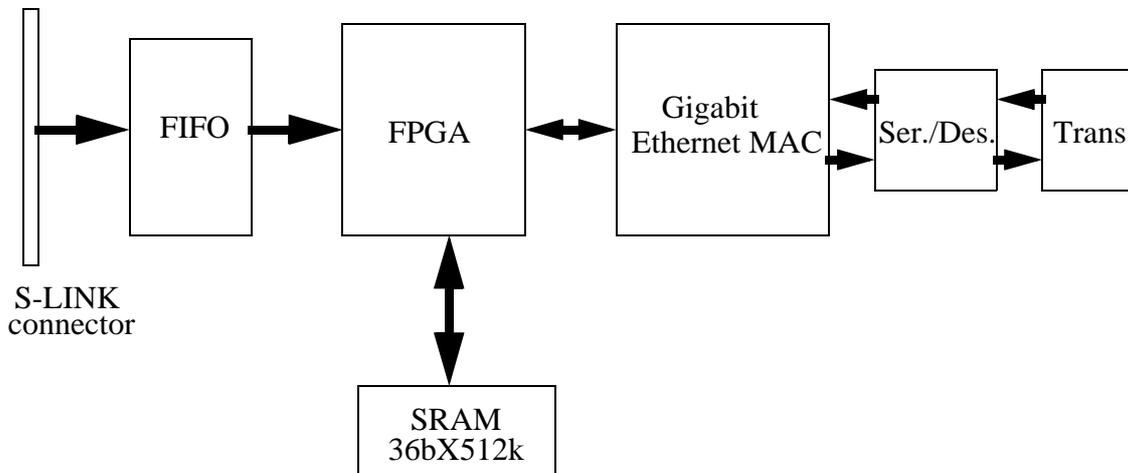


Figure 1: Block diagram of the Gigabit Ethernet LSC.

The Gigabit Ethernet MAC formats the data into Ethernet packets according to IEEE 802.3. The Physical Coding Sub-layer (PCS) of the MAC encodes the data, adds appropriate framing delimiters, and sends the 10 bit symbols to an external serialiser/deserialiser where the data is serialized. A fibre optic transceiver receives the serial bit stream and drives the optical fibre.

There is no JTAG port or other interface to an external processor to provide configuration information, but rather an on-board serial prom is used for configuration at power-up.

A return fibre drives the receiving node of the LSC to support auto-negotiation, transmission of the pause frame for flow control, and to allow the destination node to request event data from the synchronous SRAM.

The components for this prototype implementation are those selected in [2]. This selection maximises the use of expertise within the Trigger/DAQ community and hence reduces the design and implementation cycle.

3.2.1 FIFO

The FIFO is 4k words deep and is implemented by a pair of Cypress CY7C4245V-15ASC chips. The write enable is provided by the S-LINK enable (UWEN), and the strobe by the S-LINK clock (UCLK). Both of these signals originate from the S-LINK connector. The FIFO provides flow control back through the input data port when it reaches a watermark. The latter is defined on power-up via the FPGA using a value stored in the EPROM.

3.2.2 Gigabit Ethernet MAC

The Gigabit Ethernet MAC is implemented by the 8101/8104 Gigabit Ethernet Controller [3] which features independent 32-bit wide interfaces to internal transmit and receive FIFOs, sizes of 16 and 4 Kbyte respectively. There is an additional 16-bit wide bus interface to access on-chip registers and statistics counters.

The on chip-registers are initialized on power-up via the FPGA using values stored in the EPROM. Hence, the register values need to be known at the time the configuration EPROM is burned.

3.2.3 FPGA

The FPGA is implemented by a Altera 10K50V in a BGA package [4] and is configured via the on-board eprom. It provides the functionality to write event data from the FIFO to the on-board SRAM or to the MAC for transfer via Gigabit Ethernet or to the on-board SRAM, and to transfer data stored in the SRAM to the MAC for transfer via Gigabit Ethernet if requested by the target. When writing data to the MAC it is responsible for inserting the Ethernet header into the frames being sent. Depending on the size of the data block to be sent it may also fragment the data into several frames. At power-up it also initialises the 8101/8104 Gigabit Ethernet Controller.

