A PROTOTYPE ATM NETWORK FOR REAL TIME CONTROL OF THE LHC

T.Wijnands, W. Herr, P. Ribeiro, CERN, Geneva, Switzerland

Abstract

The LHC accelerator aims at injecting, accelerating and colliding beams with very well controlled beam parameters (e.g. momentum, orbit, tune and chromaticity). This is a non-trivial task since the super conducting main bending magnets will generate field errors with dynamic effects that may result in beam loss. To overcome this problem, real time control of beam parameters via the Power Converters has been proposed. This requires site wide deterministic communication of control data. In this paper we will outline some aspects of a prototype deterministic network for the LHC with a core based on ATM (Asynchronous Transfer Mode) communication technology.

1 INTRODUCTION

The LHC (Large Hadron Collider) [1] is a proton-proton collider that will be commissioned around 2005 at CERN. The LHC differs in many aspects from the presently operating accelerators at CERN since it has a smaller dynamic aperture and higher stored beam energy while being much more sensitive to beam loss. These characteristics have lead to imposing more stringent requirements on the control of beam parameters like momentum, orbit, tune and chromaticity.

During any operation cycle, beam parameters will vary due to field errors, generated by the super conducting magnets. Of particular interest to us here are field errors caused by dynamic effects that result in a decay of the integral magnetic field although the current in the magnet remains constant. This results in field errors that cannot be eliminated entirely with feed forward correction.

Site-wide real time control of critical beam parameters using beam position monitors as sensors and power converters as actuators is proposed to solve this problem. This requires communication of time critical data over a large distance (typically 5 to 7 kilometers).

In this paper, we will outline a proposal for a 2-layered real time communications infrastructure for the LHC with an upper layer that is based on ATM (Asynchronous Transfer Mode) technology. We will demonstrate that this layer can support traffic from various types of real time applications over a large distance while giving guarantees on the communications latency and on the bandwidth available to the user applications. Moreover, we will show that matching the capabilities of network and end nodes to the traffic pattern of the application can lead to very efficient and reliable deterministic communications.

2 REAL TIME APPLICATIONS

2.1 Classification

The real time control tasks for the LHC can roughly be divided in three different categories : periodic, aperiodic and interactive tasks.

Most periodic tasks will be concerned with feedback control of the beam parameters as mentioned earlier. Periodic tasks generate CBR[1] traffic and require guarantees on communications delay and jitter since otherwise machine performance degradation or damage may occur.

Examples of aperiodic tasks are the completion of a state transition of a piece of equipment or the reaching of alarm conditions. Such events generate rt-VBR[1] traffic. Fixed transfer delay is also necessary here because appropriate action must be taken within a given time.

Finally, as an example of an interactive control task, one can think of an operator asking for a status report from a piece of equipment. Interactive based events have UBR[1] traffic since transmission is initiated on a specific command or request. In this case, optimizing the use bandwidth is usually as important as minimizing the jitter (the spread in latency).

2.2 An example : Closed Orbit Control

Controlling the closed orbit in the LHC is probably the most demanding application in terms of communications. When the closed orbit of a single beam in the LHC needs correction, beam position data is collected from 512 beam position monitors (BPM) in some 248 local data acquisition crates distributed equally along the ring. There will approximately 80 bytes of data per BPM (40 kBytes total). The position data is then send to a central computer where the response matrix (maximum size 512x512) is computed using a matrix manipulation program like MICADO. Finally, the central node distributes the correction data among the 512 power converters that drive the current in the orbit corrector magnets.

Preliminary investigations suggest that it would be desirable to correct the closed orbit of the LHC beams in a closed loop with a bandwidth of the order of 0.1 Hz. It has been estimated that for this case, about 30 % or more of the overall loop-induced phase shift will be due to data communication.

[1] CBR = constant bit rate, rt-VBR = real time variable bit rate, UBR = unspecified bit rate.

3 ATM INFRASTRUCTURE

3.1 ATM : Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) [2] is a connection oriented, multiplexing, communication technique which was designed for carrying various kinds of traffic - data, voice and video - over the future B-ISDN (Broadband Integrated Services Digital Network). The main advantages of ATM over other techniques originate from the use of smaller fixed-length packets (called 'cells') of 53 bytes instead of packets which allows for fast digital switching and transmission and which make quality of service (QoS) for end users possible. QoS defines the required bandwidth and latency in terms of parameters such as Cell Rate (peak, minimum, average) and the Cell Delay Variation Tolerance.

