KUNGL TEKNISKA HÖGSKOLAN

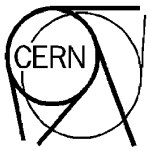

NFR

KTH, Royal Institute of Technology in Stockholm, Sweden
Section of Experimental Particle Physics

CERN, European Laboratory for Particle Physics in Geneva, Switzerland
EP-ATE, Experimental Physics, ATLAS Electronics Design Group, DAQ section

Funded by NFR, Naturveteskapliga forskningsrådet
TTA, Technotransfer AB

DEVELOPMENT OF A PROTOTYPE READ-OUT LINK FOR THE ATLAS EXPERIMENT

MASTER THESIS
BY
ERIK BRANDIN

TRITA-FYS 9106

ACKNOWLEDGMENTS

I would like to begin this report by thanking some of the people who contributed in different ways during the time it took to finish this project.

- Anders Hugnell at TTA and John Gustafsson at NFR for the financial support that made the stay at CERN possible.

- Robert McLaren, Erik Van der Bij, Zoltan Meggyesi, CERN, for technical support and letting me come into their group for my master thesis.

- My examiner, Per Carlson, and my supervisor, Mark Pearce, KTH, Stockholm, for the trust they put in me when vouching for me as a diploma worker at CERN.

- Jan Jönsson at Mitel Semiconductors AB for his interest in, and support of, this project.

- Kristina Gunne at the Swedish Office, CERN, for all the help when I came to CERN.

- Eleonor for standing by my side although being so far away.

ABSTRACT

LHC is an accelerator currently under development at CERN, built for p-p collisions at center of mass energy of 14 TeV. ATLAS is one of four detectors in LHC, and it will primarily search for the Higgs boson. The reactions looked for are extremely rare, which calls for high interaction rates. LHC will collide protons at 40.08MHz, with a multiplicity of up to 25 and this high interaction rate makes online data processing vital. Three trigger levels will in a few milliseconds reduce the event flow rate from 40.08MHz to 100Hz.

ATLAS will need high speed data links to transfer data to the final storage and in the same time serve the different trigger layers with enough information to come to a decision. Some of these link will be standardized and between trigger levels 1 and 2 there will be around 1700 of these standardized links running at Gb/s speed over distances of up to 200 meters. The current solution is the S-Link which is a CERN developed 32-bit@40MHz data link. The S-Link standard only defines a protocol and not the actual physical link which leaves it up to different developers to find the best and cheapest solution. ATLAS requires this link to have a total maximum component cost of 500 CHF, cabling not included.

This report describes the development of ODIN, an optical S-Link built on the 3.3V Hewlett-Packard G-Link chipset. The result was a single and a double ODIN, with one or two G-Link chipsets in the forward direction. The single ODIN have a maximum transfer speed of 128Mbytes/s and a component cost of 822 CHF. The double ODIN have a maximum transfer speed of 160Mbytes/s and a component cost of 1222 CHF.

TABLE OF CONTENTS

1 INTRODUCTION

1.1 ABOUT THIS REPORT

This report serves as final documentation for my Master Thesis for my education in Master of Science in Engineering Physics degree from the Royal institute of Technology in Stockholm, Sweden. The work is supposed to be at least 20 weeks of full time work, 20 credits, where approximately equal parts of studies, work and writing the report.

The work covered in this report is performed at the European Laboratory for Particle Physics, CERN, in Geneva, Switzerland from 12 April to 10 December 1999.

1.2 LAYOUT OF THIS REPORT

Chapter 1 LHC is the latest accelerator developed at CERN and ATLAS is one of four detectors that will be built. The first chapter gives a short description of LHC and ATLAS.

Chapter 2 Data transfer is one field of research and development in the ATLAS project. This chapter gives a background of the S-Link project and a short technical description of the S-Link specifications.

Chapter 3 The third chapter gives an introduction to programmable logic devices and the other main components chosen as hardware for the data link.

Chapter 4 S-Link is merely a definition of the protocol which is used to communicate with the outside world. A mapping from S-Link protocol to the protocol of the chosen serializer is needed, and this chapter describes how the S-Link specifications are met.

Chapter 5 The implementation of the FPGA is done using the language VHDL. This chapter describes some of the tools used.

Chapter 6 The last chapter concludes the project. It also states some areas where future developments can bring the S-Link project one step closer to the final goal.

1.3 LHC

The LHC, Large Hadron Collider is the next large project at CERN, the European Laboratory for particle physics. It consists of two rings and it will bring protons into collision at a center of mass (CM) energy of 14 TeV (7 TeV+7 TeV) and bring particle physics into previously undiscovered energy areas. The accelerator and its detectors are under current development and is planned to be running by 2005.

1.3.1 THE ACCELERATOR

The accelerator itself will be built in the same tunnel that now houses LEP, the Large Electron Positron Collider. LEP is now reaching the maximum CM energy of around 200 GeV (100 GeV+100 GeV), where the energy loss per turn due to synchrotron radiation approaches the energy that can be replaced by the superconducting accelerating cavities. Protons are about 1800 times heavier than the electron and will thus have less energy losses due to synchrotron radiation. The synchrotron radiation energy loss goes as $(E/m)^4$ and in LHC it will be in the order of a few ppm of the energy loss in LEP.

The LHC experiments looks for very rare reactions in proton-proton collisions. At high energies the cross section of a particle is associated with its de Broglie wavelength, and decreases like $1/E^2$. To achieve an interaction rate sufficiently high enough the accelerator must be loaded with a high density of particles compared to other accelerators. The 2 rings of the LHC will be injected with 2835 bunches of 10^{11} protons each and will at 4 collision points generate a bunch crossing rate of 40.08 MHz. At each bunch crossing up to 25 proton-proton collision are expected.

This large beam current will give rise to problems such as instability by beam-beam effect at the collision points, collective instabilities in each ring caused by a wake-filed and others. The needed magnetic field to keep the protons in the 27 km long circular tunnel will increase as the mass and energy increases. The LHC dipoles must produce stable magnetic field up to 8.36T, and this is one of the many challenging tasks that lie ahead of the LHC designers. There are many other aspects of the LHC that needs much consideration and research, for example the cryogenic system as all of the 27km of the LHC will be cooled down to 1.9K.

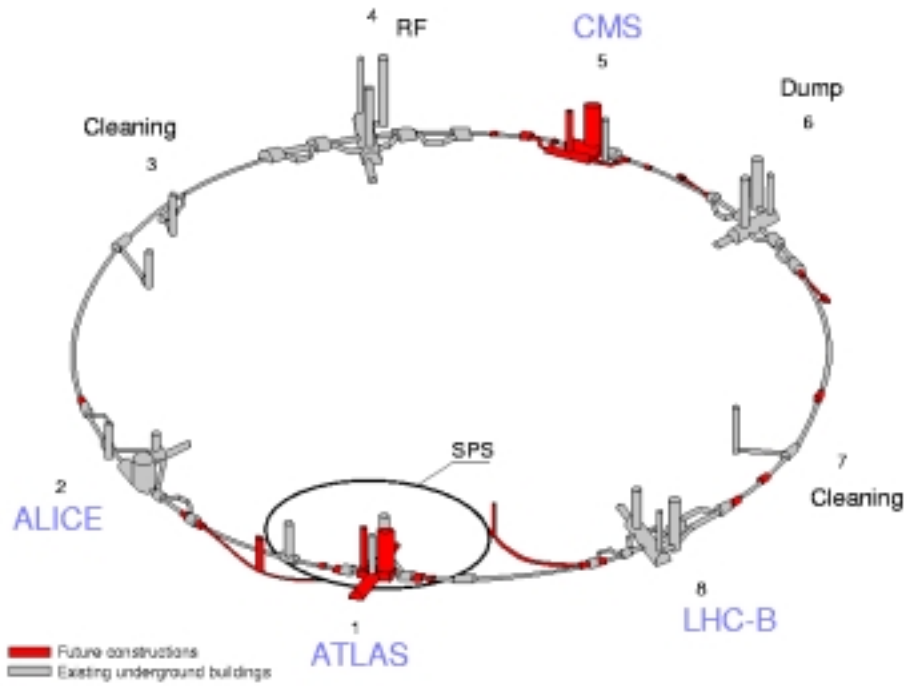

Figure 1-1 LHC and its detectors

1.3.2 THE DETECTORS

The LHC will supply proton-proton collisions to four different experiments. There are two general-purpose detectors, ATLAS and CMS, one detector designed for heavy ion collisions, ALICE, and one designed to look general b-physics, LHCb. See Figure 1-1 for an overview of the accelerator and detector sites.

1.4 ATLAS

ATLAS is the physically largest of the LHC detectors currently under development at CERN. The following subsections will describe the subsystems of the detector and the main physics issues that the detector is designed for.

1.4.1 GENERAL ATLAS DATA

The subsystems of ATLAS can be divided into three subgroups, inner detector, calorimeter and muon detector, as shown in Figure 1-2. The coordinates used to describe the systems are:

- the distance r from the beam axis,
- the angle ϕ in a plane perpendicular to the beam axis,
- the distance z from the center of the detector along the beam axis,
- the pseudorapidity η defined by η=-ln tan $\theta/2$, where θ is the angle with the beam axis, counted from the center of the detector. Large $|\eta|$ means direction close to the beam axis, and η=0 is the plane where z=0.

The high bunch-crossing rate of 40.08 MHz makes heavy demands on all parts of the detector design. The detector stretches out more that 10 meters all directions from the collision point, with its diameter of 22 meters and total length of 42 meters. None of the particles created at a bunch crossing will have time to leave the detector when the next bunch crossing occurs even if they are travelling with the speed of light. This gives rise to the pile-up effect that puts hard requirements in development of the sub detectors as each particle trajectory measured must be associated with a specific bunch crossing.

Another challenge is the level of radiation being produced within the detector, which calls for new solution of the radiation tolerant electronics that will continue operating during all of ATLAS lifetime.

At every bunch crossing around 25 proton-proton collision are anticipated. In order to distinguish different tracks from each other most of the subdetectors must be highly granular, which creates enormous amount of data, around $4 \cdot 10^{13}$ bytes/s, calculated on average event sizes of 1MBytes. For this reason both high speed data links and fast trigger system is needed.

1.4.2 THE PHYSICS

One of the main tasks of ATLAS is to search the Higgs boson, which is associated with the Higgs field believed to be responsible for giving elementary particles their mass. Depending on the mass of the Higgs boson, or if more than one Higgs particle is found would also give a hint how to proceed with the Standard Model into the area of supersymmetry. ATLAS will also be designed to look for other more exotic particles such as other supersymmetric particles and heavy W and Z-like objects.

1.4.3 INNER DETECTOR

The innermost part of the ATLAS detector is the inner detector, consisting of the Pixel detectors, Forward and Barrel Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). The design is in a 2-Tesla solenoidal field and its mission is to reconstruct track origins and to find secondary vertices. This last task is necessary to identify $b\bar{b}$-quark pair jets, which is one promising decay channel for the Higgs boson and also needed in the search for supersymmetric decays. As an example, to achieve the needed granularity the pixel detector which is the detector closest to the beam will detect charged particles with a resolution of about 10 μm in rϕ and 50 μm in z. To achieve this, the pixel detector itself uses 140 million channels for reading out collision data.

1.4.4 CALORIMETERS

Around the inner detector an electromagnetic and two hadronic calorimeters will be situated. The task for the EM calorimeter is to measure the energy of electrons and photons, with a rejection against jets of $5{\cdot}10^3$ and the main feature is a lead/liquid argon sampling calorimeter in an accordion shaped structure. Among other interesting physics channels looked for are $H \rightarrow \gamma\,\gamma$, $Z \rightarrow e^+\,e^-$, which are both mainly detected in the EM calorimeter which uses 200 000 channels.

The hadronic calorimeter measures the energy in jets and is also an essential contributor to the missing transverse energy measurement, which is vital for inference of undetectable particles such as the neutrino and other heavier neutral objects predicted by supersymmetric models. The Hadronic Tile Calorimeter is a based on iron absorber and plastic scintillating plate, the end-cap calorimeter is Cu/Liquid Argon and the forward calorimeter is a dense tungsten/liquid argon The total number of electronic channels in the hadronic calorimeter are around 10 000.

1.4.5 MUON DETECTOR

The task for the muon detector is to measure the momenta of muons and to give a muon trigger signal. The muon detector consist of three main components:

- A toroidal magnet systems using 8 superconducting magnets bends the muon trajectory. An exact measurement of the tracks together with the magnetic field gives the muon momenta.

- The muon tracker measures muon tracks in the z direction with a high precision. Up to $|\eta|=2$ the tracks are measured by drift tube chambers, resulting in a resolution of 80μm. The maximum drift time is 500 ns which gives a serious pile up effect. At $|\eta|>2$ cathode drift chambers are used, and the resolution is 60 μm.

- The muon triggers system is necessary to identify a muon signal with the correct bunch crossing. This system is composed of a resistive plate chamber and a thin gap chamber. The space resolution is of the order of cm, but the time resolution is of the order of ns, thus every muon signal can be associated with the correct bunch crossing which is important for a correct trigger decision.

