The Controlled Cell Transfer Capability

Steven Van Luinen, Zigmantas L Budrikis and Antonio Cantoni

Australian Telecommunications Research Institute (ATRI) and
CRC for Broadband Telecommunications and Networking (CRC-BTN)
Curtin University of Technology, Perth Australia
Email: {stevenv, zig, cantoni}@atri.curtin.edu.au

Abstract

A Controlled Cell Transfer (CCT) capability has been proposed to ITU for definition in Recommendation I.371. The CCT capability, if accepted, will bring the number of ATM transfer layer capabilities to five. The other four capabilities being Deterministic Bit Rate (DBR), Statistical Bit Rate (SBR), ATM Block Transfer (ABT) and Available Bit Rate (ABR). It is believed that this capability will endow the ATM network with LAN like service over the wide area. In this article, we discuss the features and mechanisms of the proposed CCT capability. We also discuss the applications that CCT is targeted at supporting and present results of simulation studies, highlighting the advantages of CCT.

1. Introduction

When ATM is initially deployed, it is believed that the network will not entirely be comprised of ATM, but rather islands of ATM and islands of traditional LANs and MANs interconnected. In such a scenario, the consequences for the ATM network is that LAN traffic will be the top use of that infrastructure. Even after acceptance by businesses and the move to a more fully interconnected ATM network, the top use of the network will still remain LAN based. It is for these reasons that there needs to be methods and capabilities by which LAN and MAN traffic can be carried efficiently across the ATM network. To ignore this requirement could be disastrous for the future of ATM. One method to tackle this scenario is to provide a dedicated transfer layer capability, which gives a LAN like service over the wide area.

In order for the transfer layer capability to give LAN like service, one must reflect on current LAN service requirements. Current LANs generally operate over dedicated links, with low latency, variable length message transfer. For instance, data sent to a printer by an ethernet host would typically be very large, but data sent between a file server and the host would not be as large. Communications between hosts on a traditional LAN are also typically bursty. In general, hosts will be both silent and active for different time periods and at differing times. At the same time, there are very few errors in LAN communications due partly to the higher layer protocols controlling re-transmissions. In this sense, LAN applications expect such low loss. Tariffing on LANs is on usage basis rather than connection time. Reasons for this are that when the hosts are idle, very little network resources are consumed. It makes more sense to therefore charge for usage in conjunction with a service fee, be it monthly or any other time scale.

Translating the above requirements into ATM over Wide Area Networks (WANs), implies that the capability should support semi-permanent or permanent connections, handle bursty data amicably and provide support for delivering variable length messages with low delays. The end-to-end delays in WANs will limit the effectiveness of re-transmission to correct for loss. Other requirements for the capability are that it should provide negligible cell loss and be tariffable on usage. The performance of the capability in achieving these goals is extremely important if it is to provide an acceptable and efficient LAN like service over WANs.

The result of the search for a suitable candidate which supports all of the requirements is the Controlled Cell Transfer (CCT) capability.

To investigate the performance of the proposed CCT capability ([1]) and to determine whether in fact it satisfies all of the service requirements, simulation studies were undertaken on an ATRI developed switch. The switch supports DBR, ABT and CCT connections thereby allowing quick comparisons to be made. Architecturally, the switch is shared memory based with two levels of priority queuing.

This paper is organised as follows: Section 2 discusses CCT and what it exactly is. Sections 3 and 4 discuss the flow control procedures used for CCT at both the VP and the VC level. We discuss the service features of CCT in section 5. A discussion of our switch implementation along with numerical results is presented in section 6. Section 8 concludes.

2. Controlled Cell Transfer Capability (CCT)

The CCT capability is an ATM layer capability. It is a transfer capability aimed at supporting LAN like services or computer communications. Expectations are that with this capability, all of the advantages of the ATM network such as quality in service and integrated applications can be extracted.

The basis of the proposed capability is a flow control mechanism. Flow control mechanisms are credit based and employed at both the VP and VC levels. For the VP level, the flow control is by means of a classical sliding credit window. Whilst at the VC level, the flow control employed is Backward Explicit Congestion Notification (BECN) messages. In both flow control cases, RM cells are used to provide the feedback path for the mechanisms. Because there can be configurations whereby a node may act as both a sender and a receiver, definitions of flow direction must be specified. Regardless of orientation (point-to-point case), the direction of flow control via RM cells is always from the receiver back to the sender. Similarly in a point-to-multipoint case, the direction is from the leaf towards the root. These directions are illustrated in Figures 1 and 2. For the formats of the RM cells, the reader is referred to either [1] or [3].

