

Real Time Communications Prototyping for LHC Controls

Author(s) / Div-Group: Thijs Wijnands SL/CO/FE

Pedro Ribeiro SL/CO/FE

Keywords: Networking, bound delays, Asynchronous Transfer Mode, WorldFip

Summary

A state of the art accelerator control network is required to meet the physical requirements of the Large Hadron Collider (LHC). The evolution of the field of the superconducting magnets cannot be predicted with sufficient accuracy for beam control. Optimum efficiency of the machine operation will therefore have to rely more on real time data channels and feedback control. The controls network must be fast and scalable and provide delivery of data with a fixed latency. One of the solutions presently under investigation is a 2 layer network consisting of a high speed backbone and a field bus (similar to the present controls network for the LEP accelerator). ATM is amongst the candidate backbone technologies while the WorldFip fieldbus is presently envisaged to provide local access. In this report benchmarking measurements on ATM are presented, followed by a brief discussion on the integration of ATM and the WorldFip fieldbus.

1 Introduction

In accelerators like LEP and the future LHC, diagnostic measurements provide experimental data that is sent to a central computing system. The central computing system executes algorithms and with the results, new references for the actuators are generated (i.e. the magnets and the RF power). The LHC will use super conducting magnets and typically 1 hour will be needed for the injection and ramping phase. The evolution of the magnetic field and beam parameters cannot entirely be foreseen in advance and more real time control will be used. In order to handle real time traffic, the controls network must fulfil the following requirements:

- Bound latency delivery of data
- Dynamic allocation of the available bandwidth
- Negotiable Quality of Service of the connection
- Scalability with minimum reduction of performance

A single layered network would be too expensive to implement and is not likely to suit the complexity of the problem. Present studies therefore concentrate on a 2-layered network using a fast backbone and a field bus to access distributed equipment. Candidate technologies for the backbone are FDDI, Fast Ethernet (switched or shared), Fast Token Ring (for existing Token Ring users) and ATM. Fieldbus technologies such as Profibus, CANbus and WorldFip are presently candidates for providing local access.

In this paper we will focus on two particular techniques, ATM and WorldFip technology. After a general introduction to ATM in chapter 2, the characteristics of permanent virtual ATM connections in a small-switched network are discussed. The overlay model, using the TCP/IP and the UDP/IP protocol stack on top of ATM, is discussed in chapter 4. A short introduction to the WorldFip fieldbus is given in chapter 5 followed by a summary of the present experience. Finally, future prospects are discussed in chapter 6 followed by a number of conclusions in chapter 7.

2 ATM technology

2.1 Introduction

ATM stands for Asynchronous Transfer Mode and is a proposed telecommunications standard for broad band ISDN. The word "asynchronous" actually means "asynchronous in time" and "asynchronous time division multiplexing" might be a more appropriate term to describe the ATM technology.

In an ATM network, a connection between two end nodes will only reserve a part of the available bandwidth when data has to be transmitted. ATM carries the data in short packets (ATM cells) of 53 bytes, each carrying identifiers of the connection they are associated with. The packets are sent as they are generated (hence the data flow is asynchronous in time) and no bandwidth is wasted for inactive connections. This is extremely advantageous for data that is generated and transmitted in short bursts. Another advantage is that the sum of peak transfers rates of all connections can easily exceed the maximum available bandwidth. However, data can be lost when all nodes burst at the same time at maximum speed.

In any ATM network, data is sent between the end nodes over Virtual Paths (VP) and Virtual Channels (VC). A single ATM interface typically supports multiple Virtual Paths, each containing several Virtual Channels. In fact, a Virtual Path simply routes the virtual channels through the network. Both VP and VC are labelled through Virtual Path Identifiers (VPI) and Virtual Channel Identifiers (VCI).

The User Network Interface (UNI) defines the interface between an end node and a switch, while the Private Network-Network Interface (PNNI, see also chapter 6) defines the interface between switches. The ATM switch has a routing table for each physical link and both the virtual path and the virtual channel of the connection can be changed. The routing table can be static in which case Permanent Virtual Connections are used or dynamic using Switched Virtual Connections.

Since ATM is designed for switching short fixed length packets in hardware, the ATM protocol stack is placed in the data link layer. The most currently used protocol is the AAL-5 layer (ATM Adaptation Layer) which is optimised for data transport.

ATM has no guaranteed delivery of data. If such a guarantee is required, transport protocols like TCP can be stacked on top of the AAL-5 layer but it should be mentioned that optical point to point ATM links can be considered as error free.

2.2 Advent of ATM

ATM is presently among the few fast networking technologies that can transport data from end node to end node within a bound latency while respecting a predefined traffic contract. For each connection, the end user can define a traffic contract. The (complex) Connection Admission Control (CAC) Algorithm will determine whether the ATM network can accept the connection

under the specified conditions or not. If 2 or more switches are present in the network, a separate traffic contract should be defined for the traffic between switches.

- The key parameters of any traffic contract are:
- Traffic Class (Constant Bit Rate, Available Bit Rate, ...)
- Traffic parameters (Peak Cell Rate, Minimum Cell Rate, ...)
- Quality of Service (Cell transfer Latency, peak to peak cell Variation, ...)