At the connection set up, a traffic contract is negotiated between the user and the network. This contract defines the expected traffic pattern in terms of a traffic class and the required Quality of Service (QoS). A call admission strategy judges whether the network capacity is sufficient to accept the new connection. Once a connection has been created, data can be transported with fixed communications latency while only a predetermined fraction of the total available bandwidth is being used. The cell rate is determined by the number of entries allocated in the transmit scheduling table.

These characteristics motivated investigating the use of the ATM networking technology for our real time control system.

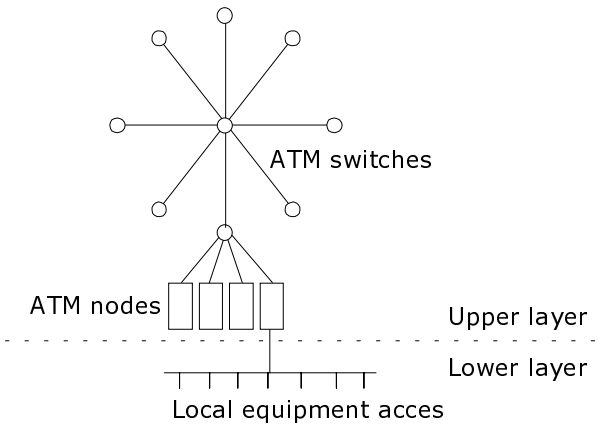

Figure 1: The 2-layered real time communications topology proposed for the LHC using ATM technology.

3.2 ATM network

An example of a 2-layered real time communications network topology for the LHC is shown in figure 1. The upper layer consists of 8 ATM switches equally spaced along the accelerator ring and a single switch in the center. The switches in the vicinity of the beam pipe accommodate several 155 Mbit/s links to the ATM nodes and a 622 Mbit/s up-link to the central switch in the LHC control room. Alternative paths to every switch (not shown here) will provide for network redundancy.

Local equipment is accessed via the second layer that will probably use a field bus technology like WorldFip [3]. WorldFip field busses have the advantage of providing a relative simple, flexible and cost-effective solution while covering LHC specific requirements concerning distance and radiation resistant copper cabling [4]. It is also possible to map a real time communication channel onto WorldFip since the bus uses a static, table driven scheduling scheme (figure 2).

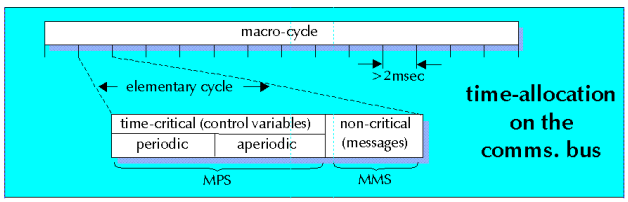

Figure 2: WorldFip Macrocycle with several elementary cycles allowing for communication of time-critical and non time-critical control data.

3.3 ATM switches & nodes

The ATM switches that have been used here have a 5 Gbit/s shared-memory switching capacity with non-blocking switching fabrics operated by a 100 MHz MIPS R4600 switch processor. All can accommodate several 155 Mbits/s links and a maximum of two 622 Mbit/s up-links.

The ATM nodes are standard VME based RIO2 8062 Power PCs at 300 MHz equipped with commercial ATM adaptor cards [5] that are plugged into to the PMC slots and accessed via the PCI bus. The PowerPCs are operated under LynxOS, a UNIX-compatible, POSIX-conformant multi-process and multi-threaded real-time operating system.

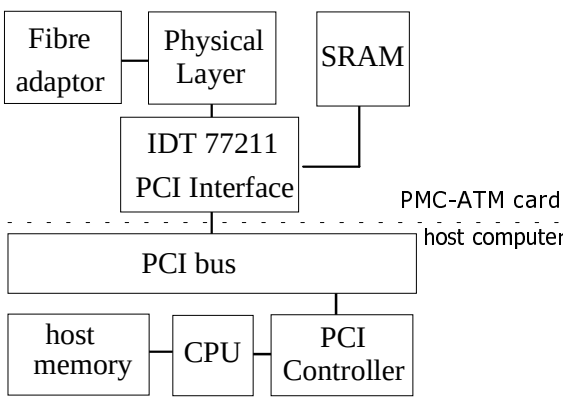

Figure 3: Layout of the PMC based, ATM Adaptor Card for the ATM nodes.

The ATM adaptor (figure 3) cards are based on the IDT NICStAR 77211 ATM chip set and the standard

software driver supports AAL-0 (raw cells) and AAL-5 formatted traffic at CBR[1] or UBR[1]. This ATM chip has a highly integrated design and performs both the AAL segmentation/reassembling and the ATM layer protocol functions. The host CPU is only used for small data movements; the remainder is using direct memory access.