All of the flow control procedures used are hop by hop mechanisms. In our definition of a hop, we accept that not all nodes will employ CCT, and so a hop may encompass multiple, non supporting CCT nodes. However, in these cases the definition of hop is still between adjacent equipment supporting CCT. As would most likely be the case, not all nodes will support or need to support the CCT capability. To account for situations like this, the definition of CCT allows CCT connections to tunnel as a VP level DBR connection through this non supporting equipment. The allowance for the capability to specify a Peak Cell Rate at the VP level, facilitates this property of the CCT capability. An example illustrating this is given in Figure 3.

In the full definition of CCT to ITU, specifications on how the flow control procedures are affected by the interface are required. On this basis, the flow control mechanisms are required at both the UNI and the NNI in order for it to be network wide. At the NNI, the VP and VC level flow controls must be employed. However, varying degrees of the VP level flow control are permitted at the UNI depending on certain configurations. In its definition, CCT allows either: the VP based credit flow control to be employed alone, or solely the GFC based flow control as specified in I.361, or even both may be employed simultaneously. In either case, the VC flow control must be applied over both interfaces.

By including VC flow control into the CCT capability, we give it the ability to throttle single VCs. In this way, we can prevent other VCs in a given VP, from being unnecessarily throttled because of a particular VC causing overload. The VP flow control provides a 'safety net' as it were, for the total CCT traffic. This is achieved by only allowing cells to leave the sending end of a hop, when there is explicit assurance that there is buffer for it at the receiving end of the hop.