For an ATM network with 2 or more switches, parameters like the Maximum Cell transfer latency are accumulated over all the hops of a call. Parameters such as the Cell Loss Ratio are attributes and checked at each interface. The Traffic parameters and the Traffic Class are defined at the end nodes (see next section) while the QoS table is defined at the switch. It may be clear that the traffic contract has many degrees of freedom. In practice, modified versions of the default traffic contracts are used which usually are the least constrictive.

ATM has found its proper place in the present networking market. Presently, about 10% of LAN and WAN networks are using ATM mainly because of its real time functionality. Bandwidth only plays a secondary role. The market share of ATM will not change in the near future, although prices are expected to decrease.

2.3 Some ATM networks presently in use

A recent application of ATM in HEP at CERN is the ATM-based Event Builder Demonstrator. Event building is the process of collecting all relevant data related to one event. The major complication in event building is the high data rate (up to 32 Tbyte/s expected for the LHCb experiment) but ATM switching fabrics are considered good candidates for implementing high performance parallel event builders for the future data acquisition systems. In [4], the feasibility of an ATM event builder is tested through simulations and implementation. It was found that small fragments (256 bytes) can be received at 38 kHz or higher (depending on the number of destinations) while for larger fragments (up to 1 kByte) the link throughput is the limiting factor. Similar experiments have been carried at other laboratories like Fermilab and at Columbia University [5,6].

The European Synchrotron Radiation Facility in Grenoble [7] has been using an ATM network since 1993 for high speed data transfer between detectors on the beamlines and a centralized computing facility. ATM technology was selected because of its scalability : 100 Mbps interfaces were already available in 1993 while 155 Mbps and 622 Mbps interfaces were announced. Moreover, interest in ATM was growing while FDDI was losing interest and became more expensive.

ATM also finds support in Controlled Nuclear Fusion Research. The network of the JET control system for example, is presently being migrated to ATM [8], while the Large Helical Device is using an ATM office network to provide physicists with experimental data at their desks [9], but not (yet) for the control system of the machine itself.

In [10], a Transatlantic ATM Pilot Project Applications for Worldwide Scientific Collaboration is proposed linking CERN, the ITER site in Garching and DESY in Germany, the INFN sites in Italy and IN2P3 in France.

The ATM network of the Bibliotheque Nationale de France (BNF) in Paris also merits mentioning. Here a "pure" ATM solution was chosen to deliver selected books to the readers in the 1800 reader halls. Some 2800 users have fibre to the desktop providing a measured bandwidth of 80 Mbit/s using IP over ATM (see chapter 4).

3 Permanent Virtual Connections

Summary : Native ATM traffic on a Permanent Virtual Channel (PVC) is investigated using two different configurations. In one set up, the data is passed directly from one end node to another through a multi mode fibre; in another one data is passed through an ATM Switch (a 5 Gbit/s CISCO Lightstream 1010). Native ATM traffic on a PVC has an average throughput of 130 Mbit/s with a minimum latency of 200 µs per cell.

3.1 ATM network

Figure 3.1 shows the two different configurations that are investigated here : the first configuration has a "back-to-back" connection between 2 end nodes (ATM 8468 PCI Mezzanine cards on RIO 8062 Power PC boards housed in VME crates, see figure 3.1a), the other configuration is a very simple ATM network consisting of 1 switch (a 5 GBit/s CISCO Lightstream 1010 ATM switch) and 2 end nodes (figure 3.1b).

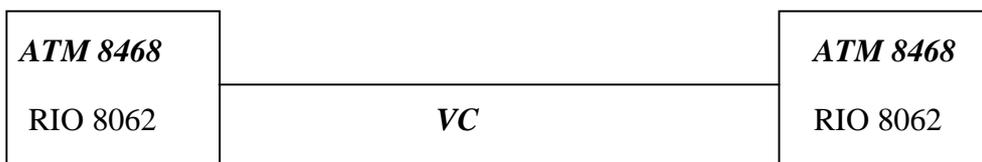


Figure 3.1a: Basic ATM Connection with a Virtual Channel (SDH/SONET multi mode fibre)

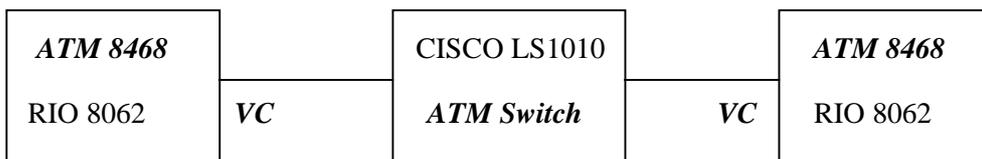


Figure 3.1b: Simple ATM network with 2 Virtual Channel Connections and one switch

Parameter	Type or value	Nature	Full Name
Protocol	AAL-5	fixed	ATM Adaptation Layer

Class	ABR/CBR	variable	Traffic Class
PCR	< 130 Mbit/s	variable	Peak Cell Rate
SDU	< 64 kByte	variable	Service Data Unit size
Timeout	60 ns	fixed	Timeout for send/receive

Table 1 : Supported traffic parameters for a direct connection between 2 ATM 8468 PCI mezzanine cards

Communication between end nodes takes place over Permanent Virtual Connections (PVCs). Any switch in the network needs to be programmed beforehand, defining the QoS parameters (maximum cell transfer latency, peak-to-peak cell delay and the maximum cell loss ratio). Here, the default QoS table of the switch was used (which is in fact the least restrictive QoS contract possible). Before exchanging data, sender and receiver must decide on the traffic class (UBR, ABR, CBR, nrt-VBR or rt-VBR) and on traffic parameters (Peak Cell Rate, Sustainable Cell Rate and their tolerances, see table I).