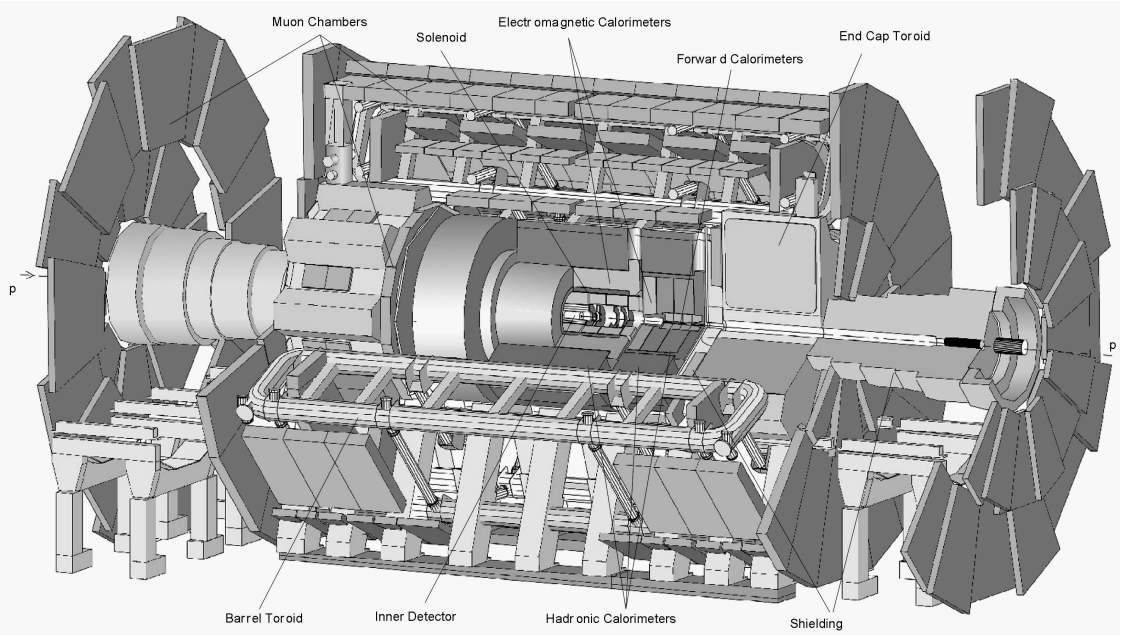

Figure 1-2 Cut-away schematic of ATLAS

1.4.6 DAQ AND TRIGGER

The data acquisition system (DAQ) is extremely important in ATLAS as the interaction rate is extremely high and the total data gathered each bunch crossing is immense. To reduce the number of events for permanent storage three levels of online event selection is implemented, see Figure 1-3. These three levels yields a $4 \cdot 10^5$-rejection rate from 40MHz down to 100Hz.

The readout system, also seen in Figure 1-3, is basically based on two links where event data is transferred. Data is sent from an on-detector site through the Front End (FE) links to the Readout Drivers (ROD) at an off-detector location around 80 meters from the detector. From the RODs data is transferred to the Readout Buffers (ROB) at a central location on the surface. These links must handle transfer distances of around 200 meters.

The level 1 trigger is implemented as hardwired processors in the detector area which makes its decision based on reduced granularity information from the muon detector and the calorimeters. For a trigger decision combinations of events are looked for, such as high-p_T muons, electrons/photons, taus/hadrons or jets. For high flexibility these selection rules can be changed depending on the current luminosity of the accelerator. The selection rules are called trigger menus [2]. These menus should cover the expected reactions as well as more exotic ones. The level 1 trigger must reach a decision in 2.5 μs, and during that time the detector data is stored in pipeline memories. Events selected by the level 1 trigger all event data, around 1 Mbytes is sent through the FE-links. A trigger acceptance rate of 100 kHz gives a total bandwidth of the read-out link of 10^5 Mbytes/s.

The level 2 trigger look at full granularity data from all detectors, including the inner detector. In order to meet the decision time requirement of 10 ms, only data from certain Regions of Interest (RoI) chosen by the level 1 trigger is used in the analysis. Events selected by level 2 trigger is transferred through the ROD-ROB links and a full event building is performed.

The event builder (EB), originally named level 3 trigger, seen in Figure 1-3 collects all event fragments from the ROBs and makes the final trigger decision based on the full event. This last step in the trigger chain is called the Event Filter (EF). Already in trigger level 2 all the data is available, and exactly where the boundary will be between LVL2 and EF triggers will be flexible in the final version for optimization during running.

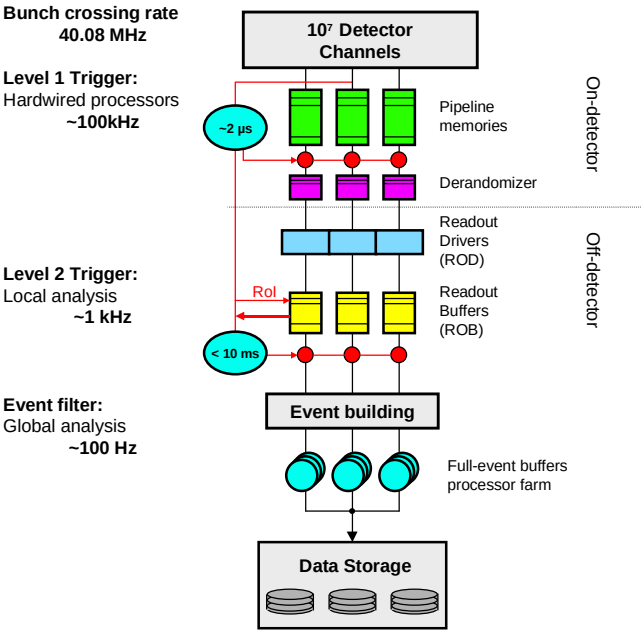

Figure 1-3 Block diagram of the Trigger/DAQ system

2 GOALS OF PROJECT

In this chapter the goal of this project will be defined and the project will be put in a wider context.

2.1 ATLAS READ OUT LINK

2.1.1 BACKGROUND

In a collaboration of the size of ATLAS there must be both a sense of flexibility for the different communities designing the subdetectors and DAQ systems and some larger standardization for easier maintenance and more stable running. In the data chain the boundary for standardization is set at the ROD-ROB data transfer [7], see figure 1-3. To the point of the RODs every detector group will be responsible for the data links and are free to develop their own solution. These links are the front end links, which are purely synchronous to the ATLAS clock at 40.08 MHz, and they transfer the raw digitized data from the subdetectors. The transmitting part of these links are placed inside the detector and they must be radiation tolerant. The protocol of the ROD output will however be standardized for all subdetectors. At the time of writing this protocol is not definitely set.

A standardization of the readout link has many advantages. It is foreseen that ATLAS will need ~1700 read out links and a standardization means cuts in non-recurring engineering costs as well as lower production cost due to a greater volume. One data format will also make the data gathering at LVL2 and EF triggers a simpler task.

The standardized readout link is expected to be running at around 1Gbit/s data link which at 40 MHz translates to 27 bit wide input, or more convenient 32 bits at 40MHz. The links should have the possibility of flow control, i.e. the receiver must be able to stop data transmission without loosing data. The transmission distance is 200m which at Gb/s speed calls for an optical solution. There are also power consumption and price requirements. Maximum power consumption is set to 7.5W per card, i.e. a total of 15W for both the transmitter and receiver card. Price is, apart from performance, the one most important factor for creating an ATLAS standard for readout links. The price goal before implementation is set to 2*250CHF in component cost. In this cost PCB, cabling and production cost is not included.

2.1.2 SHORT S-LINK SPECIFICATIONS

The current prototype for the readout link is the S-Link [4], a CERN developed data link standard which was proposed as an ATLAS read out link at the Fourth Workshop on Electronics for LHC Experiments [5].

The S-Link is a synchronous 32-bit FIFO-like data link. The S-Link specification only defines the interface, not the physical layer that the link is built upon. By doing this, the final S-Link, when ATLAS is to be assembled, can utilize future developments in different areas such as power consumption, price, size, speed etc. Furthermore, ATLAS can not rely on the

future accessibility of a certain technology, which would be the case if a chip was chosen at this early stage. The specification also includes features like:

- Error detection
- Control words intended for use in event headers
- Internal test pattern generator and checker
- Flow control (in duplex version)
- Reset function

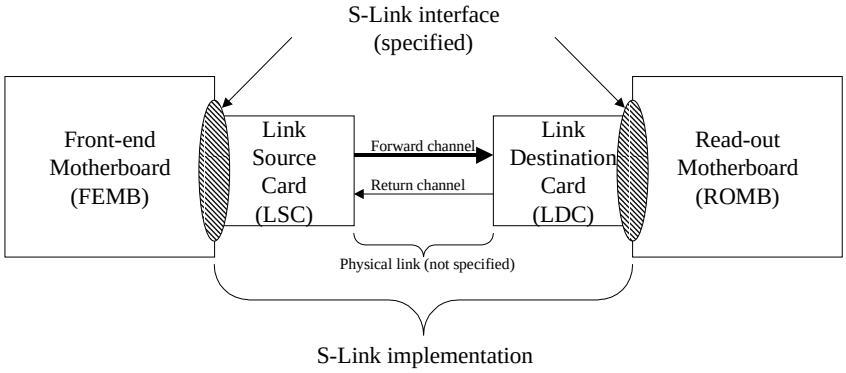

Figure 2-1 The S-Link concept

Another advantage of using S-Link as a prototype readout link is the test tools designed. These include:

- SLIDAS, an S-LINK Infinite Data Source. This card can be used to emulate the S-Link LDC as it can generate a number of preset data patterns at settable size and frequency. It can also be used as a data generator when developing an S-Link implementation.

- SLIDAD, an S-LINK Infinite Data Drain. This is the opposite of the SLIDAS, since it can be used to emulate the S-Link LSC. It can also be used to be connected the LDC in an S-Link test bench.

- SLITEST, S-LINK Tester. This board is needed when using SLIDAS and SLIDAD as an S-Link test bench.

- SLIBOX, S-LINK Extender and Break-Out Box. This box connects to an S-Link test bench to access all pins of the S-Link connector.

- S2P2, S-LINK to P2 adapter. The S2P2 adapter is a VME board that can carry an LSC or an LDC, which connects to a VME P2 connector. With this setup an S-Link implementation can be tested under more real circumstances.

These tools are available and useful for both detector groups since SLIDAS and SLIDAD emulates the S-Link interface for the ROMB and the FEMB, and S-Link developers as the SLIDAS and SLIDAD mounted on a SLITEST emulates both ends of the interface.

S-Link products already have users at other CERN experiments (NA48 and Compass) as well as outside CERN (ASDEX, Megacam and AT&T Laboratories) [6].

2.2 GOAL OF PROJECT

The goal of this project is to design an optical S-LINK for a lower cost, running at a greater speed and lower power consumption than the previously existing S-Links. This will be achieved by taking advantage of the recent developments in serializers-deserializer chipsets, optical transceivers and programmable logic devices (PLDs or FPGAs).

2.3 PREVIOUS LINKS — OVERVIEW

Numerous links have previously been designed. Currently available links are the E-Slink and the FCS-Link2 [6].

The E-Slink is a parallel electrical S-Link, which is synchronous data transfer mode, i.e. transmission clock equals double user supplied clock, gives a maximum data transfer rate of 160Mbytes/s, 32bits@40MHz, over 10m cables. In asynchronous mode the maximum transfer rate is 80Mbytes/s over 25m cables. ATLAS requires read out links over a distance of 200m and the space for cables also makes this solution inappropriate. The specified power consumption is 3.5W (typical) up to 5W (maximum) per card.

The FCS-Link2 is a serial optical or electrical link based on the Fiber Channel (FC) serializer/deserializer chipset. Maximum data transmission rate is 103Mbytes/s over up to 500m of optical cables. Apart from the low transmission frequency there are a few other drawbacks, such as price and power consumption. The FC protocol also makes the link inappropriate for sending data blocks of small size, since the FC protocol needs to frame all data, or at lower frequencies than 27 MHz, which is the internal FC transmission frequency. The power consumption is 5W, maximum 6W per card.

3 DESCRIPTION OF HARDWARE

3.1 PHYSICAL TRANMISSION MEDIA

Today electronic designs often go towards lowering the power supply voltage, hence lowering the power dissipation. The 5V power supply standard of yesterday is now changed to a 3.3V power supply standard. Some chips are even 2.5V and in the future even lower power supplies are possible.

The Hewlett Packard low-power G-Link chipset (HDMP- 1032/1034)[3] is a 3.3V 16-bit serializer/deserializer which is suitable for point to point data links, such as the S-Link[4]. This chip-set is built on the previous 5V G-Link (HDMP-1022/1024) with some functionality removed in order to reduce price and power consumption. The new chip-set also has one new feature called 'Enhanced Simplex Mode' which 'provides an extra level of protection to guard against improper word alignment caused by static valid code-filed bits embedded within the data-field' [3].

The G-link (including the previous versions) uses CIMT encoding (Conditional Inversion Master Transition), which ensures DC balance of the serial line, equal numbers of 1's and 0's sent over the serial line. This is important as the receiver is AC-coupled and a non DC-balanced serial transmission would lead to a charge build-up in the receiver, and the sensitivity would go down. A 4-bit code field appends the 16-bit data, and the code field always include a transition between the middle bits, called the master transition, shown in Figure 3-1. This transition is used by the receiver chip for phase locking the incoming data stream and for word alignment when phase lock is achieved.