Figure 1: Backward Flow Control - Point to point Connection

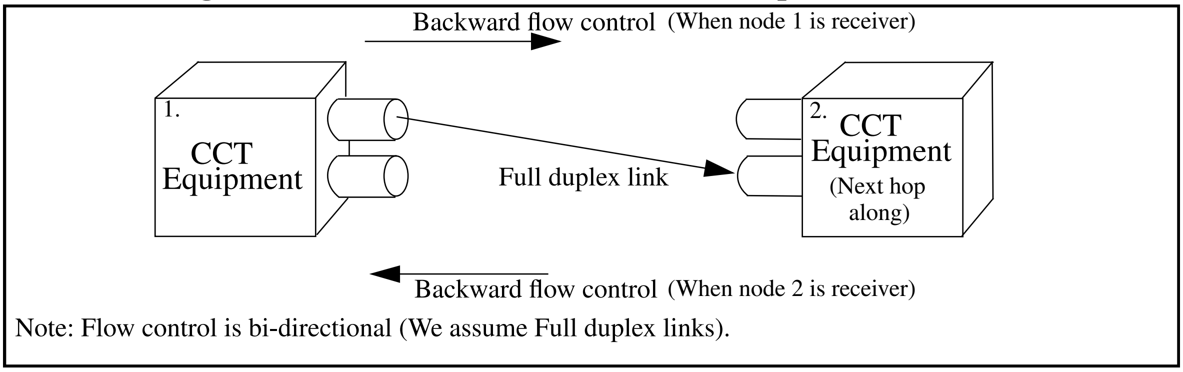

Figure 2: Backward Flow Control - Point to Multipoint Connection

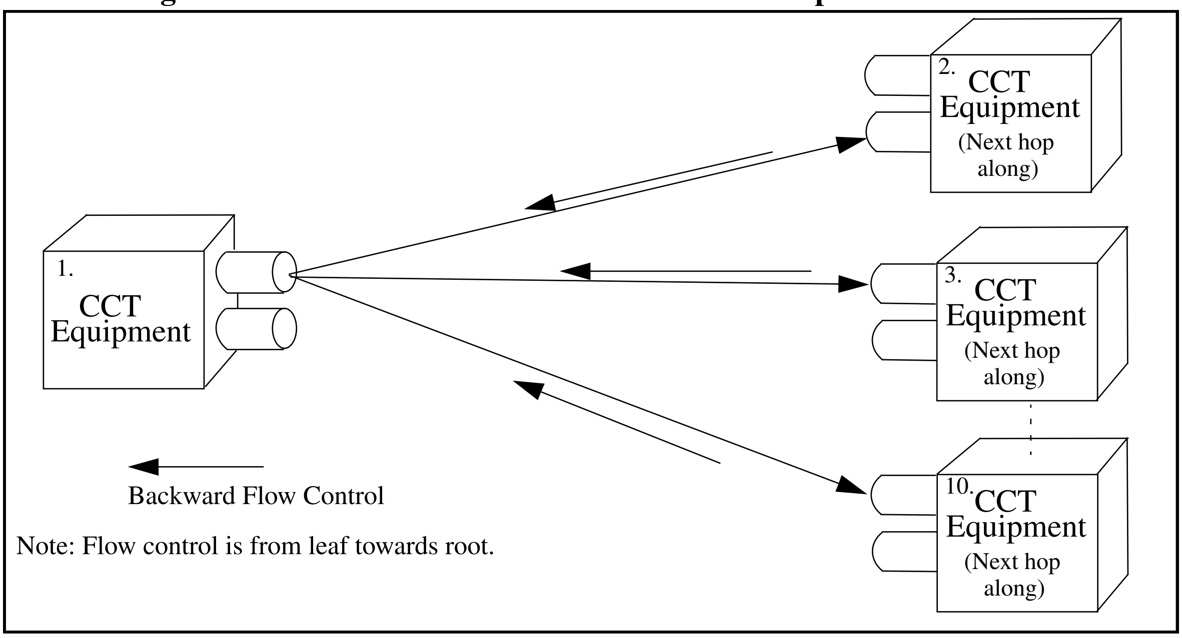

At any node, the CCT capability is of a controlled ATM virtual path (VP), containing multiple controlled virtual channels (VCs). When a QoS is requested by a connection, there is no commitment given in terms of cell delay, cell delay variation or sustainable cell rate. This is because CCT cannot make any guarantees on these parameters due to it's statistical nature. However, commitment can be made on CLR, so long as the VPC abides by the traffic contract negotiated at set-up time. If this is the case, the cell loss ratio on connections implementing the CCT capability, can be kept to effectively zero.

Figure 3: DBR Tunnelling

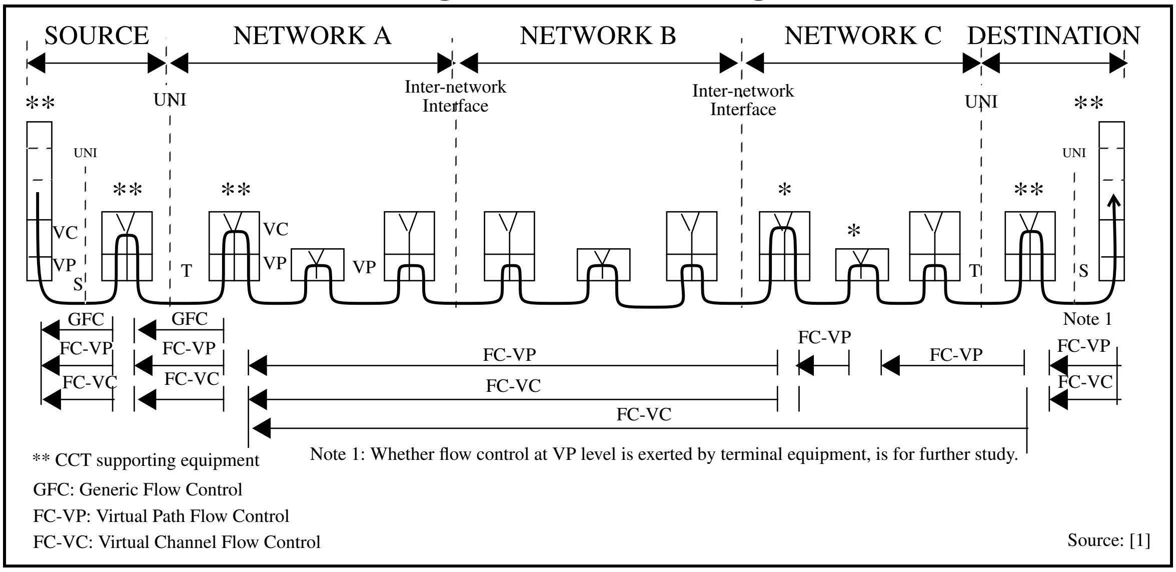

3. VP Flow Control Procedures

By definition, there are three options for CCT on the type of flow control implemented at the UNI. The three include solely VP level flow control, solely GFC, or both the GFC and VP level flow control together. In this section, both the GFC and VP level flow control procedures are presented. The procedures for the VC level flow control is left to section 4.