3.2 Native AAL-5

The sending (receiving) time is defined as the time required for the sending (receiving) the data from user space. In what follow GPS timing modules were used to time emission, reception and latency. Emission (reception) is defined as sending/receiving data from user space. The latency of a packet is defined as the time this packet is under way from user space to user space on different end nodes.

Available Bit Rate

Figure 3.2 shows the sending and receiving of 1000 data packets of varying size over a PVC using ABR/AAL5 and the configuration shown in figure 3.1a ("back-to-back" or Tx-Rx). In an initial phase, data is emitted over the PCI bus faster then it is emitted over the fibre and so the Tx-FIFO queue of the sending end node is filled up. When the Tx-queue is full, data is emitted at a constant rate of 130 Mbit/s. Sending smaller data packets takes less time and so the Tx-FIFO queue is filled up quicker and less time is required to get to a constant throughput.

It should be noted that the maximum throughput of 130 Mbit/s is only achieved when the maximum ATM Protocol Data Unit (PDU) size of 64 kByte is used. In this case, the overhead introduced by the hardware represents only a very small fraction of the time required to send 64 kBytes over a 155.52 Mbit/s link.

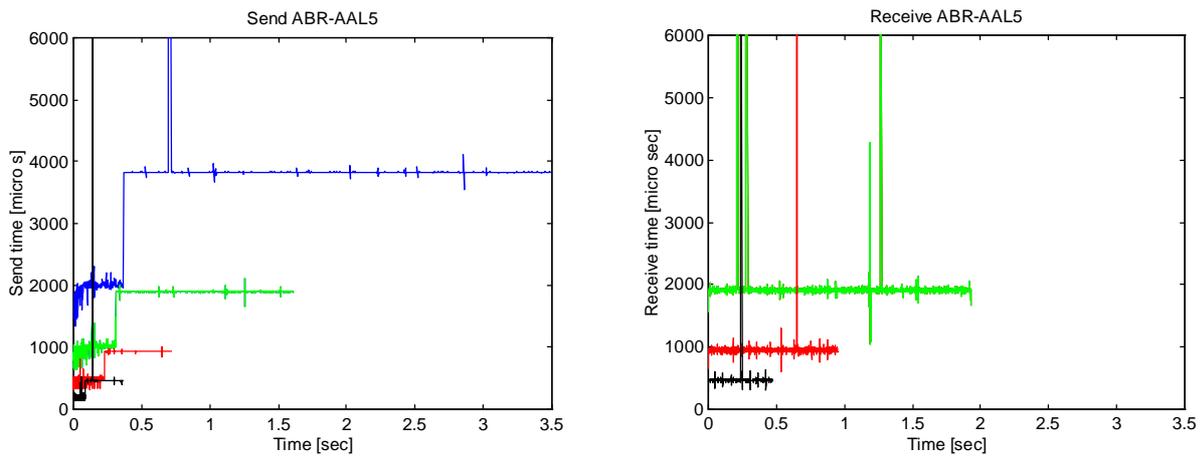


Figure 3.2 : Time for sending (left) and receiving (right) 1000 data packets of varying size (64,32,16 and 8 kByte) over a direct ATM-link using AAL5-ABR.

Occasional large variations (2-5 times the average value) of the throughput in figure 3.2 are due to the interrupts not related to the ATM device. It is well known [11] that LynxOS handles interrupts in a particular manner providing a single entry point for all possible candidates while no priority levels can be set. Interrupts thus pile up until being processed in an occasional burst, leading to a sharp decrease of the throughput. Further detailed studies are required to determine the cause of the interrupt pile up.

