

Atlantis MultiRob Scenario

The ATLANTIS system is described in the attached document **robscenario.pdf** and can be viewed as a standard PC with a number of FPGA co-processors. However due to the compactPCI board format a larger number of co-processors is supported and each of these is capable of handling up to 4 electrical links @160MB/s (M-Links, compatible to one specific S-Link implementation).

The three basic functions of a Rob Complex (as seen by Level-2) RobIn, RSI, and RobC(ontroller) are usually (in the present PP studies) distributed in a way that RSI and RobC are handled by a single host cpu while the RobIn function is handled by a dedicated co-processor (of any kind).

By adding more than one RobIn module per cpu one creates what we call a MultiRob. The benefit one can gain from a MultiRob-Complex can be limited by a number of factors:

1. Bandwidth of the host bus serving the input links (from RobIn's)
2. Bandwidth of the host bus serving the output link (to RSI) [if shared with input bus]
3. Bandwidth of the network link attached to the RSI
4. Realtime response characteristics of the host CPU (probably dominated by OS issues)
5. Processing power for managing input links (buffer management, preprocessing etc.)

Most of the items above (additionally) suffer from the high message rates we expect for the RobComplex. The bandwidth of the input links is not considered critical in the scope of the RobComplex.

The basic goals of the Atlantis MultiRob project are to

- Demonstrate a full MultiRob Complex, with real input and output,
- Measure the benefit of local (on the co-processors, and eventually via a private bus) managing and preprocessing (=> multi RobIn), compared to a single RobIn or a collection of single RobIn's on a single bus,
- Investigate the effect of the reduced message rates allowed by a RoI based rather than RoI-Fragment based traffic (for both RobC and RSI).

ROB Scenario on Atlantis

1 Introduction

This document describes a scenario for a ROB complex on the FPGA processor system ATLANTIS. The system consists of ROB-IN, ROB-Controller and ROB-OUT (RSI). Section 2 gives a summary about the requirements of the trigger system to a ROB Complex. Section 3 describes the setup of a ATLANTIS based ROB. Section 4 presents the ROB-IN based on the ATLANTIS I/O Board, Section 5 the ROB-Controller and Section 6 the ROB to switch interface.

2 Requirements

This section is a summary of requirements for a ROB complex regarding the different subdetectors, the Level 2 trigger and the Event Filter.

2.1 Data from the ROD's

The fragment size coming from the ROD's for the different subdetectors is shown in table 1. Assuming a LVL1 acceptance rate of 75 kHz the data rates between ROD and ROB-IN can be computed.

Subdetector	Typical Fragment Size /event	Data Rate per ROB-IN	Number of Links
TRT	8kbits	75MB/s	256
SCT	12.5kbits	117MB/s	96
Pixel	max 6.1kbits	57.2MB/s	84
EM Calorimeter	14kbits	134MB/s	760
Hadronic Calorimeter	7kbit	67.5MB/s	64
Precision Muon	2-6 kbit	18,8MB/s-56.3MB/s	192
Trigger Muon	3kbit	28.7MB/s	48

Table 1: Eventsize and data rates transferred from ROD to the one ROB-IN at high luminosity [2]

2.2 Data to LVL2 and EF

The data rates at the ROB output to the LVL2 trigger are different for the several subdetectors. This depends on the trigger menu which is used to examine the event. To get an overview about the data rates between ROB and LVL2 refer to the Modelling parameters of Jos Vermeulen[3]: For the different subdetectors the data rates of this model are summarized in Table 2 and 3.

The data rate to the Event Filter depends on the LVL2 acceptance rate. The full data of every accepted event is transmitted. Assuming a LVL2 acceptance rate of 1.4 kHz [4] leads to the data rates in table 4.

Subdetector	Maximal data rate per ROB-IN
TRT	<0.5MB/s
SCT	1.5MB/s
Pixel	1.1MB/s
EM Calorimeter	7.9MB/s
Hadronic Calorimeter	3MB/s
Precision Muon	<0.5MB/s
Trigger Muon	<0.5MB/s

Table 2: Maximal data rates at high luminosity from modell[3]

Subdetector	Maximal data rate per ROB-IN
TRT	6.5MB/s
SCT	18.6MB/s
Pixel	12MB/s
EM Calorimeter	5.9MB/s
Hadronic Calorimeter	3.9MB/s
Precision Muon	<0.5MB/s
Trigger Muon	<0.5MB/s

Table 3: Maximal data rates at low luminosity from modell[3]

Subdetector	Data rate to the EF per ROB-IN
TRT	1.4 MB/s
SCT	2.2 MB/s
Pixel	1.1 MB/s
EM Calorimeter	2.45 MB/s
Hadronic Calorimeter	1.2MB/s
Precision Muon	0.4 - 1MB/s
Trigger Muon	0.5 MB/s

Table 4: Data rates to the Event Filter

3 Setup for a ATLANTIS based ROB Complex

The backbone of the ATLANTIS system is a Compact PCI crate with up to 14 PCI slots. One of these slots is always used for the host CPU which is a normal PC ¹. At this time we use a 8 slot PCI backplane. So 7 slots are left to take Atlantis I/O boards. These act as four port ROB-IN's.