4 EXPERIMENTAL RESULTS

As indicators for the performance of the ATM network and end nodes we have considered latency and throughput. Latency is always present due to soft and hardware overheads at the host and due to a limited line speed. Figure 4 shows the latency for 2 different sized messages using native AAL-5 traffic over a 155 Mbit/s ATM link in a configuration without a switch.

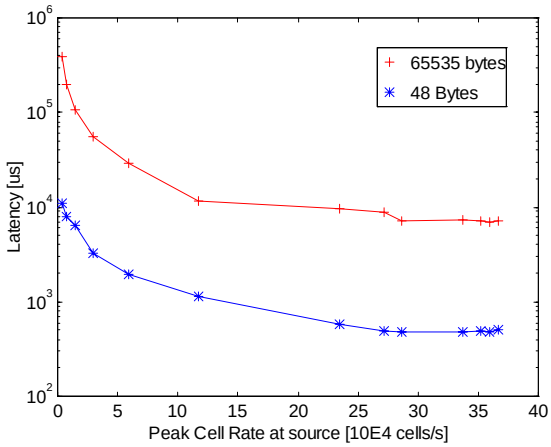

Figure 4: Communication latency as a function of the Peak Cell Rate at the source (AAL-5, 155 Mbit/s link)

While the peak cell rate is set to a small value, the hosts experience difficulties in sequencing the cell emission efficiently. In addition, it will take a significant amount of time before large messages (broken up in many cells) have been put on the link by the host.

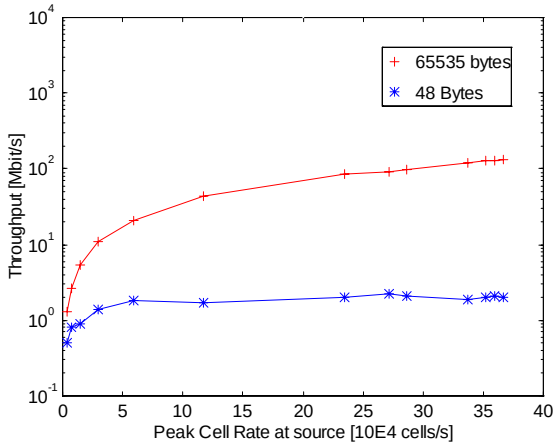

Figure 5: Throughput as a function of the Peak Cell Rate at the source (AAL-5, 155 Mbit/s link)

At cell rates of 60 % or more of the maximum, the soft and hardware overheads and the line speed increasingly dominate the latency. In this case, increasing the cell rate does not decrease the latency any further.

Figure 5 shows the throughput for traffic with identical sized messages. For large messages there is a linear relationship between cell rate and throughput as one would expect. For small messages however, it is not possible to get an improved throughput by increasing the cell rate. When small messages of few cells are used, the throughput is no longer controlled by the cell rate. The time it takes to put a cell on the line is now equal to the soft and hardware overhead at the host.

It should be noted however, that the host can prepare data for transmission at rates of the order of 500 Mbit/s which is far superior to SDH-STM1 line speed. This creates a potential risk of data loss.

Similar experiments have been carried out in configurations with one or two ATM switches. Crossing a single ATM switch increases the latency by about 20 μs and this is independent of the message size. However, the switch delay plays only a role of importance when small messages are send.

5 CONCLUSIONS

Efficient and reliable operation of the LHC will probably require a site-wide real time control system to correct field errors with dynamic effects during operation. This has motivated the construction and investigation of a prototype real time network with an ATM core.

It has been shown here that ATM technology is appropriate for providing deterministic communication links for the real time applications such as those expected for the LHC accelerator.

Very efficient and reliable connections can be created across the network if the application traffic profiles are known beforehand. The communications requirements should then be formulated in terms of traffic contracts that are compatible with the network resources and the performance of the ATM hosts.

REFERENCES

[1] *LHC Study Group*, "The Large Hadron Collider Conceptual Design", CERN report, LHC Disvision, CERN/AC/95-05, October 1995.

[2] D. Ginsburg, "ATM - Solutions for enterprise networking", Addison-Wesly, Addison Wesly Longman, Harlow, England, 1996.

[3] http://www.worldfip.org.

[4] R. Rausch *et al*, "On-line radiation test facility for industrial equipment needed for the Large Hadron Collider at CERN", this conference.

[5] http://www.ces.ch