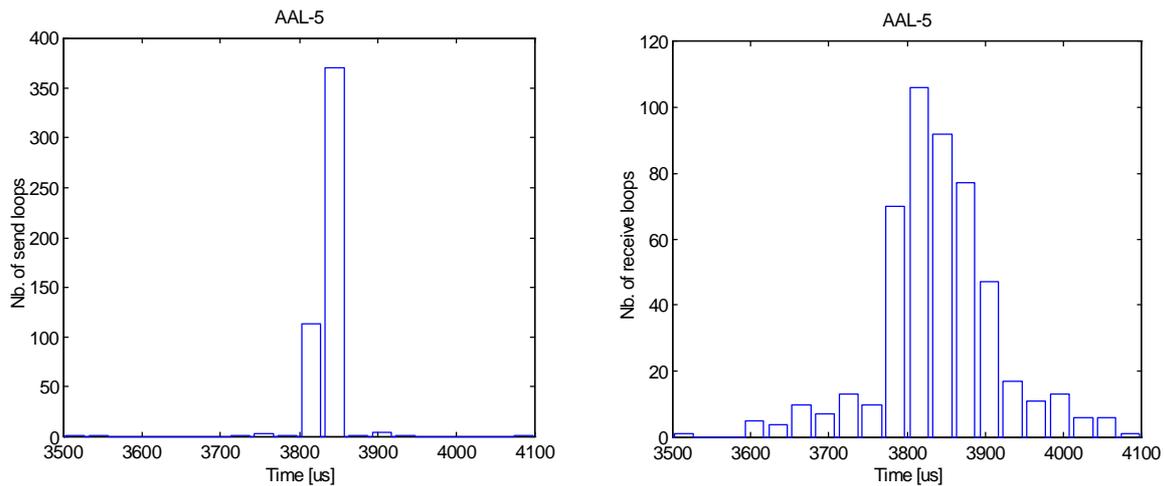


Figure 3.3: Histograms of the throughput for sending (left) and receiving (right) over a direct ATM-link using AAL5-ABR (64 kByte packets, bin size 30 μ s)

Figure 3.3 shows histograms of the data presented in figure 3.2 for the case of sending/receiving 64 kByte (occasional bursts in the ATM throughput omitted). Both histograms have a bin size of 30 μ s and contain 500 data points.

To send an ATM cell, the cell payload of 48 bytes is transferred from host memory to the Tx-FIFO queue via a DMA transfer and then emitted on the fibre via the UTOPIA bus interface. The standard deviation in sending is thus low. Receiving an ATM cell is more complex : when the cells arrive in the Rx-FIFO queue (315 cells), the NicStar performs DMA transfers and fills up host memory space. Corresponding addresses are written to the receive-status-queue which is served when the end of a PDU is detected and an interrupt is generated. The host has to do some processing when giving the data to the user that creates a spread.

Packet size [bytes]	Mean [us]		Standard deviation [us]		Theory [us]	
	send	receive	send	receive	time	overhead
65535	3833	3841	21	84	3722	118
32768	1900	1908	19	36	1861	47
16384	936	944	7	37	931	13
8192	461	462	5	16	461	1

Table II : Sending/receiving data buffers of varying size over a PVC using AAL5-ABR.

Table II summarises the data shown in the previous figures supposing no additional interrupts are generated, or that some priority setting of interrupts is being used. The theoretical time indicated in the last column, was calculated assuming a 10 % overhead for the ATM header.

Constant Bit Rate

Figure 3.4 shows a similar experiment, emitting 1000 data packets of varying size over a PVC, only now using Constant Bit Rate (CBR). The Tx-FIFO queue for ABR traffic is 8 times bigger than for so that a constant data flow is rapidly obtained (figure 3.4 left, note the logarithmic scale). Occasional large variations are again due to interrupt bursts.

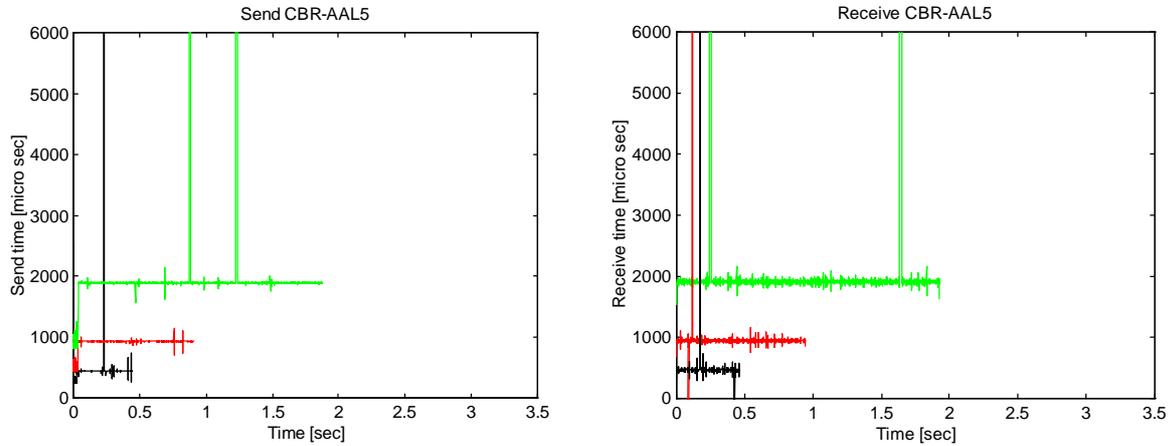


Figure 3.4 : Time for sending (left) and receiving (right) 1000 data packets of varying size (64,32,16 and 8 kByte) over a direct ATM-link using AAL5-CBR.

From the summary in table III it can be concluded that sending and receiving CBR is very similar to sending ABR. For 64 kByte PDUs, the bandwidth is close to the theoretical bandwidth of 155.52 Mbit/s if the header overhead is subtracted. When interrupt pile up is avoided, AAL5-CBR traffic is well suited for real-time feedback control applications that require a bound latency (see also next section).