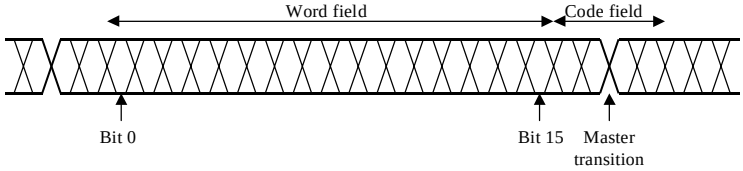

Figure 3-1 G-Link coding scheme

The G-Link also features one flag bit and one control bit. When setting the control bit only 14 bits are transmitted, which limits the use of G-Link control words.

When using a 16-bit transmission chip-set for a 32-bit protocol, a decision must be made if speed or cost is the most important. That is, should the 32-bit data be split to two 16-bit channels, or should the data be multiplexed to 16 bits thus using 2 clock cycles for transmission. Obviously the first solution yields a faster data link, and the latter a cheaper link. The maximum frequency that is guaranteed by the G-Link data sheet is 70MHz. This gives a maximum frequency of 35MHz at 32-bit datawidth. As price and speed are the two major concerns of the final solution two version of the link will be designed, one with 2 G-

Links[1] reaching the maximum data rate of 160Mbytes/s, and one with 1 G-Link[1] reaching only 140Mbytes/s[2].

3.2 PROTOCOL CHIP

3.2.1 PROGRAMMABLE DEVICES

In order to map the S-Link protocol onto the G-Link protocol one needs a protocol chip, primarily for framing and error detection. For this purpose a programmable logic device (PLD) is suitable. Programmable logic circuits have many advantages over traditional ASIC designs. Some of these are:

- Programmable
- Shorter development cycle
- Cost effective, especially in low volume
- Good testability
- Large market, many choices of vendors
- Last-minute design changes
- New versions only needs update of PROM

3.2.2 SHORT PLD HISTORY

It is difficult, almost impossible to give a correct overview of the history of programmable logic devices. The development goes extremely fast. Different vendors have chosen different solutions but using the same name. Others choose the same basic architecture but using different names.

When the first PLDs was introduced, in mid 1970s, they only included a few gates and larger designs often included many chips [10],[11]. The Programmable Array Logic (PAL) was a commonly used PLD consisting of a programmable AND-plane followed by a fixed OR-plane. Since then the development has been fast. A decade later, in 1985, Xilinx introduced the first programmable gate-array, and the name Field Programmable Gate Array (FPGA) is today commonly (not always correctly) used for most types of programmable devices. A Gate Array in a broader sense is a type of ASIC in which the transistors, gates, and other active circuit elements are fixed on a wafer called a "master slice".

Some vendors, such as Altera, uses Look-Up Tables with 4 inputs instead of fixed gates, and the programming in these devices are related to the configuration of these LUTs and the routing of the signals in between. These devices are by Altera called Complex Programmable Logic Device (CPLD) although it has many similarities with the gate array architecture.

---

[1] The number of G-Links specified is in the forward channel. As the link will be a duplex link one G-Link chipset is used for the return channel.

[2] This is not good enough for a final-solution S-Link, but for most users this transfer speed is high enough as it does exceed the 1Gbit/s transfer speed required for an ATLAS read-out link.

Today's programmable chips include typically up to a few 100K typical gates[3] and up to 1000 I/O pins which makes them a better competitor to traditional ASICs.

3.2.3 THE STATE MACHINE

One main component of many digital designs is the state machine. In a sense this is a way to make the parallelism of a digital circuit behave in a sequential way. In section 4.5 the reset cycle of the link will be described by discussing the state machines used to explain the behavior, thus a brief introduction might be needed.

A state machine is a set of states of which one is occupied at each given time and a set of rules when to move from one state to another. A state is only a set of registers that can only be 0 or 1. A common example is a non-resetting '1011' pattern detector shown as a state diagram in Figure 3-2. Note that the actual register values that is used to implement this state machine are not shown in the diagram. Each state is given a symbolic name shown in the top half of each state and this is the current detected sequence. The input signal to the state machine is shown in the transition between the states. Depending on the input being either 1 or 0, the next state is determined. No matter which state is currently occupied, an input sequence '1011' will always end up in the right-most state. The bottom of each state is the output from the state machine and the output '1' signals that '1011' is detected.

The output of the state machine below only depends on the current state, and this is called a Moore state machine. The other kind of state machine is the Mealy state machine where the output is associated with a state transition instead of a certain state.

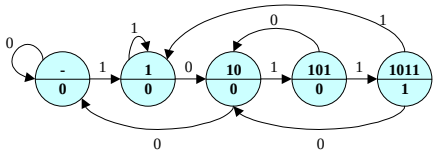

Figure 3-2 Moore State machine of a sequence detector

3.2.4 CHOOSING DEVICE

When choosing an FPGA vendor there are a few decisions to be made. Some are one time programmable while other are reprogrammable. Size versus speed is another important decision to make as well as the need of embedded RAM or built-in LVDS serializer. The decision often comes down to previous experiences with different vendors, and this narrowed the options to Actel, Altera and Xilinx. Altera was chosen mainly for the ability to reprogram the FPGAs which gives more room for debugging the final system, and because of previous experiences with these devices. Size and speed grade of the chip is the next decision to make and as both the price and power consumption goes up for larger and faster

[3] Since most PLDs today don't use gates as a main part of the architecture the size of a PLD is measured in *typical gates*. The exact corresponding number of maximum usable gates depends on the design. Most vendors also specifies maximum and minimum number of typical gates for each chip.

devices and the smallest possible device is preferred. At an early stage of the design the most logic-demanding block of the design, the input FIFO memory of the transmitter card, showed that an Altera FLEX 10K30 was necessary. 30% of the logic elements were used by a 8x33bit asynchronous FIFO, and Altera does not recommend usage of above 70% for internal routing reasons. It was decided to use a 10K30PQFP240, where 240 stands for the number of pins, and of those are 189 I/O pins user accessible, which are connected to the internal logic through the I/O Elements (IOE).

3.2.5 ALTERA FLEX ARCHITECTURE

The traditional FPGA architecture includes segmented signal lines with interconnecting switching matrices, i.e. short signal lines where signals have to be routed through sometimes more than one switch. This architecture makes routing easy, but delays are added in each interconnect which adds to the total signal delay, hence lowering maximum frequency. Altera FLEX 10K architecture instead uses what is called a Fast Track Interconnect routing which are continuos horizontal (Row Interconnect) and vertical (Channel Interconnect) signal lines, see Figure 3-3.

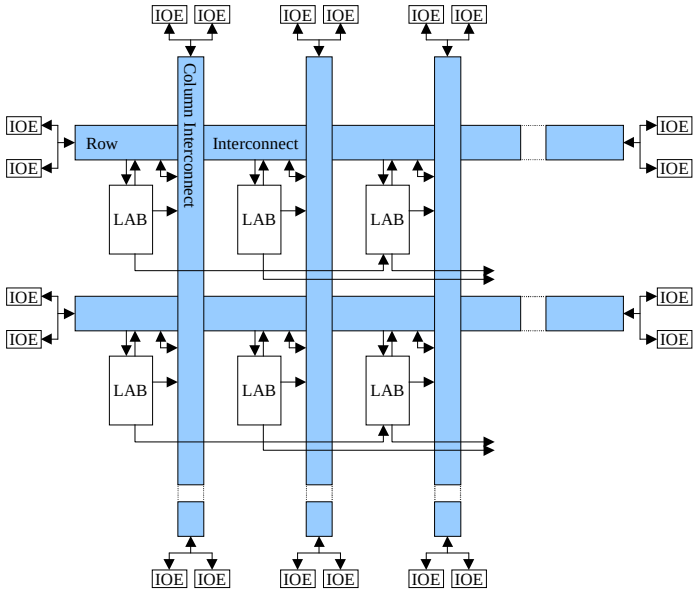

Figure 3-3 Altera Flex 10K Fast Track Interconnect

Each Logic Array Block (LAB) consists of 8 logic elements (LE), which are the smallest unit of logic, see Figure 3-4. The LE consists of a look-up table, a programmable flip flop and some dedicated signal lines The main part of the programming of the device concerns the configuration of, and routing the output from, these LEs. The output from a LAB can be routed to both the row and column interconnects, but the input to the LAB is only connected to the row channel. Figure 3-3 also shows signals routed directly in-between

LABs These are the cascade and carry chains, which by its direct routing is suitable for larger adders (carry chain) and functions with a wide fan-in[4] (cascade chain). Furthermore, every LE can exchange access to the column channel with an adjacent LAB. This signal line is not shown in the figure.

There are two ways to connect a signal from an LAB to an interconnect, either directly or through a register (not shown in the figures). This is an important feature since most digital designs are synchronous, i.e. signals are stored in registers that are clocked by one master clock. The by-pass of a register is mainly used in large combinatorial networks and the faster carry or cascade chains for some reason cannot be used.

Apart from the 22 (26 for larger devices) input signals to the LAB local interconnect there is 6 global input signals that are connected to every logic element of the device. Two of these are dedicated clock inputs, and the other four are dedicated inputs. The latter can also be accessed internally, and can be used as internally generated clocks or reset signals which often can be useful.

As each row interconnect does not contain enough channels to drive every look up table on a specific row and 100% usage of the logic is usually not possible. Fitting the logic in a device can be visualized as a giant jig saw puzzle, where all signals must be correctly interconnected, thus the 70% usage of the LEs is recommended by Altera.

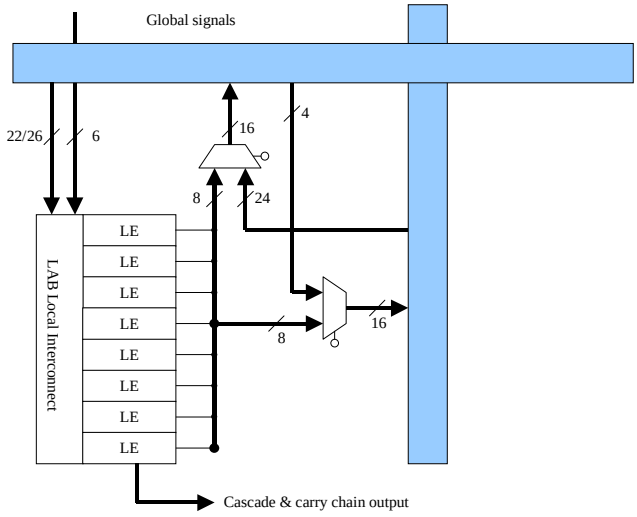

Figure 3-4 Close-up of LAB

[4] Fan-in is the input signals to a function.

3.3 OPTICAL TRANSMITTER AND RECEIVER

As stated in chapter 2, price reduction is a major task of the S-Link project. The optical transceivers have always been a large part of the total cost. In an attempt to reduce the total price a new optical transceiver is tested. The VF-45 fiber-optic connector, also called Volition, does not use a ferrule and this considerably lowers the price.

One aspect of this solution is that the bandwidth of the return channel of the duplex S-Link is greater than what is required, and for that reason a tailor made solution is better. This could include the use of VCSEL arrays as optical transmitter, such as the Mitel 4D769. This solution would also require a laser driver, which doesn't make this solution cost effective for one card. However if several S-Link were implemented on one single board, this solution would be one strong candidate as all channels in an array would be used. One must also include the laser driver and the PIN-diode in the cost calculations. This solution will be further investigated in spring 2000 [6] as a board with Mitel 4D769 VCSEL array with two Microcosm MC2062 VCSEL driver will be built as an daugther board.

4 IMPLEMENTING A PROTOCOL

After deciding which physical media will be used for the serial transmission one has to invent a protocol that the source and destination card will use to communicate. It is not likely that the protocol serializer/deserializer chipset, in this case the low power G-Link will follow the protocol defined to the user, the S-Link protocol. There is also a need for better error detection and flow control than the G-LINK offers.

The project was given the name ODIN, for **O**ptical **D**ual G-L**in**k.

4.1 DATA SPLITTING

As stated in section 3.1 there are two versions of the link. For the single G-Link version the data splitting scheme is obvious. The 32-bit S-Link word is split and sent as two consecutive G-Link words. If the most significant bits are sent first, the least significant bits must be pipelined in a register to send the next clock cycle. The transmission frequency is limited by the G-Link and the fiber optics and the Altera.

For the double G-Link version two approaches exists. The same scheme as the single G-Link version can be used, alternating between the G-Links. This way an S-Link word will be sent over one specific channel. The other approach would be splitting an S-Link word into two G-Link words and be sent in parallel in one clock cycle. One channel would always send the most significant bits while the other channel would send the least significant bits. The latter solution is tempting, since it does not require any pipelining of data waiting to be transmitted. However, the first solution is more similar to the one used by the single G-Link version, and would yield the least number of differences between the versions. The first solution is therefore chosen.

Control words are sent the same way, except for the few differences discussed in the following chapters.

4.2 ERROR DETECTION

Almost all transmissions of digital data require some kind of error *detecting* or error *correcting* capabilities. Which one is chosen depends largely upon the application. The data that is sent through the ROD-ROB links, where S-Link is the prototype, will be rejected with a factor of 100 by the LVL2 trigger, and if retransmission would be decided by the link itself the performance of the readout links would degraded when errors are encountered. But there is still need to check the data for errors. Because of this, the S-LINK specifies an error detection [4], which is up to the designer of the actual link to choose.