In any case, the procedures discussed below must be adhered to if the flow control is implemented across the UNI or NNI. It should also be noted that VP level flow control is mandatory across the NNI, but is optional across the UNI. VC (Virtual Channel) level procedures, which are mandatory across the entire network, are left to section 4.

3.1. GFC at the UNI

The GFC is an ATM layer function at the UNI and is defined in ITU Recommendation I.361. GFC is asymmetrical, which allows it to distinguish between controlling and controlled equipment. The function is on controlled connections, and is implemented on the aggregate traffic flowing from controlled equipment towards controlling equipment. There is the possibility that the control may be in two independent groups, controlled group A connections and controlled group B connections. For an illustration of which equipment is the controlled equipment and which is the controlling equipment, refer to Figure 4.

The ITU Recommendations only define procedures for the controlled equipment, leaving the procedures for the controlling equipment up to the implementer to define. The defined procedures are briefly outlined below. (Also presented in [2]).

Figure 4: Definition of Controlled and Controlling Equipment for GFC

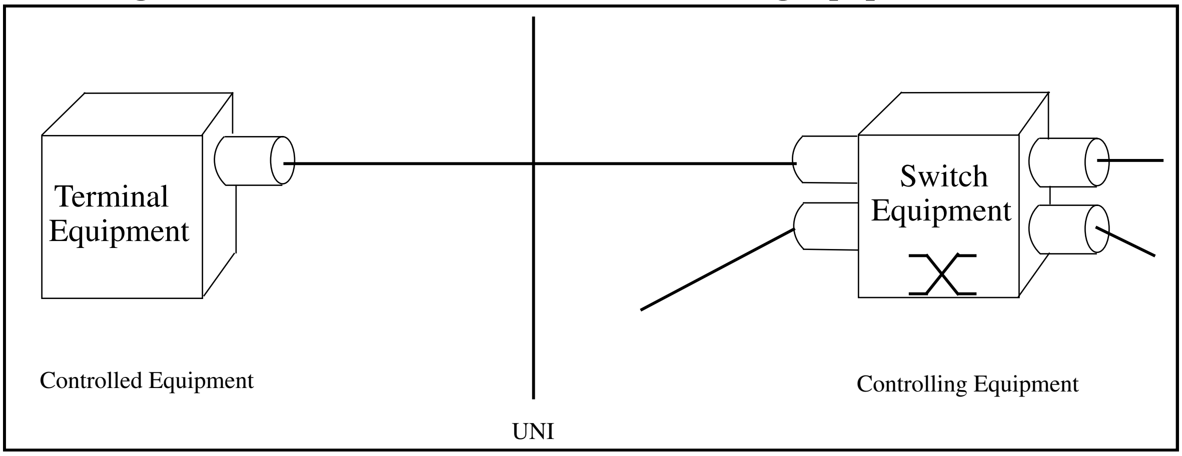

For the controlled equipment, a credit counter is maintained. This counter is known as *GO_CNTR*. For the case of two groups being supported, the two counters are *GO_CNTR-A* and *GO_CNTR-B* respectively.

On start up, the *GO_CNTRs* are reset. i.e.

$$GO\_CNTR = 0 \qquad (EQ\ 1)$$

Whenever a SET signal is received by the controlled equipment, (Or SET_A, SET_B in the two group case) the *GO_CNTR* is set to the value *GO_VALUE*. This parameter is controlled by network management, and has a default value of 1.