A privat bus system connects all I/O boards independent of the PCI Bus. Over this second bus they communicate with a data rate up to 0.5GB/s with each other. The Host CPU acts as a ROB-Controller for all ROB-IN's. The ROB ouput port (RSI) is a network interface card (Ethernet). The whole setup can be seen in figure 1

4 The ROB-IN

The ATLANTIS I/O board acts as a four port ROB-IN in the presented ROB Complex. The main task of the I/O board is the buffering and preprocessing of data recieved from the ROD's. It acts as ROB-IN. A schematic structure can be seen in figure 2.

¹In the actual configuration a 200MHz Pentium System with 128MB memory, harddisk, video adapter and fast ethernet.




Figure 1: The ATLANTIS based ROB Complex

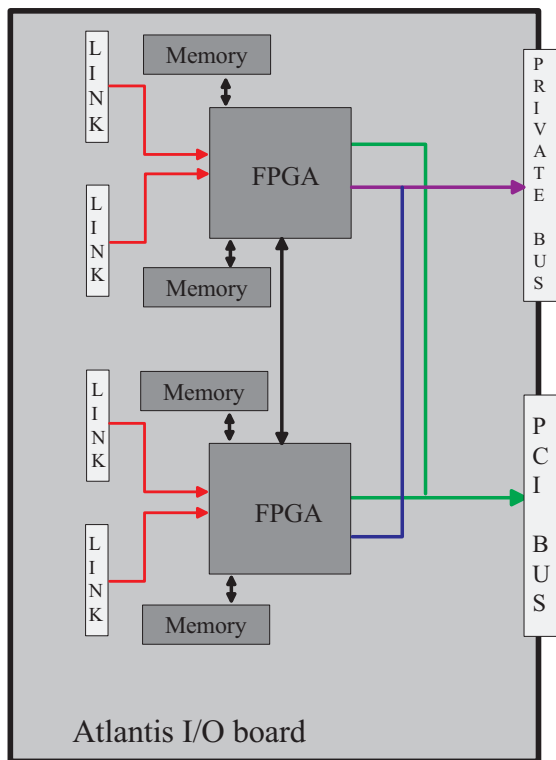


Figure 2: The ATLANTIS I/O board

4.1 Input links

The board is equipped with up to four MLink input ports. MLink is a LVDS based interface card which is compatible to the electrical SLink interface (SCSI). The difference to SLink is a mechanically smaller interface card. Optional two standard SLink interface cards can

be used with the ATLANTIS I/O Board instead of the MLink cards.

With 40 MHz SLink clock the maximal speed per input interface is 160MB/s which is the upper limit of the specification at this time.

4.2 FPGA's

Two FPGA's on the I/O board are processing the incoming and outgoing data. Two input channels are dedicated to one FPGA. The design for one FPGA contains two input processes with memory management. So two FPGA's have four processes for storing data coming from 4 input links. The second task for the FPGA's is the preprocessing and the output to the PCI bus. One output task is used for the whole I/O board.

4.3 Memory

The memory consists of four banks dual port SRAM with a datawidth of 32 bit, two per FPGA, which are independent accessible for each input process. The memory size depends on the event fragment size coming from the ROD's, the event rate, the LVL2 processing time and the LVL2 acceptance rate with the EF readout time[1]:

$$MEM = EFSZ \cdot (ER \cdot LVL2_{pt} + LVL2_{ar} \cdot EF) \cdot SAVE$$

EFSZ : Event fragment size
ER : Eventrate from ROD
LVL2_{pt} : Level 2 processing time
LVL2_{ar} : Level 2 accept rate
EF : Time between LVL2 accept and arrival at EF
SAVE : Safety factor

This leads to a memory size of 1 MB. For the I/O board are 4-8 MB per input link provided.

The maximum transfer rate to memory is 160 MHz (40MHz clock speed).

4.4 Output to PCI

The output process reads the requested event out of the buffer memory and preprocesses it. This data package is send over the PCI Bus to the Host CPU.

If a requested ROI requires the data of several input ports, the output process merges the data to one big package and passes it over the PCI bus to the Host CPU. This reduces the number of small packages on the PCI bus and increases the PCI data rate. It reduces also the number of messages from the ROB Controller to request data, because only one ROB-IN board must be accessed. Measurements with the microEnable ROB-IN, whose PCI interface is the same as on ATLANTIS Boards, show that at low data package sizes the PCI bandwidth increases dramatically with the package size.

With a standard PCI bus (32 bit, 33 MHz) each of the seven I/O boards and output processes get a maximal bandwidth of 19 MB/s.