4.2.1 DATA BLOCK

The S-LINK specification leaves it up to the designer of the link to choose either word by word or block error detection. Neither of these are available by the G-LINK. Even though the G-LINK has an error bit this only includes errors in the 4 framing bits and not the 16 bits data field.

Word by word error detection requires additional data to be sent on each word, for example as parity bits. The G-LINK only offers one flag bit that would be convenient for this purpose, although 1 parity bit for 16 data bits might be considered too weak as a main error detection. As we shall later see, this flag bit will be used for other purposes, consequently another error detection is required.

There are several options for achieving block error detection, such as Cyclic Redundancy Check (CRC), Hamming Code and Reed-Solomon Coding [8]. The CRC error detection offers good balance between easy implementation and good error detection capability. The CRC error detection uses a given polynomial of the N^{th} power, N+1 bits, with which the data stream is modulo-2 divided by. A modulo polynomial has no carry operation between places, i.e. each place is computed separately, which makes the division easier. The quotient is disregarded and the remainder of N bits, also called checksum, is appended to the data stream. When the data with the appended reminder is divided by the same polynomial, the reminder should be zero if no errors occurred.

A short example, with polynomial x^3+x^2+1 (which translates into the binary number '1101' when 2 is used as the base x) and data '110010', is shown in Figure 4-1. A few things should be noticed. For a polynomial of the N^{th} power, N zeroes must be added to the data to calculate the remainder. Also, "a number goes into another number" when the first digit of the divisor is greater or equal than the first digit of the divided. In this case the transferred data would be '110010100', which is evenly mod 2 dividable with '1101'.

```
                    1 0 0 1 0 0
  1 1 0 1 | 1 1 0 0 1 0 0 0 0
          - 1 1 0 1
            1 1 0 0
          - 1 1 0 1
                    1 0 0
```

Figure 4-1 Modulo 2 division example

CRC is easiest implemented in hardware by a shift register with xor feedback doing the modulo 2 subtraction (XOR is also used for addition). shows the same example as above implemented in shift registers and XOR feedback. After the message '110010' is shifted through, the shift register contains '100', if the shift register is initialized with all 0's. If this remainder also is shifted through the register contains '000', i.e. no errors in transmission. Note that this scheme is slightly modified and does not require additional 0's to be appended. This improves the speed performance of the error detection as the reminder is calculated when the last data bit is sent (or received on the other side of the link).

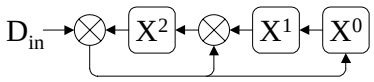

Figure 4-2 CRC shift register

The CRC algorithm works just as well regardless of the initialized value, although the remainder changes. The only criteria is that the receiving side initializes the register to the same value. Initializing to 0's will make the CRC algorithm blind to 0's at the beginning of the data. This is why many CRC implementations initializes the register to non-zero values such as all 1's.

A commonly used CRC code is the CRC-CCITT, $x^{16}+x^{12}+x^5+1$. It was first used by IBM for their 8-inch floppy disks, and is now used in almost all floppy disk controller devices, although invisible for the user. This CRC is especially convenient since it leaves a 16-bit reminder that is easily appended on a stream sent by a G-Link. The shift register representation of the CRC-CCITT algorithm is shown in Figure 4-3.

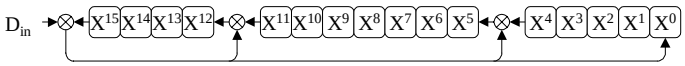

Figure 4-3 CRC-CCITT shift register

The shift register approach is suitable for serial data as the reminder is computed after each data bit. For an M-bit parallel data M shifts must be calculated in terms of the M bits data and the current N bits reminder. By assigning the register and data symbolic values one can easily calculate the value of the register after an M-bit parallel word. If the initial register values are named R[15..0] and the data D[15..0], the register value R_M[15..0] after the M-bit word can be calculated using the equations in Table 4-1. The M-bit word is shifted with bit 0 first to yield these equations. As a result the remainder must be reversed before appended to data stream.

4.2.2 CONTROL WORDS

The S-Link protocol specifies a control word error bit, apart from data error bit. This calls for an error detection on control words distinguished from the CRC error detection on data. CRC is not suitable for very small data volume. By appending a 16 bit CRC checksum to 28 bits data would obviously make little sense. Instead parity is used. By sending a 28-bit S-Link control word as two 16-bit G-Link words, 4 bits are free for parity. Even parity is chosen, merely to make a control word consisting of all zeroes to have all zeroes even with parity bits appended.

4.3 FRAMING

A G-Link does not require any form of block framing when sending data. The CRC however needs some kind of framing to distinguish one data block from another because the CRC shift register must be reset. This framing will be transparent for the user.

R_M 0 = R0⊕R4⊕R8⊕R11⊕R12⊕D3⊕D4⊕D7⊕D11⊕D15
R_M 1 = R1⊕R5⊕R9⊕R12⊕R13⊕D2⊕D3⊕D6⊕D10⊕D14
R_M 2 = R2⊕R6⊕R10⊕R13⊕R14⊕D1⊕D2⊕D5⊕D9⊕D13
R_M 3 = R3⊕R7⊕R11⊕R14⊕R15⊕D0⊕D1⊕D4⊕D8⊕D12
R_M 4 = R4⊕R8⊕R12⊕R15⊕D0⊕D3⊕D7⊕D11
R_M 5 = R0⊕R4⊕R5⊕R8⊕R9⊕R11⊕R12⊕R13⊕D2⊕D3⊕D4⊕D6⊕D7⊕D10⊕D11⊕D15
R_M 6 = R1⊕R5⊕R6⊕R9⊕R10⊕R12⊕R13⊕R14⊕D1⊕D2⊕D3⊕D5⊕D6⊕D9⊕D10⊕D14
R_M 7 = R2⊕R6⊕R7⊕R10⊕R11⊕R13⊕R14⊕R15⊕D0⊕D1⊕D2⊕D4⊕D5⊕D8⊕D9⊕D13
R_M 8 = R3⊕R7⊕R8⊕R11⊕R12⊕R14⊕R15⊕D0⊕D1⊕D3⊕D4⊕D7⊕D8⊕D12
R_M 9 = R4⊕R8⊕R9⊕R12⊕R13⊕R15⊕D0⊕D2⊕D3⊕D6⊕D7⊕D11
R_M 10 = R5⊕R9⊕R10⊕R13⊕R14⊕D1⊕D2⊕D5⊕D6⊕D10
R_M 11 = R6⊕R10⊕R11⊕R14⊕R15⊕D0⊕D1⊕D4⊕D5⊕D9
R_M 12 = R0⊕R4⊕ R7⊕R8⊕R15⊕D0⊕D7⊕D8⊕D11⊕D15
R_M 13 = R1⊕R5⊕R8⊕R9⊕D6⊕D7⊕D10⊕D14
R_M 14 = R2⊕R6⊕R9⊕R10⊕D5⊕D6⊕D9⊕D13
R_M 15 = R3⊕R7⊕R10⊕R11⊕D4⊕D5⊕D8⊕D12

⊕ denotes the XOR operation

Table 4-1 Equations for 16-bit look-ahead for CRC-CCITT

4.3.1 DATA WORDS AND CRC

The receiving side of the link must be able to identify data words from control words. Setting the flag bit high for data word fragments solves this. There is no transmitted information about most or least significant bits, since most significant bits always precede the least significant bits. Loss of one G-Link word would in this case mean corrupt data, but both the CRC and the G-Link itself would identify this as a transmission error and realignment would occur after the next CRC or idle word.

The flag bit is also set high for the CRC checksum. This way the CRC logic only has to look at the flag bit for CRC calculations.

4.3.2 CONTROL WORDS

Control words are sent in basically the same way as data words. The only difference is that flag bit is low during sending of control words. This achieves two goals, it tells the CRC logic that this data should not be included in the CRC stream and it also tells the decoding logic of the receiver that it is an S-Link control word

4.3.3 INTERNAL COMMANDS

In a data link there is always a need for some internal command, such as link initialization, reset signals and link state information, such as link down. These signals are not visible for the link user. For this purpose the G-Link control word is suitable. A full reference of the internal commands can be found in the ODIN hardware specification, appendix B.

4.3.4 TRANSMISSION EXAMPLES

Table 4-2 shows a transmission example where a data block of 2 data words (DW1 and DW2) are sent, framed by control words to mark begin and end of fragment (CW1 and CW2). Note that this example is from the single channel ODIN. For a double channel ODIN only the internal command and the CRC checksum is sent over both channels, see appendix C.

G-Link data	Flag bit	Data bit[5]	Control bit[5]
CW1[31..16]	0	1	0
CW1[15..4]+parity	0	1	0
DW1[31..16]	1	1	0
DW1[15..0]	1	1	0
DW2[31..16]	1	1	0
DW2[15..0]	1	1	0
Internal command	-	0	1
CRC checksum	1	1	0
CW2[31..16]	0	1	0
CW2[15..4]+parity	0	1	0

Table 4-2 Transmission example

4.4 RETURN CHANNEL

The Duplex S-Link features 4 return lines and 1 flow control to stop the source card from sending data. In addition there is also a need for internal commands in the return channel. These features add up to 7 channels, and it is obviously not essential to use a 16-bit high speed serializer/deserializer for this purpose. However by using the G-Link for the return channel also, the same PCB can be used for both the LSC and the LDC, which lowers the total production cost.

4.4.1 ERROR DETECTION

There is no need for an advanced error detection scheme on the return channel, as this would only add complexity in the design. An even parity bit is added to each bit, i.e. every bit is sent twice. If any error is encountered on the return channel, the entire word is discarded, even for the parity bits that are correct.

---

[5] This is the G-Link data and control word bit, not to be confused with the S-Link data and control words.

4.5 RESET CYCLE

The reset function is a vital part of a data link. If an error is encountered it is preferred that the system around the link knows about the problem to be able to make a decision how to handle it. For example a link might go down due to a transient error such radiation, voltage drop, disconnected optical fiber etc. The reset function must be easy to use and also reliable and preferably fast.

4.5.1 CURRENT PROBLEMS

When an S-Link encounters an error it is specified that this state should be latched until the user initiates a reset cycle to prevent links to clear an error by itself. The S-Link specification also defines a card reset, i.e. both cards must be reset to clear an error latch. In addition, if a reset cycle is activated when a card is up, this should cause the other card to go down. The link should also come up when powered up and cables are connected. All these features combined are somewhat contradicting, and the implementation of the reset function is not elementary.

4.5.2 NEW SOLUTION

The aim is to make the reset a secure operation that will work the same way regardless of which state the link is in. The first step was to redefine the card reset to a link reset instead of a card reset. The reset should reset both cards, regardless of which sides the reset pin was active or which state the link was in before the reset. It should be backward compatible, so a system designed to work with any S-Link should work with the new link reset.

The solution to the problem is to construct simple, yet functional, state machines for both sides of the link, see figures 4-1 and 4-2. On power up both state machines enters Power-state. From there, LSC start to send reset commands to LDC 'semi-continuously', one command followed by 7 idle words. This helps the G-Link receiver to acquire lock faster and it eliminates the risk of wrong word alignment due to 'static valid code field embedded in data field' [3]. Remote Link Up (rlup) is an internally generated signal that is set low only when a Remote Link down is received. There is a 21-bit counter on the G-Link transmitter locked and receiver ready signals that generates the hp_up signal. This is to filter out up to 1 ms 'glitches' on the receiver that might occur if the fiber is connected when link is powered up. The counter in the LSC on RESET → UP is only to assure a link down of at least 4 clock cycles when LSC is reset, according to S-Link specification.

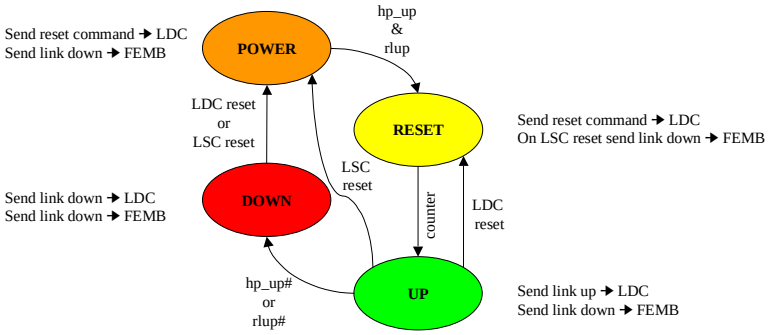

Figure 4-4 LSC state machine

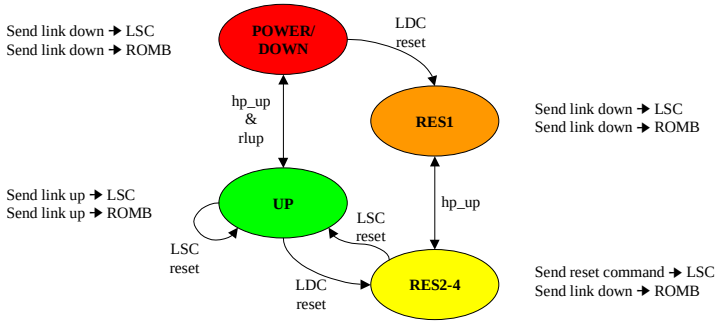

Figure 4-5 LDC state machine

5 IMPLEMENTATION

5.1 PCB

The schematic layout of the PCB was done by Erik Van der Bij, and the PCB design was done at CERN. The schematics can be found at URL: http://www.cern.ch/HSI/S-LINK/odin

5.2 HARDWARE / VHDL

The Altera protocol chip was programmed mainly using the Altera MaxPlus software package and VHDL [9]. VHDL (VHSIC Hardware Description Language) is becoming increasingly popular as a way to capture complex digital electronic circuits for both simulation and synthesis. Digital circuits captured using VHDL can be easily simulated, are more likely to be synthesizable into different target technologies, and can be archived for later modification and reuse. VHDL features 3 levels of abstraction, which makes it useful for any part of the design cycle:

- Behavioral. Consists of a set of instructions that are executed in sequence to perform some task.
- RTL, Register transfer level. This describes how data is transformed as it is passed from register to register.
- Gate level. Describes each gate in the design including signal delays.