If the controlled equipment wishes to send a cell on a controlled connection, then it may do so provided there is no uncontrolled cell, (Cell not controlled by GFC procedures), waiting and the appropriate *GO_CNTR* is not zero. On sending of a cell on a controlled connection, the appropriate *GO_CNTR* is decremented by one:

$$GO\_CNTR = GO\_CNTR - 1 \qquad (EQ\ 2)$$

When the controlled equipment sends a cell on a controlled connection, it must set the appropriate bits of the GFC field in the cell's header. Bit 4 is always set if the cell is a controlled cell. Bit 2 is set equal to one if the cell is on a controlled ATM connection (One group option), or is on a controlled group A connection (Two group option). Bit 3 will be set equal to one only if the cell is on a controlled group B connection, and this will occur only when the two group option is being exercised. Bit 1 is always set to zero by controlled equipment.

As a consequence of the procedures defined above, the controlling equipment requires a buffer to store the controlled cells as they arrive. The size of the buffer depends on the Round Trip Delay

(RTD), the size of *GO_VALUE* and a buffer threshold (T) value. A definition of the threshold value is:

the value of buffer fill in the controlling equipment, which when exceeded, causes the controlling equipment to send NULLs instead of SETs to controlled equipment.

Therefore, the total buffer capacity required which avoids overflow is:

$$B = RTD + T + GO\_VALUE \qquad (EQ\ 3)$$

where T and RTD are in units of cell period at the link rate at the UNI

For the controlled cell transfer to run at maximum link rate, it is clear that *(T + GO_VALUE)* must at least be equal to RTD. Hence, to avoid cell loss and allow for the maximum rate, the buffer must at least be of size:

$$B \geq 2 \times RTD \qquad (EQ\ 4)$$

i.e The buffer must be at least twice the round trip delay, in cell periods, if we are to allow maximum link rate for the controlled traffic.

3.2. Flow Control on VP

The flow control to be implemented on the VP, is credit based. It is a hop by hop control, with VP RM cells providing the return control between sender and receiver (local hop definition- Figures 1 and 2). Traffic flow comprised of RM cells will be in both directions, because the use of full duplex links is assumed. In all cases, the flow control is from the receiver back to the sender. VC control is independent of the VP control, and uses BECN techniques. This control is discussed in section 4.

The flow control procedures operate as follows (Can also be found in [2]):

In a similar way that GFC at UNI has a credit counter, so does the sending end in the VP flow control. The sending end's credit counter starts with an initial value of I_c credits. i.e.

$$GO\_CNTR = I_c \qquad (EQ\ 5)$$

The equipment sends cells on the VP only while *GO_CNTR > 0*, and for each cell sent, *GO_CNTR* is decremented by 1:

$$GO\_CNTR = G0\_CNTR - 1 \qquad (EQ\ 6)$$

When the sending end receives an RM cell indicating the return of N_c credits, the *GO_CNTR* is incremented by N_c:

$$GO\_CNTR = GO\_CNTR + N_c \qquad (EQ\ 7)$$

However, an RM cell is only sent by the receiving end when it has successfully received and forwarded on N_c cells. We can now determine what the maximum possible rate is in terms of I_c, N_c and RTD.

If $I_c \geq RTD + N_c$, then we can operate at link rate.

However, if $I_c < (RTD + N_c)$, then the maximum rate for the controlled traffic is:

$$R_{max} = \frac{I_c}{RTD + N_c} \tag{EQ 8}$$

The size of the buffer required at the receiving end which avoids cell loss and allows the VP to run at full rate, and the initial credit value at the sending end are clearly, $RTD + N_c$.

It is significant that the size of the buffer required is actually dependent upon I_c, a network controlled parameter. It is not dependent upon a physical quantity such as RTD, and is therefore not swayed by variances in such a parameter. The reason for this dependence is because I_c is the maximum number of cells that can be in transit any one time from a particular source. Therefore the buffer at the receiving end requires only enough room for I_c cells.

4. VC Flow Control Procedures

Independent VC flow control is provided using Backward Explicit Congestion Notification (BECN) messages. These messages are sent using VC RM cells, and are sent in the direction from receiver to sender. Being a hop to hop flow control mechanism, the sending and receiving ends are one hop away from each other. The flow control process operates as follows:

1. When the receiving end decides that a particular VC is congested, an RM cell is sent backwards to the sending end. Specifically, the RM cell is sent on the VC that the congestion occurs. The RM cell indicates congestion and sends with it a parameter, Ts, for which suspension of transmission is requested.