3.2.4 Serialiser/Deserialiser

The serialiser/Deserialiser is implemented by a HDMP-1636 Physical Layer Device from Hewlett-Packard [5].

3.2.5 Fibre Optic Transceiver

The fibre optic transceiver is implemented by a HFBR-53D5 from Agilent [6].

3.2.6 Clock domains

There are two clock domains: S-LINK input clock rate of 40 MHz and data rate to the 8101/8104 of 125 MHz.

3.3 Control Words

Data arriving in the FIFO via the S-LINK connector are ROD event fragments framed by two S-LINK control words [7]. These event fragments may be either transferred directly to the MAC from the input FIFO for transfer via the Gigabit Ethernet or written to the on-board SRAM the data are retained in the S-LINK format with the concatenated control bits on the header and trailer words. These control words are removed by the FPGA. The framing of the ROD event fragments is provided by the Ethernet frame.

3.4 Flow control

In this prototype implementation there are potentially two sources which “drive” the S-LINK Link FIFO Full (LFF) line: the watermark reached within the FIFO; the watermark reached within the 8101/8104 transmit FIFO. In this prototype exercise, the LFF line will be driven low when the watermark is reached in the 8101/8104 transmit FIFO.

3.5 Link Reset

There are two links that will need to be reset: the gigabit Ethernet link and the input S-LINK connector. With respect to the gigabit Ethernet link, the reset will be accomplished by request-

ing auto-negotiation. As for the input S-LINK connector the reset procedure will be that specified in the S-LINK interface specification.

4 Testing

Testing of the prototype LSC will be done by placing the card on a Simple PCI to S-LINK PMC which, in turn, will be placed on a CES RIO2 running the LynxOS operating system. At the receiver end a COTS network interface, such as the Alteon, shall be used to receive the Ethernet Frames. Software to support this equipment already exists.

The initial testing will be performed at Argonne with the aim of ensuring correct functionality of the LSC and the ROL. The tests will concentrate on verifying: link reset, flow control and data transmission. The cards will then be sent to CERN for further testing and integration with the current implementation of the Read-Out Subsystem.

5 Time scales

The hardware implementation described in this document is straightforward as it leverages expertise within the Trigger and DAQ community, it does not involve any significant development program. The design of the card will be set down in ORCAD, the CAD work will be done in ALLEGRO and four or five prototype cards will be fabricated.

The configuration code for the FPGA will be developed and simulated in Altera MAX+II for the Altera 10K. In the first instance no use will be made of the on-board memory. The rest of the circuit involving the 8101 MAC, HDMP1636 serialiser/deserialiser, and HFBR5305 will be taken from the ARCHES II development [2].

Work on this prototype LSC is expected to begin in July, therefore: ORCAD design should be done during July; the CAD work and fabrication should be done by September; first prototypes should begin testing in October.

6 Component Cost

The component cost, i.e. small quantities, for the prototype card is given in the table below.

Component	Cost/\$
8101/8104 MAC	50
HDMP-1636	9
HFBR-53D5	130
Altera10k50vbc356-4	104
eprom	3
fifos	64
S-LINK connector	4
SRAM	75
PCB ^a	200
Total	639

a. Estimation

7 References

- [1]The S-LINK Interface Specification, <http://www.cern.ch/HSI/s-link/spec/spec> CERN, 1997.
- [2]Report on the Gigabit-Ethernet To HS-Link Adapter Port
- [3]LSI logic 8101/8104 Gigabit Ethernet Controller datasheet, <http://www.lsilogic.com/tech-lib/techdocs/networking/index.html>
- [4]FLEX 10K Embedded Programmable Logic Family Data, Altera. <http://www.altera.com/document/ds/dsf10k.pdf>
- [5]HDMP-1636 Transceiver, data sheet, Hewlett-Packard, <http://www.hp.com/HP-COMP/gigabit/geprod.html>
- [6]HFBR-53d5 850nm Duplex SC VCSEL Transceiver, data sheet, Agilent, <http://www.semiconductor.agilent.com/>
- [7]The event format in the ATLAS DAQ/EF prototype -1, Atlas internal note, ATL-DAQ-98-129 (1998).