For implementation the register transfer level is normally used. The RTL standard defines a subset of the VHDL language, a subset which differs slightly between different synthesis environments. The Gate level is also possible to implement in hardware, but seldom used.

5.2.1 LINK SOURCE CARD

The Link Source Card (LSC) was programmed mainly by Zoltan Meggyesi. The source card exists in two version, single or double G-Links in the forward channel. In the single version a 64MHz on-board oscillator is used as transmission clock, and in the double channel a 40 MHz oscillator is used.

5.2.2 LINK DESTINATION CARD

The Link Destination Card (LDC) also exists in two versions, single and double forward channel G-Links. The double channel G-Link version is synchronous in the forward direction, i.e. the recovered clock from the G-Link is also used as the clock that is sent to ROMB as LCLK, and this clock is 40MHz. In the single version a 64 MHz clock is used as transmission clock, and since S-Link specifies the maximum LCKL frequency to be 40 MHz this clock cannot be used as LCKL. To achieve this the recovered transmission clock is divided by 2 internally in the Altera and this divided clock is used as LCLK at a frequency of 32 MHz. In this case, there is an asynchronous element in the forward channel of the LDC,

which is illustrated in figure 6-1. The divided clock is used internally as a global signal, as discussed in 3.2.5. By using the divided clock as a global signal timing performance is improved as the global signals has lower clock skew[6].

17-bit wide G-LINK data is clocked at the rising edge of the recovered clock RX_CLK and decoded into 33-bit S-LINK data (LD), which is clocked out at the rising edge of the divided clock LCKL. To ensure proper setup and hold time, LD must be stable for two RX_CLK cycles with a positive LCLK edge in the middle. At a rising edge of RX_CLK, LCKL is sampled and when high, LD is assigned its new value. When low, LD is assigned the old value. At the positive edge of LCKL the setup time (Tsu_LD) and hold time (Th_LD) is ensured with a good margin. All values in figure 5-1 is taken from MaxPlus simulations and are worst-case values.

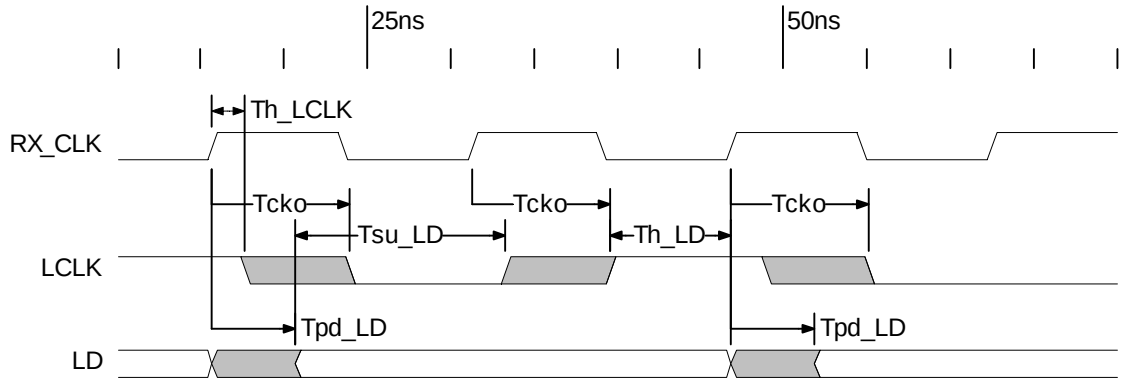

Row		Name	Formula	Min	Max	Margin	Comment
1	D	Tcko	[2,8.3]	2	8.3		Clock to output time XCLK->LCLK
2	D	Tpd_LD	[0,5]	0	5		Propagation delay time for data
3	C	Tsu_LD	[7,]	7		<5.63,>	Setup time for data sampled at rising edge of LCLK
4	C	Th_LD	[2,]	2		<5.33,>	Hold time for data sampled at rising edge of LCLK
5	C	Th_LCLK	[2,]	2		<0,>	Hold time for sampling LCLK state

Figure 5-1 Timing issues in LDC

5.3 SIMULATIONS

Both LSC and LDC was simulated separately as well as together prior to implementation using the Altera MaxPlus waveform simulation tool. Simulations are a way to check the logic for logical errors. This does not mean that all errors show up in simulations. Especially when new components are being used the actual hardware seldom behave the way it is expected to.

[6] Clock skew is the difference in time when the signal arrives at a logic element

5.4 TESTING

5.4.1 TESTING EQUIPMENT

The ODIN cards were tested with commercially available S-Link product. The SLIDAS was used as a data source and was mounted on a SLITEST board together with a ODIN LSC. The SLIDAS is programmed to send a variety of different test data such as walking one/zero, 00/FF, AA/55, and random data in block sizes of up to 8KB and clock rates from 1MHz to 40MHz. The internal S-Link test mode was also tested via a switch on the SLIDAS.

The receiving side consisted of a SLIDAD data drain mounted on a SLITEST together with a ODIN LDC. The SLIDAD does not feature an error detection apart from the error detection in the S-Link. Thus for better testing capabilities, the output from the LDC is also connected to a HP logic analyzer. The logic analyzer was programmed to detect the patterns transmitted by the SLIDAS as well as the S-Link test mode.

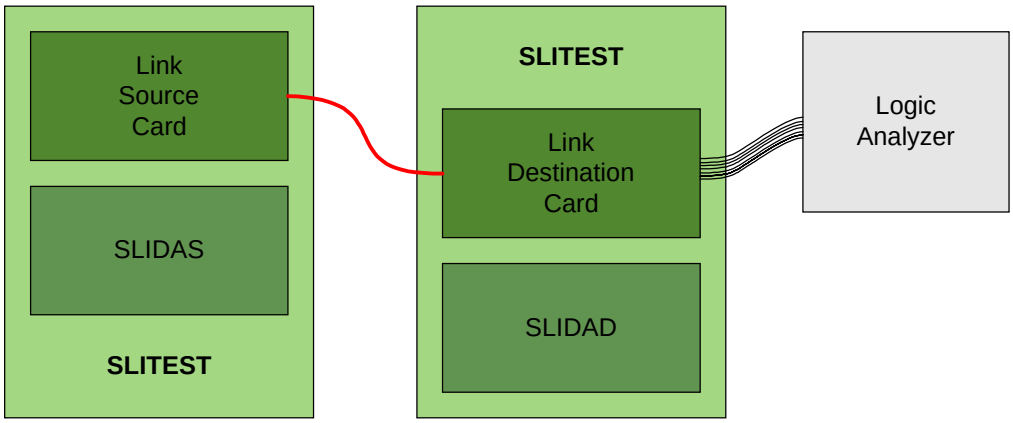

Figure 5-2 Test bench setup

Pulse generators with frequency in the range Hz – KHz were also connected to each side of the link. On the SLIDAS this starts and stops the data transfer. On the SLIDAD each pulse generated a flow control signal (actually each pulse turns the flow control off). These pulse generators are not shown in the figure above.

5.5 RESULTS

The ODIN S-Link Interface proved to work in the lab environment described above with no errors enountered. No cooling was needed, and the power consumption was typically 3.6W at 3.3V, although it differed for different test patterns and different transmission frequencies. The current drawn by the boards were did not differ much from the previous S-Link boards described in section 2.3. The lower power consumption comes almost entirely from the lower voltage.

The receiver sometimes lost synchronization while receiving the static AA/55 test data. This error was reproduced by HP in their test lab, but at the time of this report no solution had

been brought to my attention. By deactiveting the enhanced simplex mode mentioned in section 3.1 was a work-around solving this specific problem was found.

The G-Link specification does not mention any specific requirements that the receiver needs to synchronize to incoming data. However the receiver would not always lock correctly when sending. This problem was solved by introducing the 'semi-continuously' reset word described in section 4.5.2.

The maximum data transfer speed of the single channel ODIN is 128MBytes/s and for the double channel ODIN 160Mbytes/s.

For a more complete summary of the ODIN cards see the ODIN S-Link Hardware specification in appenix B.

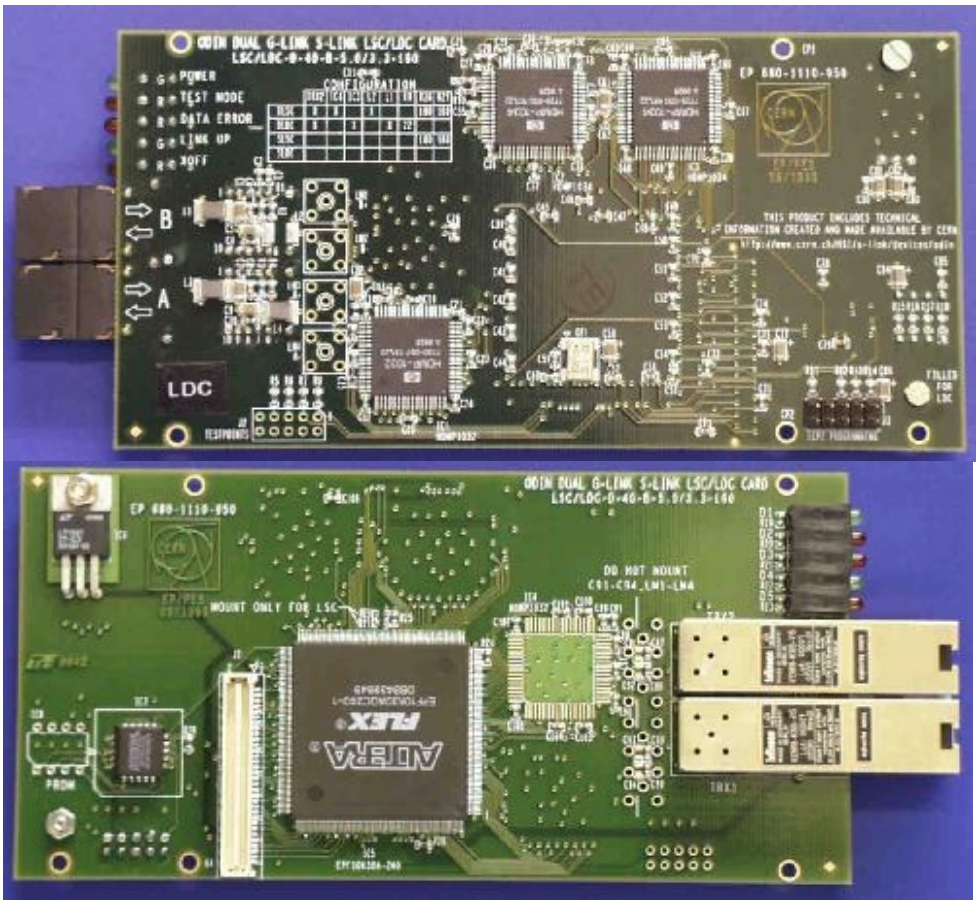

Figure 5-3 Pictures of the Dual ODIN Destination card

6 CONCLUSIONS

6.1 THIS PROJECT

The goal of the project was to design an optical data link according to the S-Link specification at a lower cost and running at greater transfer speed than previously designed S-Links. The aim was not to design the final data link for implantation in ATLAS, but to use newly developed commercial product in order to get one step closer to the final goal for ATLAS readout links.

The total price tag of the single Odin ended at 411CHF per card. Figure 7-1 shows the estimated price evolution for S-Links, with the Single Odin inserted. Further cost reduction must be made to reach 250CHF per card in component costs. The greatest achievement in this project, cost wise, was the new optical transceivers, which was a substantial part of the cost reduction.