Packet size [bytes]	Mean [us]		Standard deviation [us]		Theory [us]	
	send	receive	send	receive	time	overhead
65535	3824	3840	38	77	3722	118
32768	1888	1972	31	37	1861	111
16384	926	954	3	21	931	23
8192	463	465	2	30	465	0

Table III : Sending/receiving data buffers of varying size over a PVC using AAL5-CBR.

Latency

Figure 3.5 shows measurements of the latency in the previous experiments. For both CBR and ABR, the latency is increasing as a function of time when Tx and Rx Buffers fill up. Since there is 8 times more queuing memory available for ABR than for CBR, the initial phase

takes longer. Moreover, in a steady state situation, the latency for ABR is on average 8 times bigger. Table VI(a) summarises latency measurements when sending large PDUs.

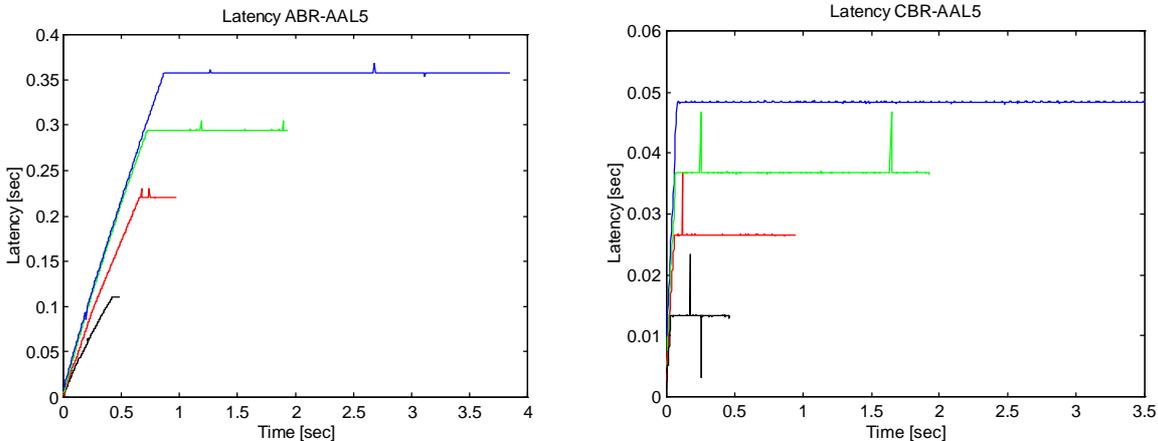


Figure 3.5 : Latency for when sending 1000 data packets of varying size (64,32,16 and 8 kByte) over a direct ATM-link using AAL5-ABR(left) and AAL5-CBR(right).

Packet size [bytes]	Latency ABR [ms]	Latency CBR [ms]
65535	358	48
32768	294	37
16384	220	26
8192	110	13
4096	55	6

Table VI : Latency when transmitting large data buffers of varying size over a PVC using AAL5.

The ATM send-receive mechanism can be studied in detail when small (48 bytes) PDUs are used. For single cells, the latency is more or less constant since the available memory space is never entirely filled (figure 3.6).

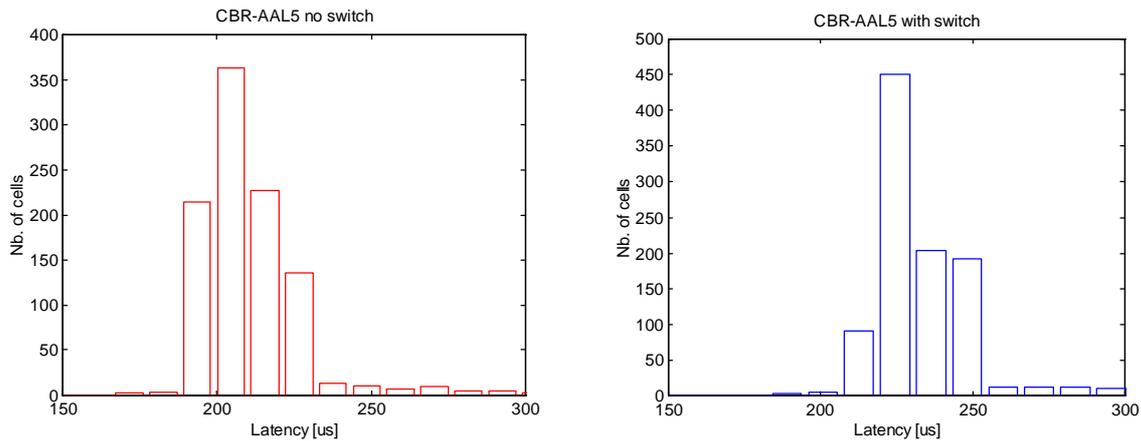


Figure 3.6 : Histograms showing latency for single ATM cells (48 bytes) using CBR-AAL5 (left) measurements in the configuration shown in figure 3.1a (right) measurements of the latency using the configuration with the ATM switch shown in figure 3.1b.