This is not sufficient for the most subdetectors particularly with low luminosity (table 4 and 3). So the number of I/O Boards should be matched to the required output datarates. The PCI bus will be one bottleneck of the ATLANTIS based ROB Complex. Using a 64 bit bus with 66 MHz increases the bandwidth of the PCI bus to 528 MB in future systems and eliminate the PCI bottleneck.

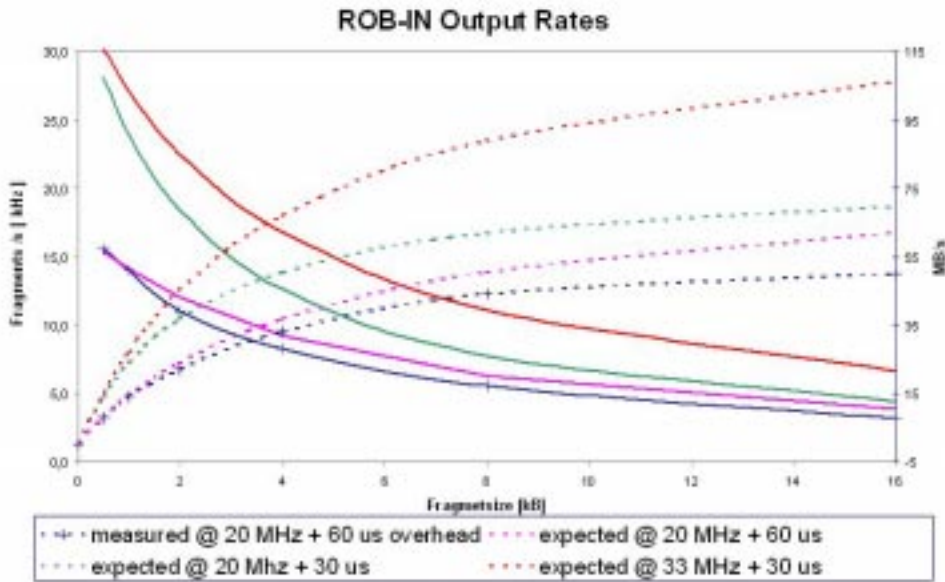


Figure 3: PCI output bandwidth vers. data size measured with microEnable

4.5 Pre-processing

Pre-processing is done in the FPGA's before passing the data to the PCI bus. This reduces the data volume and decreases the requirements for the following instances. For the TRT Rob a suggestion for a data format can be found in [5]. Also measurements have been done in [5] with a ROB-IN based on microEnable. An other suggestion for a data format can be found in [2] which reduces the data volume with factor 0.7.

A preprocessor algorithm for the SCT has been done in [6].

4.6 The private bus

The private bus is a connection between the ATLANTIS boards. It is not connected to the Host CPU. The bus has no specific protocol or protocol engine. It is a connection between two neighbored I/O Boards which gives the FPGA's the possibility to exchange signals.

The private bus is used for collecting data from neighbored I/O boards. This is useful if data from more than one input port is requested and the data can be found on two or more I/O boards. In that case one I/O board has the possibility to collect all data and pass one big data package over the PCI bus instead of several small packages. Also the ROB Controller passes only one message to one I/O board. Both increases the bandwidth of the PCI Bus and reduces the messages between ROB Controller and ROB-IN's.

5 The ROB-Controller

The host CPU of the ATLANTIS system acts as ROB-Controller. The host CPU answers requests from outside to the ROB and dispatches them to the different I/O boards. It analyses the requested ROI's and requests the matching event data from the different boards.

6 ROB to Switch Interface

The Rob to Switch Interface connects the ROB Complex with the LVL2 Trigger, the Eventfilter and the Run Control. The actual ATLANTIS Host CPU has a build in 100 Mbit ethernet adapter. This is sufficient for the first studies of the system. But with four I/O Boards the bandwidth of that adapter is too low. In future systems a Gigabit ethernet adapter or a 622Mbit ATM adapter is more reasonable.

References

- [1] O. Boyle, R. McLaren. ROB-IN Input Module For The Demonstrator ROB, ECP Division, CERN, 1996.
- [2] P. Clarke, F. J. Wickens. Detector and Read-Out Specification, and Buffer-ROI Relations, for Level-2 Studies. Atlas DAQ Note Final DRAFT 9/1999.
- [3] Jos Vermeulen. Modelling slides, <http://www.nikhef.nl/pub/experiments/atlas/daq/Modelling-2-6-99/Presentation-June-99-update.pdf>.
- [4] Atlas Technical Proposal Chapter 5. CERN, 1994.
- [5] Ruediger Rissmann. Implementierungen von Vorverarbeitungsalgorithmen für den ATLAS Level 2 Trigger auf dem FPGA-Prozessor microEnable. University of Heidelberg, 1999.
- [6] R. Dankers, J. Bains. A Data Preparation Algorithm for the Precision Tracker LVL2 FEX. ATL-DAQ-99-001, CERN, 1999.