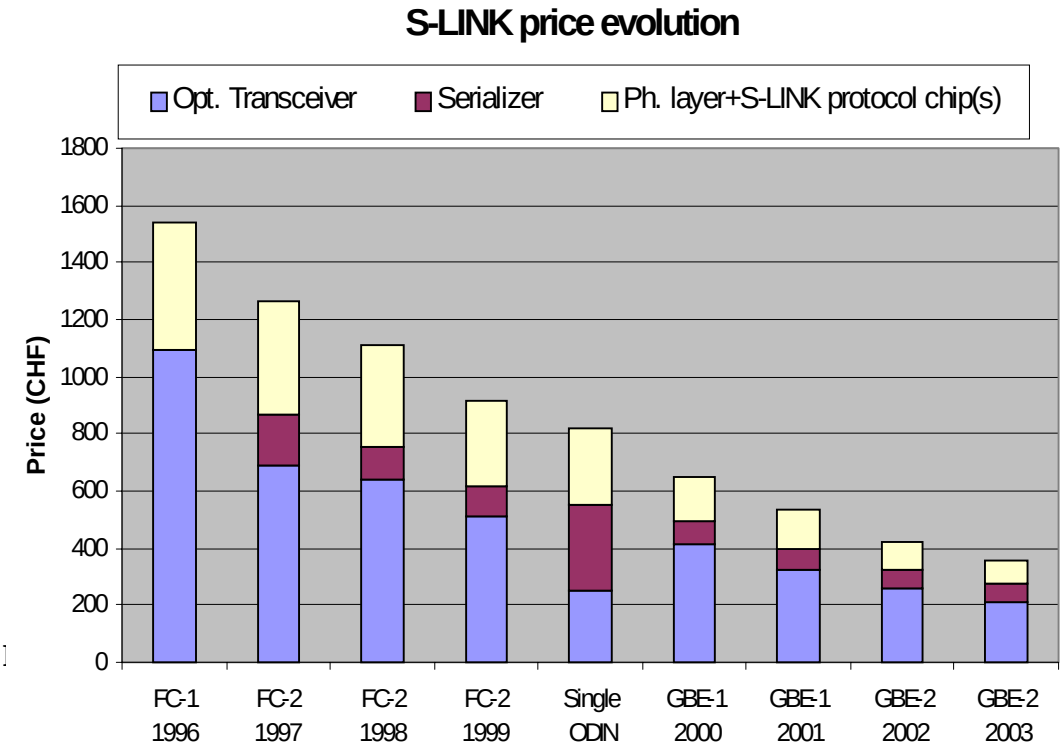

6.2 FUTURE WORK

Future cuts in cost can be made mainly in the serializer part. Even the single channel Odin consists of a total of 4 G-Links, which adds up to 300 CHF, almost half of the total cost. New PLD devices, such as the Altera APEX, include low speed LVDS serializers, which can be used for the return channel. A part of the cost reduction when going from a total of 4 G-Links to 2 will go into the higher price of the PLD, but there is still a total cost reduction.

Another future source of cost reduction is in the PLD. The two major parameters determining the price of a PLD is size, i.e. the number of usable gates, and speed grade. As stated in section 3.2.4 the current size is needed. By optimizing the design of the PLD however, a lower rated (speed grade wise) chip can be used.

7 REFERENCES

[1] ATLAS Technical Proposal, CERN/LHCC/94-43, December 1994

[2] ATLAS Technical Design Report , ATLAS TDR-12, June 1998

[3] Hewlett Packard HDMP-1032/1034 (G-Link) data sheet

[4] S-Link Specification,
`http://www.cern.ch/HSI/s-link/spec/spec/s-link.ps`

[5] S-LINK, a prototype of the ATLAS Read-out link, H.C. van der Bij et al. Presented at the Fourth Workshop on Electronics for LHC Experiments, Rome, 21-25 September 1998

[6] S-Link home page, `http://www.cern.ch/HSI/s-link/`

[7] ATLAS Technical Progress Report, CERN/LHCC/98-16, June 1998

[8] W.Stallings, "Data and Computer Communications", Third Edition, New York, Macmillan Inc.,1987

[9] Andrew Rushton, "VHDL for Logic Synthesis", Second Edition, John Wiley & Sons, Inc., 1998

[10] Lars-Hugo Hemert, "Digitala kretsar", Studentlitteratur

[11] Rabaey, " Digital Integrated circuits, A design perspective", Prentice Hall International, 1996.

Appendix A. Acronyms

ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal Lhc ApparatuS
CRC	Cyclic Redundancy Code
CRCC	Cyclic Redundancy Code Checksum
FPGA	Field Programmable Gate Array
G-LINK	Hewlett Packard serializer/deserializer chipset
LDC	Link Destination Card
LSC	Link Source Card
LEP	Large Electron Positron collider
LHC	Large Hadron Collider
PCB	Printed Circuit Board
PLA	Programmable Logic Array
PLD	Programmable Logic Device
SLIDAS	S-LInk DAta Source
SLIDAD	S-Link Data Drain
S-LINK	CERN standard for a 32bit@40MHz link
VHDL	VHSIC Hardware Description Language
VHSIC	Very High Speed Integrated Circuit

Appendix B. ODIN S-Link Hardware specification

Introduction

This document is intended as a designers note on the ODIN S-LINK interface and is for anyone who has to understand the internal functions of ODIN, or for anyone who wants to know more details than the ODIN data sheet features. The reader should have a profound knowledge

Features

Main Features

- Duplex version - **D**
- Max. 40 MHz User Clock (UCLK) - **40**
- Block basis error reporting - **B**
- 5.0V or 3.3 V supply voltage - **5.0/3.3**
- Maximum 160Mbytes/s - **160**
- S-LINK code - **LSC/LDC-D-40-B-5.0/3.3-160**

Single version features

- 32 MHz link clock
- 128 Mbyte/s maximum data rate
- 64 Mbytes/s maximum control word rate
- 8 MHz return line sampling

Double version features

- 40 MHz link clock
- 160 Mbyte/s data rate
- 80 Mbytes/s maximum control word rate
- 5 MHz return line sampling

Additional features

- Improved reset protocol
- 32 bit data width only

Forward Channel

The low power G-Link is run from an on-board oscillator since the receiver chip needs a reference clock for locking onto serial data stream. The G-Link features 17-bit parallel input, including flag bit. It also features a possibility to send a 14-bit control word. This section describes how the 33-bit S-LINK protocol is mapped to the 17-bit G-Link protocol. It also describes the internal commands the ODIN LSC and ODIN LDC use for communication.

Data Words

The G-Link protocol is defined on a word-basis only. Therefor there is no need to pack the data into frames when transferring and to reduce loss of bandwidth no extra framing protocol other than what the S-LINK control words specifies is used.

The 32-bit S-LINK data word is sent as two consecutive 16 bit G-LINK words. Flag bit is high for data words and no other odd/even word indicator is used. The most significant bits, LD[31..16], are sent first, followed by least significant bits, LD[15..0], on the next clock cycle. No idle words between msb and lsb are allowed. Idles between S-LINK words are allowed but not necessary.

See below for transmission example including control words and CRC

Data demultiplexing

In the Double ODIN version, two G-Links are used in the forward channel. The data splitting protocol is described in figure 1

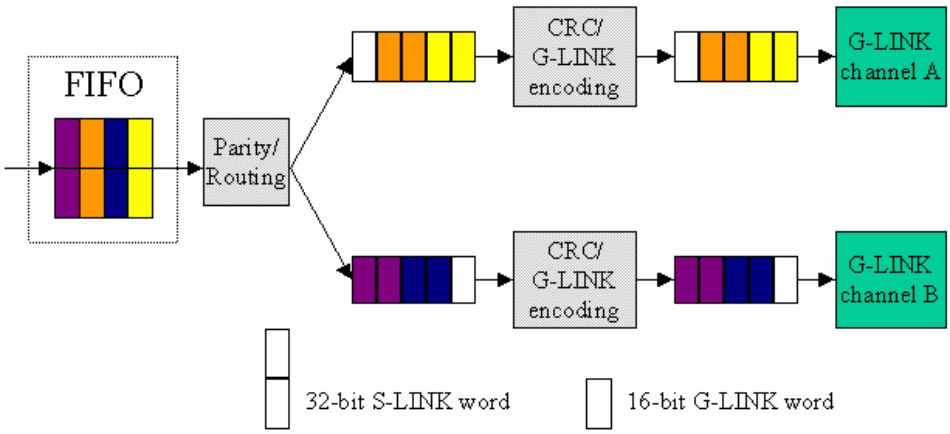

Figure 1. Data splitting in double G-LINK version

S-LINK data is written into the FIFO (left 2x4 block) at positive edge of UCLK, typically 40 MHz. The routing logic reads from the FIFO at positive edge of XCLK, the on-board oscillator, 40MHz.. As there is no framing it is important that not only upper and lower half of each word is sent in a defined order, but also that data from the A and B channels that are sent the same cycle also arrives at the LDC the same cycle. The cable length difference must

be small to meet this requirement. The B channel receiver in the double G-Link version is in PASS-mode, see HDMP-1032/34 for details, and this introduces an extra latency of 0.4 cycles, equivalent to ~2 meters of optical fiber. If cable lengths are different the B-channel should have the shorter cable.

In the single G-LINK configuration, only the A-channel (upper) is available, and the routing logic uses the 64 MHz on-board oscillator, reading at every other cycle. In this mode, data is written to the G-LINK at 64 MHz. On the single ODIN LDC the 64MHz receiver clock must me frequency divided to 32 MHz to meet S-LINK requirements of maximum 40 MHz LCLK.

CRC

CRC - Cyclic redundancy code.
CRCC - Cyclic redundancy code checksum, 16 bit checksum appended to data stream.

The 16 bit CRC-CCITT i.e. the polynomial $X^{16}+X^{12}+X^5+1$, is used as error detection for data words. A CRCC is sent in the following cases:

- Before a control word, if data has been transmitted.
- After average 1024 S-LINK dataword (4kB) of data per channel.

Note that a CRCC is not be sent when the LSC FIFO is empty. When UCLK is slightly below 40 MHz, or 32 MHz in single G-LINK configuration, this would cause a dramatic decrease in over all speed as seen in previous S-LINK implementations.

A CRCC is always sent over both channels and will always be preceded by a CRCC command. The error signal on the control words is ORed from the error signals from both channels, and the individual error signals are shown on LD[3..2] for diagnostic reasons. Since there is no begin-CRC-stream it is important that the CRC works in a similar manner on both cards. The CRC should be reset on RESET commands and CRC commands sent over the forward channel. When resetting a CRC the quotient can be set to arbitrary values and the ODIN link the value of all ones is chosen for simplicity and the ability to count commencing zeroes, which a CRC preset with all zeroes does not have. It is common practice to reset CRC quotient to '1'.

Since the error detection is only visible in control words, every transmission of data should be ended with a control word, and a CRC checksum is sent and CRC logic on both cards are reset. Figure 2 shows an example when a data block of 4-8kB of data is sent over one channel, followed by a control word.

Figure 2. Transmission example with CRC

Control Words

An S-LINK control word (28-bit word) is sent as a 2 G-LINK data words with flag bit low for both cycles. If data have been sent since last control word, a CRC checksum is sent over all available channels.

Even parity is used as error detection for control words. Even parity is chosen so that a control word containing all zeros (for example in test mode) will have all parity bits also zero. Since control word is sent as 2 G-Link data words, 4 unused bits are available. These are LD[3..0] and these are

LD[31..4]	LD[3]	LD[2]	LD[1]	LD[0]
Control Word	Even parity bit for LD[31..25]	Even parity bit for LD[24..18]	Even parity bit for LD[17..11]	Even parity bit for LD[10..4]

Figure 3. Control Word with parity bits

The control word, including the parity bits, are now sent in the same manner as data word, with LD[31..16] followed by LD[15..4]+parity the next cycle.

At LDC, parity bits are checked, and according to S-LINK specification, these lowest four bits are reserved for error detection report.

Command Words

Internal S-LINK commands are sent as G-LINK control words. The control bits have the following meaning: (bit 14 and 15 used internally by the G-LINK). If an error is detected in any bit, the LDC should disregard the word. Rx[13..10] are reserved for future use and current receiver disregards these bits.

Rx[9..0]	Symbol	Command
"0000000011"	CRCC	Next word a CRCC
"0000001100"	TON	Test mode on
"0000110000"	TOFF	Test mode off
"0011000000"	RLDWN	LSC down
"1100000000"	RRES	Remote reset

Figure 5. Internal ODIN commands

G-Link Flag bit

The G-LINK flag bit is used to distinguish between control word and data word. For data words flag bit is high during transmission, and for control words flag bit is low. For CRC the flag bit is high during the transmission of the CRCC, so flag bit also indicates that the received word is a part of the CRC data stream.

Transmission	TX_DATA	TX_CNTL	TX_FLAG
Data word	high	low	high
Control word	high	low	low
ODIN Command	low	high	n/a
CRC checksum	high	low	high

Figure 6. Overview of G-LINK control bits.

Test mode

The ODIN uses the standard S-LINK test pattern, the walking bit pattern. The test mode is started by sending a 'test-mode-on' command and ended with a 'test-mode-off' . The return channel is working as normal in test mode. Each test mode cycle is commenced by a test-mode-on command

Transmission example

Here is a longer transmission example for double channel version. For single channel version the main difference is that only one CRC is sent. The example contains the following data:

- Beginning of fragment Control word, CW1
- 4 S-LINK Data Words DW1-DW4
- 2 CRC command+CRCC, one per channel
- End of fragment Control word, CW2

Channel A				Channel B			
TX/RX[15..0]	**FLAG**	**DATA**	**CNTL**	**TX/RX[15..0]**	**FLAG**	**DATA**	**CNTL**
CW1[31..16]	0	1	0	x	x	x	x
CW2[15..4]+par[3..0]	0	1	0	DW1[31..16]	1	1	0
DW2[31..16]	1	1	0	DW1[15..0]	1	1	0
DW2[15..0]	1	1	0	DW3[31..16]	1	1	0
DW4[31..16]	1	1	0	DW3[15..0]	1	1	0
DW4[15..0]	1	1	0	CRC command	-	0	1
CRC command	-	0	1	CRCC	1	1	0
CRCC	1	1	0	CW1[31..16]	0	1	0
x	x	x	x	CW2[15..4]+par[3..0]	0	1	0

Figure 7. Longer transmission example

Return Channel

The return channel also uses a G-LINK as physical layer. All words are sent as data word, no flag bit used. For error detection every bit is sent twice. If an error is detected in the word the word is discarded. No flag bit or control word is used.