The latency for a single ATM cell is about 200 μ s, very large compared to the cell line speed (2.7 μ s for 48 bytes over a 155 Mbit/s link). The sending and receiving mechanisms have not yet been studied in detail, but preparing the a single cell in the Tx-FIFO queue takes 30 μ s while 10 μ s are required to empty the Rx-FIFO queue. In addition, the UTOPIA bus interface has been recognised for causing delays and IDT has recently introduced a new interface bus [12,13] that reduces the amount of data on the bus.

Figure 3.6 also shows the single cell latency in the configuration with the switch (see figure 3.1b). The single cell latency is increased by 20 μ s.

In summary, Reliable transfer of data at CBR or ABR over Permanent Virtual Connections has been achieved. Using native AAL5, the maximum throughput of 130 Mbit/s is obtained for PDUs of 64 kByte. In this case, the latency is dominated by the process of copying data from and to host memory space. For small sized PDUs or for single ATM cells (48 bytes), the latency is about 200 μ s that is large compared to the single cell transfer time (2.7 μ s). The ATM switch that was used here added another 20 μ s to the total delay. Detailed studies of the hardware and the low-level software are needed to make this issue more precise. The possible optimisation of the ATM LynxOS drivers is presently investigated.

4 Internet Protocol over ATM

SUMMARY: Measurements of throughput, latency and reliability of TCP/IP and UDP/IP traffic over AAL5 are presented. UDP/IP over ATM has low latency and a throughput of 80 Mbit/s has been achieved. TCP/IP over ATM has a throughput of 55-60 Mbit/s but is less suitable for real time applications. Both protocol stacks require careful tuning of buffer space, transfer windows and data flow.

4.1 Motivation

The large majority of the present day LANs and WANs are based on Ethernet and use IP based protocols. An isolated ATM network without an Ethernet connection is therefore not acceptable. Any ATM network should be able to accommodate IP traffic and IP routing and forwarding in particular. There are two fundamentally different modes of operation in which IP over ATM is implemented [14]:

In native operation, IP address resolution is used to map the network layer address directly into an ATM address. In this peer model, the ATM switch becomes an IP router and ATM addresses are obtained via IP address resolution mechanisms. In the peer model, the protocol stack has an AAL-5 header on top of a standard UDP/IP header.

In LAN Emulation (LANE), a local area network is emulated on top of the ATM network. The ATM network is transparent to the end user that will only notice an increase of the available bandwidth. The protocol stack larger then in the peer model mentioned above since it has an AAL-5 header on top of a UDP/IP/Ethernet header. In this overlay model, there is no relation between ATM and IP addresses or between ATM routes and IP routes.

Routing is of no particular interest to the benchmarking tests that are presented here. In what follows, we therefore make use of a simplified peer model in which the TCP and the UDP/IP stack are statically connected to the AAL-5 layer. All TCP/IP and UDP/IP traffic is directed to a single virtual path without multiplexing PDUs.

4.2 Software Implementation

The encapsulation of the IP traffic over ATM AAL-5 makes use of a technique called "Virtual Channel (VC) multiplexing"[15], which is in fact the most straightforward approach to implement IP over ATM. With VC multiplexing any set of protocols can be multiplexed over a single Virtual Channel. An 8 byte header is added to the data packet before being encapsulated in an AAL-5 Service Data Unit. The IP layer is thus considered as the start of the Virtual Channel where IP packets are placed into AAL-5 Service Data Units.

The set up of the IP over ATM application is purely connection oriented. First the output of the IP stack is directed to the ATM interface by creating an atm device [16]. This defines the underlying AAL-5 traffic contract of the connection. The ATM interface is then configured like any other networking interface. As many links as desired can be now be created, each with their own specific traffic contract. For the user, the ATM layer is completely transparent. If data has to be sent from end node, a socket is opened and a VC set up. When all data has been sent from the end node, the VC is torn down. Increase of performance can be achieved by suppressing check summing and reducing the amount of copying within the host memory space.

VC multiplexing results in minimal bandwidth and processing overheads, but has a sincere drawback. On a particular VC, one will never find ATM cells of different IP packets

interleaved. The packet structure is thus maintained, even though the payload is distributed over several 48 byte ATM cells. Contrary to the fundamental idea of cell switching, ATM cells on a particular channel have a predefined order. In other words, sending packets over a network based on cell switching will inevitably lead to loss of performance.

4.3 TCP/IP and UDP/IP over ATM

UDP provides a connectionless “best effort” delivery of data without guaranteed arrival of data sequencing or flow control. Since the optical ATM link can be regarded as error free, loss of datagrams is often the result of overflowing buffer space at the end nodes. For UDP traffic, it is therefore important to tune the available memory space at the end nodes in such a way that a steady flow of data can be handled without excessive queuing. The maximum buffer size for sending and receiving UDP buffers under LynxOS is limited to 43690 bytes while the maximum PDU size is 9216 bytes. If data is sent faster than it is received these buffers gradually fill up. Eventually, there will be data loss when the buffers are full.