Bits	**When "11"**
[1..0]	RL[0] high
[3..2]	RL[1] high
[5..4]	RL[2] high
[7..6]	RL[3] high
[9..8]	Flow control, XOFF
[11..10]	LDC down

[13..12]	Remote reset
[15..14]	Reserved

Figure 8. Return channel data and commands

Reset protocol

The reset protocol is changed from the S-LINK specified *card reset* to a *link reset*. When either side detects a problem with the link, i.e. the receiver(s) do not acquire lock on serial data, a link down command is sent to the other card, making it go down. If URESET# is set low on either side, a Remote reset command (RRC) is sent to the other side, making both cards go up when the error is cleared.

LSC Power-up and Reset State machine

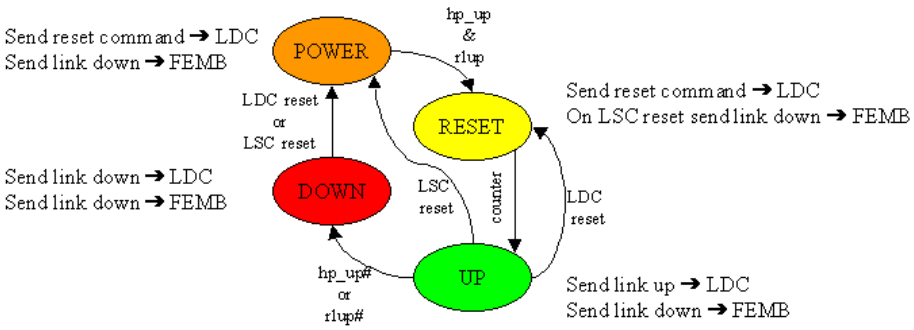

LSC powers up in POWER state. From there it waits for a hp_up, 21-bit filter signal on Transmitter locked and Receiver ready signal, and rlup, indicating that a command other than link down is received on return channel. Counter in RESET-UP transition is reset after a LSC reset and assures link down to be 4 clock cycles at a LSC reset.

When link is up and LDC is reset, LSC goes into RESET state to answer reset command. This is necessary since CRC and LDC is reset only on this LSC reset command. When answering reset command the reset counter is not started, and LSC will stay up during LDC initiated reset cycle.

In POWER and DOWN state, 1 command word is sent followed by 7 idles. This is to make LDC G-Link receivers to acquire faster and more secure lock, as it prevents improper word alignment caused by static valid code field embedded within the data field that the command words have..

LDC Power-up and Reset state machine

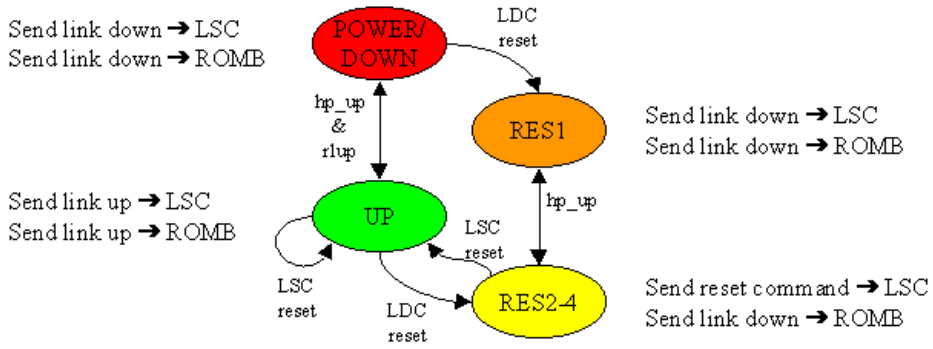

LDC powers up in POWER/DOWN state. Here the LDC waits for hp_up and rlup, and is continuously sending link down command to LSC. By doing this, it is assured that LDC will come up before LSC in power up sequence, and no data written to LSC will be lost. When LDC is in UP state, a LSC reset will do nothing to the state machine, but CRC error latches and Test mode state is cleared.

Return channel always send 1 word followed by 7 idles to make LSC G-Link receiver to acquire faster and more secure lock.

Power-up sequence, LSC powered up first

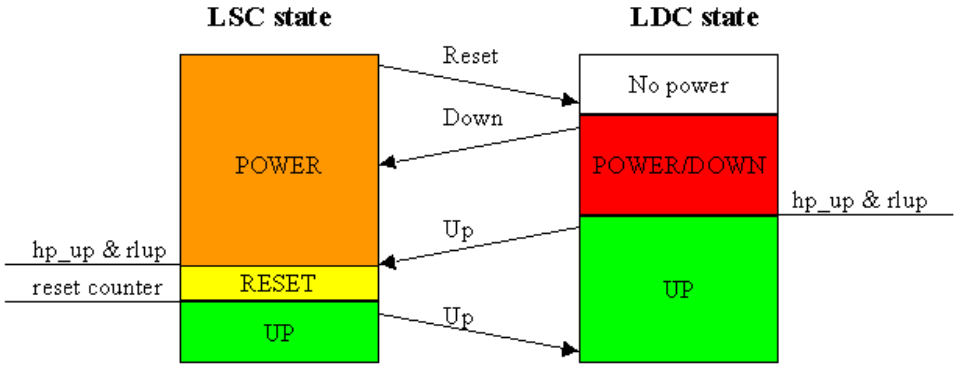

Power-up sequence, LDC powered up first

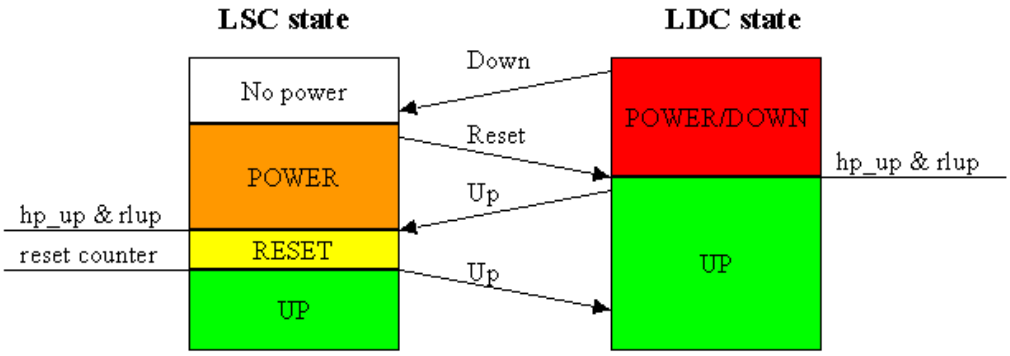

Clock domains

LSC

LSC has three more or less independent clock domains. **UCLK** is the user clock, maximum 40MHz,and this is decoupled from the internal logic through the asynchronous FIFO. However, at LSC reset a counter is run from this clock to assure 4 cycles of LDOWN# low, which is specified in the S-LINK spec. **XCLK** is the transmission clock from most of the LSC logic runs from. **RXCLK** is the recovered return channel clock, and since return lines are specified to be asynchronous even relative each other, no fancy synchronization is needed when going to XCLK domain. XCLK and RXCLK have the same frequency, but originate from different oscillators, so these will drift relative each other.

LDC

Single channel LDC has three clock domains, double channel only two. These are **RXCLK**, the recovered clock from the forward channel. G-Link data logic runs from this clock. In the single channel version the transmission clock is 64MHz, and too high frequency for **LCLK**, and its frequency is divide by 2. In double channel version RXCLK and LCLK is the same clock, including phase. The last clock domain is XCLK from which the state Power and Reset state machine is run.

The figure below describes the synchronization from RXCLK to LCKL domains in the single channel LDC. The SGMUX must keep the output, LD, stable for two cycles of RX_CLK and is only allowed to change when LCLK is high. Timing designer shows that the setup and hold times are valid.

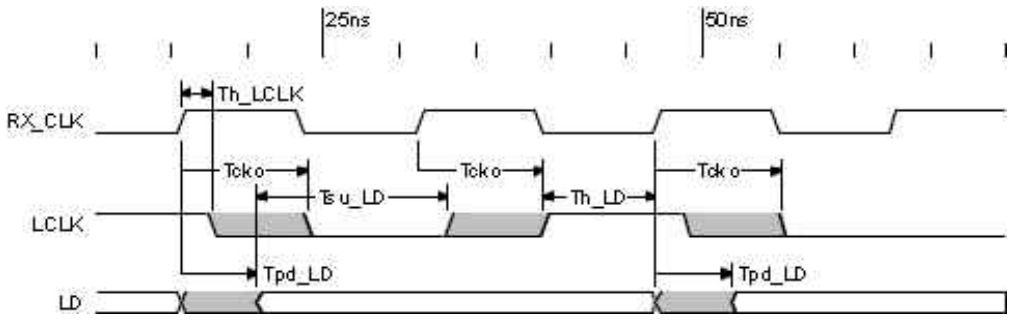

Symbol	Description	Min	Max	Units
RX_CLK	clock frequency		64	MHz
divclk	divided clock frequency		32	MHz
Tpd_LD	Data propagation delay time	0	5	ns
Tcko	Clock to output time for LCLK	2	8.3	ns
Tsu	LD setup time	7		ns
Th	LD hold time	2		ns
Th_LCKL	hold time for sampling LCLK	2		ns

Analyzed with timing designer gives a buffer at LD setup time of 5.6 ns, and at hold time of 5.3 ns. The values of Tpdck and Tcko are taken from MaxPlus timing analyzer. If major changes are done to the design these calculations should be re-evaluated.

Bugs - Room for improvement

PCB

Loop-filter capacitors:
C27, C63, C13 and C95 should never be mounted according to final G-Link specification.

Serial clock signal termination:
R9, R10, R22, R24, R25 should be 68 ohm instead of 22ohm to reduce overshoot in clock signal.
Measurements clock signal quality should be made for further optimizing this resistor value.

Motherboard need a dedicated clock signal for LSC. Better signal termination on LSC UCLK could give better performance.

Swapped cables

In current version, LDC will come up even when cables are swapped, but LSC will stay down. This is unfortunate since LDC Link Up led is more accessible than the LSC Link Up led. One solution to this problem would be to specify different reset signals for channel A and channel B. Thus neither channel would see a reset if the cables are swapped.

LFF# and LDOWN#

User is only supposed to write data when LFF# is high, i.e. FIFO not full. When link is down LFF# should be high to assure that no data is written to LSC. Probably easy to implement by an AND-gate between LFF# and LDOWN# internally in LSC FIFO.

Error recovery with different fiber lengths

In double channel version fibers must be of the same length within 1 meter. If there is greater length difference, it is possible that the slave receiver (B channel) data is delayed one clock cycle. The performance in this case could be greatly improved by easy means. In the dbmux block of the LDC, the state machine always look at channel A first for data. If channel B data is delayed there will be data available the same clock edge, and then should B channel data be chosen, and A-channel data should be registered and sent next clock cycle.

Error reporting

Use of LD[3..2] in control words (signals error in channel A or channel B) should be removed.
File to change: PARRX.VHD

Test mode

Flow control during test mode does not work when UCLK is <10MHz as LSC FIFO gets stuck..

Sometimes in the double channel version, when leaving test mode, a control word (UCTRL# low) 01000000 is written before link goes up. The error probably lies in the FIFO logic assignments for dbldc. The error is no longer reproducible.

Return channel does not feature a test mode. Should be implemented.

Transmission frequency

One of the ideas of ODIN is to get the highest transmission rate and still be running within all the specs. XCLK could go up to 70MHz without violating any specifications. This should be tried.