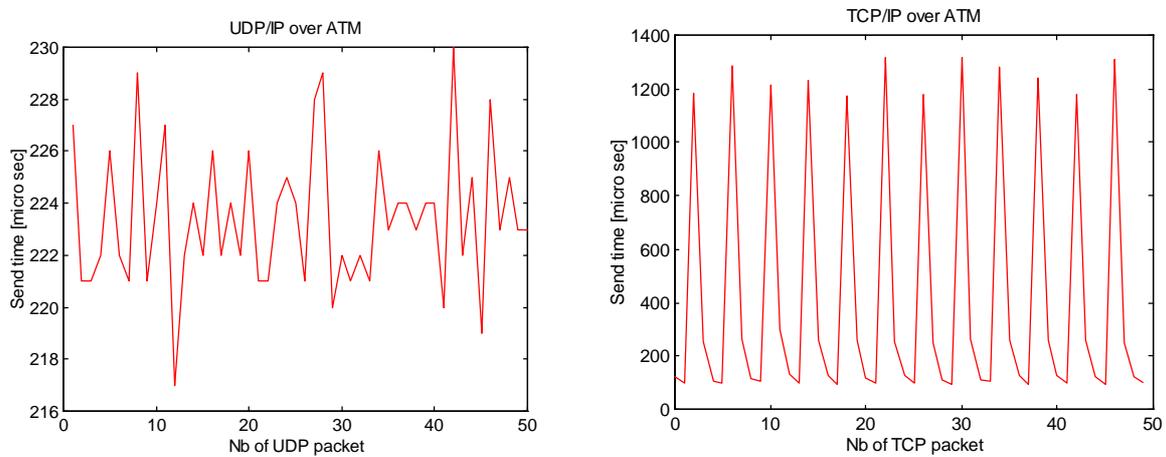


Figure 4.1: Traffic sample while sending 2048 PDUs of 4096 byte over a direct ATM-link using UDP/IP over ATM (left) and TCP/IP over ATM (right).

A throughput of 120 Mbit/s is obtained when sending 2048 PDUs of 4 kByte (check summed). Apart from the occasional overhead, the data flow is continuous with a small spread. At the receiving end node, the throughput is only of the order of 31 Mbit/s (see table V).

Protocol	Send [Mbit/s]	Rec. [Mbit/s]	Loss
UDP	122	31	23 %
TCP	54	54	0 %

Table V : Performance of TCP/IP and UDP/IP over ATM using checksumming.

Figure 4.1 also shows a sample of the TCP/IP traffic over ATM when sending 2048 PDUs of 4096 bytes each. Again tuning is required in order to get the maximum performance. A constant flow of data of about 50 Mbit/s is obtained (with checksumming).

These benchmarking measurements were repeated using the ATM switch (see figure 3.1b). In none of the cases, a significant change of the performance could be observed.

In summary, VC multiplexing is an efficient technique to emit UDP/IP and TCP/IP based traffic over ATM. The price to pay for using these protocol stacks is a reduction of the data throughput compared to native ATM. Similar to Ethernet, datagrams are occasionally lost and in the case of UDP/IP. Careful tuning of sender and receiver buffer sizes is needed in order to provide a constant data throughput. Alternatively, the number of PDUs and their size can be adapted to the available memory space, as was done in the experiments present here.

The TCP/IP/AAL5 protocol stack is unsuitable for real time data traffic since data delivery within a certain fixed latency cannot be guaranteed. The UDP/IP/AAL5 protocol stack might be used if the occasional loss of data can somehow be accounted for.

5 WorldFip fieldbus

5.1 Technology

The classical Field Bus advantages of simplified cabling, ease of deployment and operation [17] are extended by powerful mechanisms integrating the basic bus mode of operation with the data exchanges. With WorldFip [18] it is possible to determine at a specific moment in time the data that is being transmitted over the bus. The bus becomes then a predictable element in a communication chain where the exchanges must occur within specific timeslots. In addition to these important characteristics there is a full support for redundancy both at physical layer and also at the data link and application layers.

For the LHC, a 2-layered network based on ATM and WorldFip technology can be envisaged although for some special applications, the use of only one of the two technologies might be sufficient. The combination of the two technologies by means of Real Time Operating Systems based gateways presents a powerful solution to a range of distributed applications requiring controlled communications latency.

5.2 WorldFip Protocol

This multidrop bus offers a one-producer multiple-consumers transaction model. Every piece of data exchanged over the bus has a unique identifier known to all the stations. A Bus Arbiter (BA) entity will coordinate all the exchanges by issuing a precise sequence of requests for data associated with the known identifiers. This sequence is known as the periodic variable exchange.

The BA will end its cycle by giving the possibility to the stations to transmit non-periodic information. This information can be in the form of non-periodic variables or messages. A user application can be responsible for refreshing a variable or it can be interested in consuming one of the distributed ones. The Data-Link Layer will notify each registered application every time a specific variable is updated or requires to be refreshed. Using this mechanism (plus a set of flags) will enable the system to do the bookkeeping of two important data characteristics, the freshness and the promptness.

6 Future Developments

Several issues will have to be addressed in the future in order to judge on the suitability of ATM and WorldFip technology for real-time data communication in the LHC.

First, the experiments presented here should be repeated in an experimental environment. This will allow studying the real time behaviour of the system and the ATM interface when other operational hardware (like the TG8 timing module, private Ethernet cards, a DSP or supplementary power PC) is present. Additional measurements should be carried out when the CPU is significantly loaded with (for example) floating point operations.

Second, the performance measurements should be carried out with up to date hardware and software. Both the ATM chip (the IDT77201) and the LynxOS (version 2.4) have become obsolete now. This comment does not concern the CISCO ATM switch since the new ATM hardware has not yet been implemented in the CISCO switching fabrics.

Finally, dynamic allocation of bandwidth via signalling should be studied using individual setting of the QoS parameters for each connection. Dynamic bandwidth allocation is called signalling and the type of connection that is set up is called a Switched Virtual Circuit (SVC) instead of a PVC. The signalling mechanism for ATM is using the SAAL (Signalling ATM Adaption Layer): a protocol stack that runs on top of the physical ATM layer. Signalling stacks are commercially available. Signalling is required to make an ATM network flexible and scalable and configuration management for a non-trivial network is facilitated.

7 Conclusions

Performance measurements of native AAL5 traffic over Permanent Virtual Connections have been presented. Traffic was generated on RIO Power PCs running LynxOS 2.4 and emitted via the PCI bus to the ATM mezzanine based on the IDT77201 ATM chip. Both CBR and ABR (or "best effort") traffic over permanent connection attains a throughput of 133 Mbit/s for 64 kByte Protocol Data Units (PDUs). For smaller sized PDUs, the throughput is reduced due to low level software and hardware overheads. For a single ATM cell of 12 words (48 bytes), the latency is about 200 μ s and 20 μ s are added every time the data cross a CISCO LS1010 ATM switch. Further detailed studies are required to determine the individual contributions and whether latency reduction is possible. The static ATM configurations that were used are convenient only for small-scale applications. Large-scale application of ATM technology requires a dynamic set up using Switched Virtual Connections.

VC multiplexing is an efficient technique to emit UDP/IP and TCP/IP based traffic over ATM although there is a reduction of the data throughput. Datagrams are occasionally badly sequenced in the case of UDP/IP. Careful tuning of sender and receiver buffer sizes is needed in order to provide a constant data throughput. Alternatively, the number of PDUs and their size can be adapted to the available memory space, as was done in the experiments present here.

The WorldFip technology has only been addressed on a conceptual level. Specific distributed applications in an operational environment should be carried out in order to estimate its applicability as a level 2 networking technology. This requires the development of a level 1 and level interface. The LynxOS operating system enables implementation of a gateway between the level 1 and 2 networks whilst preserving the network characteristics and performance.

The combination of WorldFip technology and ATM by means of Real Time operating systems such as LynxOS could provide a powerful solution to a range of distributed applications requiring controlled communications latency.

8 References

- [1] Gainsburg, D., ATM Solutions for enterprise internetworking, Addison-Wesley, T.J. Press, Cornwall, 1996.
- [2] McDyson, D.E *et al.*, ATM theory and application, McGraw-Hill Series on Comp. Comm., NY, 1995.
- [3] Ebrahim, Z., A brief Tutorial on ATM, March 1992, <http://av.avc.ucl.ac.uk/tltp/atm.html>
- [4] Costa, M. *et al.*, Results from an ATM based Event Builder Demonstrator,
http://pcvlsi5.cern.ch/MicDig/rd31/demonstrator_1.html
- [5] Holmgren, D. *et al.*, http://www.hep.net/chep98/index_abstracts.html
- [6] Cole, B. *et al.*, http://www.hep.net/chep98/index_abstracts.html
- [7] Lebayle, B. , <http://www.esrf.fr/computing/cs/sysadmin/network/atm/atm.html>
- [8] Blacklar, K. *et al.*, http://www.hep.net/chep98/index_abstracts.html
- [9] Fujikawa, M. *et al.* ICALEPCS conference, Beijing China, 3-7 September, 1997.
- [10] Newman, H. *et al.*, <http://nicewww.cern.ch/~davidw/public/ATMapps.htm>
- [11] Wiesel, A., Improvement of Interrupt response time and Handling, Creative Electronic Systems, S.A., Geneva, Switzerland, April 1997.
- [12] <http://www.idt.com/docs/4521.pdf>
- [13] http://www.idt.com/news/May98/5_21_98_1.html
- [14] Stevens, R.W, TCP/IP Illustrated, Volume II, Addison-Wesley, Professional Computing series, May 1994.
- [15] Cole, R, Shur, D., RFC 1932, IP over ATM, a framework document, April 1996.
- [16] Heianen, J. RFC 1483, Multiprotocol Encapsulation over ATM-AAL5, July 1993.
- [17] Rausch, R., Technical Overview of Fieldbusses, Proceedings ESONE workshop, 18-19 March 1998, CERN, Geneva, Switzerland.
- [18] Azevedo, J., The WorldFip Protocol Manual, 28 November 1996.