Appendix C. ODIN S-Link Interface data sheet

LSC/LDC-D-40-B-5.0/3.3-160

Version 1.0

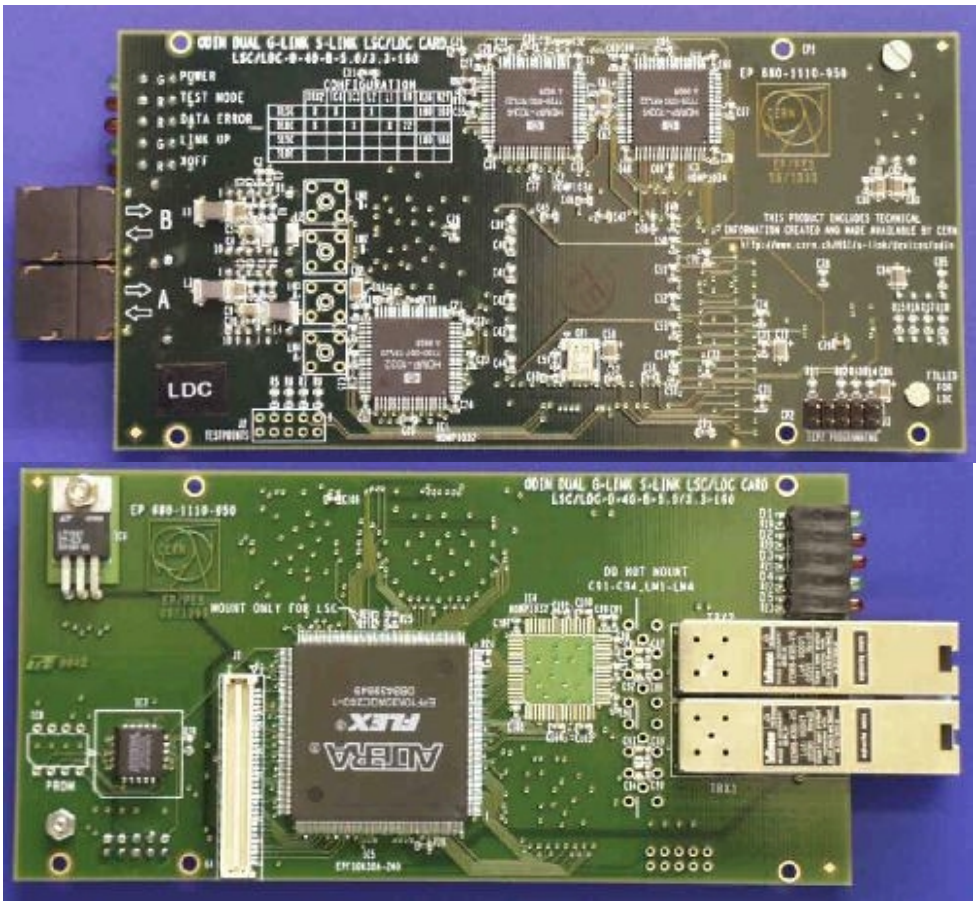

1. Introduction

The ODIN S-LINK is a standard duplex S-LINK that uses the low power G-Link chip-set (HDMP-1032/34) as physical layer. Its components are the Link Source (LSC), the Link Destination (LDC) cards and a fiber optic cable with Volition duplex connectors. More information, including this document, about the ODIN Interface can be found on the World Wide Web at:

http://www.cern.ch/HSI/s-link/devices/odin/

The reader of this document should know the basics of the S-LINK specification. This data sheet only point out a few important features of ODIN and makes a few clarifications. The S-Link specification can be downloaded from the Worldwide Web at:

http://www.cern.ch/HSI/s-link/spec/

2. Main Features

ODIN exists in two different operation versions, single and double channel, where one or two G-Link chip sets are used in the forward channel. These two versions cannot be mixed; i.e. a single channel LSC will not work with a double channel LDC and vice versa. ODIN also exists in two different voltage versions, 3.3V and 5V. See figure 2-1 for block diagram of the double channel version and the usage of the two optical fibers.

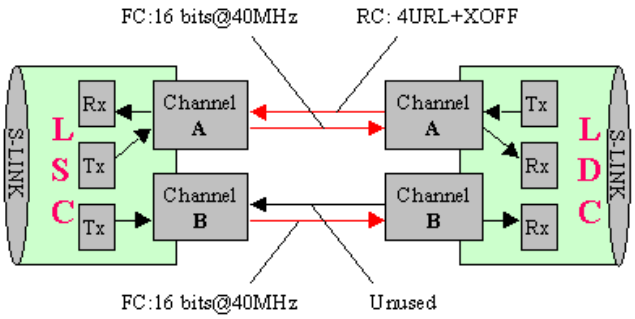

Figure 2-1 Double channel ODIN

The S-LINK code for the ODIN cards is:

LSC/LDC-D-40-B-5.0/3.3-160

The main features of the ODIN S-LINK implementation are:

General features

- Duplex S-Link
- 32 bit data width

- UCLK up to 40 MHz
- Block basis error reporting on data words
- Word-by-word error reporting on control words
- 5V tolerant input signal levels
- Option of 3.3V or 5.0V power input PCB versions
- Power consumption typically 5.5W, max. 7W at 5.0V, 3.6W, max. 4.5W at 3.3V
- Autonomous link synchronization and maintenance of physical link
- Improved S-LINK Reset protocol
- Flow control is provided both in data mode and self test modes
- Optical output, max. cable length with 50μm multimode cables: 550m

Double channel version

- 40 MHz LCLK
- 160 Mbytes/s maximum data transfer rate
- 80 Mbytes/s transmission rate for control information (UCTRL# low)
- 5 MHz sampling rate for the return lines
- Maximum 1m fiber length difference

Single channel version

- 32 MHz LCLK
- 128 Mbytes/s maximum data transfer rate
- 64 Mbytes/s transmission rate for control information (UCTRL# low)
- 8 MHz sampling rate for the return lines

3. Installation

Fix the ODIN S-LINK cards with screws to standoff pillars and front panel mounting holes to avoid mechanical stress and contact problems. For proper operation the board should be fixed by all four screws. Avoid mechanical stress on the cards during mounting. Do not plug in and out the cards on the motherboard when this is powered on.

Before powering up, check that the voltage version of the ODIN is the same as the motherboard. 5V motherboard should have a voltage keying pin described in S-LINK specification [1] and the 3.3V ODIN should not have a voltage hole to eliminate the chance of mounting 3.3V S-LINK on a 5V motherboard. Mounting 5V ODIN on a 3.3V motherboard is physically possible, but will not work.

Double ODIN features two optical connectors, marked A and B, and cables must not be switched. Maximum length difference is 1 meter. If cables are different in length channel B should have the shorter cable for better word alignment tolerance.

Connect the optical fiber(s) and after powering up the cards the link goes up immediately and Power and Link up LED's go on. There is no need for a reset at power up as the cards go up when powered on and fibers are connected.

4. Operating Conditions

Symbol	**Description**	**Min**	**Typical**	**Max.**	**Units**
$I_{cc,d}$	Current drawn, single ODIN	700	1000	1200	mA
$I_{cc,d}$	Current drawn, double ODIN	1000	1200	1400	mA
V_{cc} (3.3V board)	Voltage	3.1	3.3	3.5	V
V_{cc} (5V board)	Voltage	4.5	5.0	6.0	V
T_{op}	Temperature	0	25	70	C

Table 4-1 ODIN operating conditions

5. Timing characteristics

Table 5-1 gives the required timing parameter for LSC for proper operation. See [1] for explanation of the parameters.

Symbol	Description	Min	Max	Units
t_{DS}	Data Set-up time	10		ns
t_{DH}	Data Hold time	1		ns
t_{ENS}	Enable Set-up time	10		ns
t_{ENH}	Enable Hold time	1		ns
t_{WFF}	Write Clock to Full Flag		12	ns
t_{CLK}	Clock Cycle time	25		ns
t_{CH}	Clock High time	11		ns
t_{cl}	Clock Low time	11		ns

Table 5-1 LSC timing parameters

Table 5-2 gives the guaranteed timing parameters for LDC. Suffix -single relates to single channel ODIN with 64 MHz transmission frequency and 32 MHz LCKL. Suffix -double relates to double channel ODIN with 40 MHz transmission frequency and 40 MHz LCLK.

Symbol	Description	Min	Max	Units
t_{DS}	Data Set-up time	10		ns
t_{DH}	Data Hold time	1		ns
t_{ENS}	Enable Set-up time	10		ns
t_{ENH}	Enable Hold time	1		ns
t_{WFF}	Write Clock to Full Flag		12	ns
$t_{CLK\text{-}single}$	Clock Cycle time	31		ns
$t_{CH\text{-}single}$	Clock High time	13		ns
$t_{CL\text{-}single}$	Clock Low time	13		ns

$t_{CLK\text{-}double}$	Clock Cycle time	25		ns
$t_{CH\text{-}double}$	Clock High time	11		ns
$t_{CL\text{-}double}$	Clock Low time	11		ns

Table 5-2 LDC timing parameters

6. Optical characteristics

Symbol	Description	Min	Typical	Max.	Units
	Link length (50/125 microns)	2		550	m
	Link length (62.5/125 microns)	2		260	m
BER	Bit Error Rate			10^{-12}	
lambda	Center wavelength	830	850	860	nm

Table 6-1 Optical characteristics

7. LED Indicators

LED symbol	Color	Function at LSC	Function at LDC
PWR	green	Power On	Power On
TST	red	Self test Mode	Self test Mode
ERR	red	-	Data Error
UP	green	Link Up	Link Up
XOF	red	Link Full Flag	Flow control active

Table 7-1 LSC and LDC LED indicators

8. User Data Width Lines

The UDW input lines are unused. If ODIN is set to 16 or 8-bit mode the effective transmission rate will go down accordingly since all 32 bits always are transferred.

9. Data Transfer

The ODIN S-LINK will transfer all input data when up, Power and Link Up LED's are on and LDOWN# lines are high on both sides. If data is written to single channel LSC faster than 32 MHz the LFF# flag will go low and the XOFF led goes on at the LSC, even if there is no flow control sent from LDC. This is because maximum data transmission rate is 128 Mbytes/s. For maximum transfer rate, 2 more words should be written to LSC after LFF# goes low.

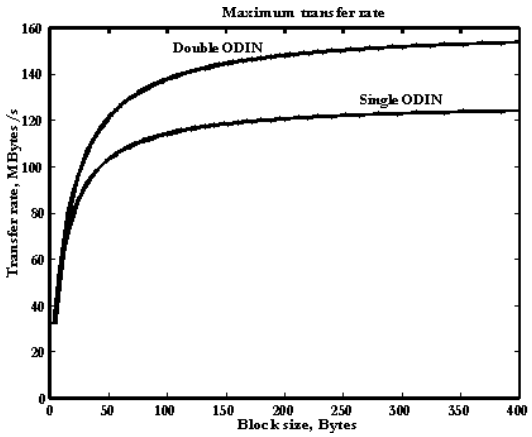

Figure 9-1 Maximum transfer rate at different data block sizes with 40 MHz UCLK

For small data block sizes the LFF# line will also become active as error detection and S-LINK control words take up bandwidth. Figure 9-1 gives maximum average transfer rate for small block sizes. The block size refers to the ATLAS event data format where one Begin of Fragment control word and one End of Fragment control word frames the data words. These transfer rates given are maximum values and to reach the maximum value 2 more words must be written to LSC after LFF# goes low. Also note that the graph shows the average values while the nominal value inside a block is the maximum transfer rate.

10. S-LINK Reset

ODIN S-LINK features an improved reset protocol, which is backward compatible. The user may reset the whole link from either side, or even both sides. The card reset is changed into a link reset to make the reset an easy and reliable operation regardless of usage.

If the link is down prior to the reset cycle, LDC will come up before LSC, regardless of which side the reset is performed. This eliminates the chance of data written to LSC being lost because of LDC being down.

If the link is up prior to the reset cycle, the card where URESET# line is asserted will go down according to [1]. The other side will be reset without going down.

It is not recommended to perform a reset while writing data to the LSC, as this will cause data loss.

11. Test Mode

Self-test mode fully complies with S-LINK specification. UTDO# line is sampled at link reset.

12. Flow Control

The Read out buffer size can be calculated from the formula given in [1]:

$$RBS = \left(\frac{LDC_{rt} + [L \times UFD \times 2] + LSC_{rt}}{DTR} \right) \quad [1]$$

Where:

RBS = Read out Buffer size (words)
LDC_{rt} = LDC reaction time to send XOFF after UXOFF# goes low (ns).
LSC_{rt} = LSC reaction time to stop transmitting data after XOFF received (ns).
L = Length of S-LINK (m).
UFD = Unit Fibre Delay - time for light to travel 1m in fibers (approx. 6 ns/m).
DTR = Data Transfer Rate (ns/word).

Data words in LSC and LDC pipelines are included in LSC_{rt} and LDC_{rt} respectively.

Symbol	**Single ODIN**	**Double ODIN**
LDC_{rt}	350	550
LSC_{rt}	200	300
UFD	6	6
DTR	31.25	25

Table 12-1 Flow control parameters

The values for the ODIN boards are given in table 12-1, which yields the formulas:

$$RBS_s = 20 + 5L/13 \text{ (Single ODIN)}$$

$$RBS_d = 40 + L/2 \text{ (Double ODIN)}$$

13. Return Lines

Return Lines are functional during data transfer and in test modes. LRL[3..0] lines will stay unaltered when the link is down. An internal parity checking logic ensures proper operation of the Return Lines. The sampling rate of the return lines is 8 MHz for the single channel version and 5 MHz for the double channel version.

14. Error Detection

ODIN features a CRC-based block error detection and errors are reported in the following control word, as specified in [1]. There is also a weak word-by-word error detection, but this only looks at the internal data format and G-Link error bit, and does not look at the actual data

S-LINK control words uses parity bits as error detection in order to separate data errors from control word errors.

15. Link Down Function

The following events may result the link down signal to be asserted:

1. Reset cycle
2. Self-test mode
3. Local reset, i.e. the LDOWN# is asserted only on the card where URESET# is set low.
4. LSC or LDC is not powered up
5. Broken optical link
6. Fatal error occurred

The three first cases are covered by the S-LINK specifications. In the other cases the link down should be latched until a reset cycle. This is true except on LSC power down as the LSC is powered up in a reset state. On LDC power down, LSC will go down caused by the return channel loosing synchronization. This link down state will be latched until cleared by a reset at either side.

16. Known Bugs

For UCLK in low MHz range, <10 MHz, flow control does not work in Test Mode, as this sometimes will cause a link down. Going out of test mode and perform a link reset clears this state.

References

[1]. O. Boyle, R. McLaren, E. van der Bij, "The S-LINK Interface Specification", http://www.cern.ch/HSI/s-link/spec/spec/ CERN, 1997.