2. On receipt of an RM cell notifying congestion, the sending end stops all transmissions of cells on that particular VC. Transmission is stopped until a Decongestion notification message is received or time Ts expires. Time Ts is a necessary element of the VC level flow control because it adds robustness to the flow control mechanism. On occasions when a Decongestion notification message is lost, the terminal can recover from its stopped state through the expiration of Ts. An appropriate value of Ts to use is of course implementation specific, but an example as quoted in [1] is that it could be an integer value of 10 milliseconds. The above can be summarised as:

No cell is sent if: (i) A decongestion message has not been received and (ii) $t \in [t_a, t_a + T_s]$*, where* t_a *is the time when the RM cell was received.*

Determining when to send congestion and decongestion notification messages, is implementation specific.

5. Services Achieved by CCT

From the above discussion, it is clear that CCT satisfies the service requirements specified in section 1. Firstly, because of the credit based flow control mechanisms used, it can guarantee zero cell loss. At the same time, again due to the flow control mechanisms, tariffing can be achieved on a usage basis. The basic reasons for the above two points stem from the credit based mechanisms used. Put simply, no cell can be transmitted to the downstream node without explicit assurance that there is buffer available for it. In this way, no cell can be in transit without there being room for it and consequently no cell can be lost. Similarly, because a cell can only be transmitted when there is a place for it at the other end, we have control over the number of cells sent. Therefore, we can easily tally up the number of cells sent and tariff on that basis. Recall also that in section 3.2, we can determine the exact amount of buffer required by changing the number of credits allocated. We can therefore have even more control over cell loss.

Because no messaging or other CCT specific information flows down the connection whilst the source is idle, very little network resources are used. Semi-permanent connections and permanent connections therefore make good sense. Note however that all of the capabilities can support semi-permanent connections, but it doesn't always make sense to do it that way. For example, it would be rather inefficient to have a semi-permanent DBR connection because the cost would be astounding. This would not be the case with CCT.

Finally, as the simulation studies in section 6 will show, CCT supports variable length bursty data amicably whilst also giving low latency. It outperforms any of the other capabilities due primarily to their servicing aspects.

6. Simulation Studies

Simulations have been undertaken to determine the performance of the CCT capability compared to the DBR and ABT capabilities under bursty, variable length message sources. An ATRI developed switch was used to simulate the capabilities, and its operation is discussed in section 6.1 below. The simulation configurations are discussed in section 6.2 with the numerical results presented in section 6.3. Note that all of the results discussed below have been published in [5].

6.1 Switch Implementation

A shared memory based switch with two levels of priority queuing was designed at ATRI for CCT simulation purposes. A feature of the switch designed is that it performs the Cell Transfer Protocol (CTP), thereby allowing transfers to be both fair and efficient. CTP operates on the basis of providing each input port with a set amount of controlled cells that it may queue in the switch. This set amount is known as QUOTA, and once an input port reaches this QUOTA, a STOP control is applied to it. When STOP control is applied, the input port enters controlled mode and cannot place any more cells into the switch. As cells from this input port are transmitted out of the switch, the input port may replenish these transmitted cells with more cells, providing that the quota is not exceeded.

Within the switch, cells are served in FIFO order at two priorities. The highest priority queue is for uncontrolled cells (DBR) and management cells, and the second is for controlled cells (CCT cells) and ABT cells. Only when no uncontrolled cells are waiting will the controlled cells be serviced. To gain a better insight into the operation of the switch, two flow diagrams illustrating the algorithms used for each input port and output port, are presented below in Figures 5 and 6. Figure 5 illustrates the algorithm for each input port in the switch while Figure 6 illustrates the algorithm for each switch output port. Things to note about the two flow charts are: The flow for each side (input port and output port flow) takes one cell period to accomplish, Controlled cells are only sent to the switch when the STOP control is not being applied (i.e. The input port has not exhausted its quota), A counter for each input port takes care of the quota being reached, A multicast counter tells us how many times the cell must be transmitted before it can be dropped (For multicast situations), There is only one space available in the line out card for controlled cells and inport_id holds the input port number on which the cell arrived on. The inport_id allows us to decrement the appropriate controlled cell counter for the switch (CTP). To be complete, a diagrammatic structure of the switch is illustrated in Figure 7.

Figure 5: Input Port Algorithm Flow Diagram

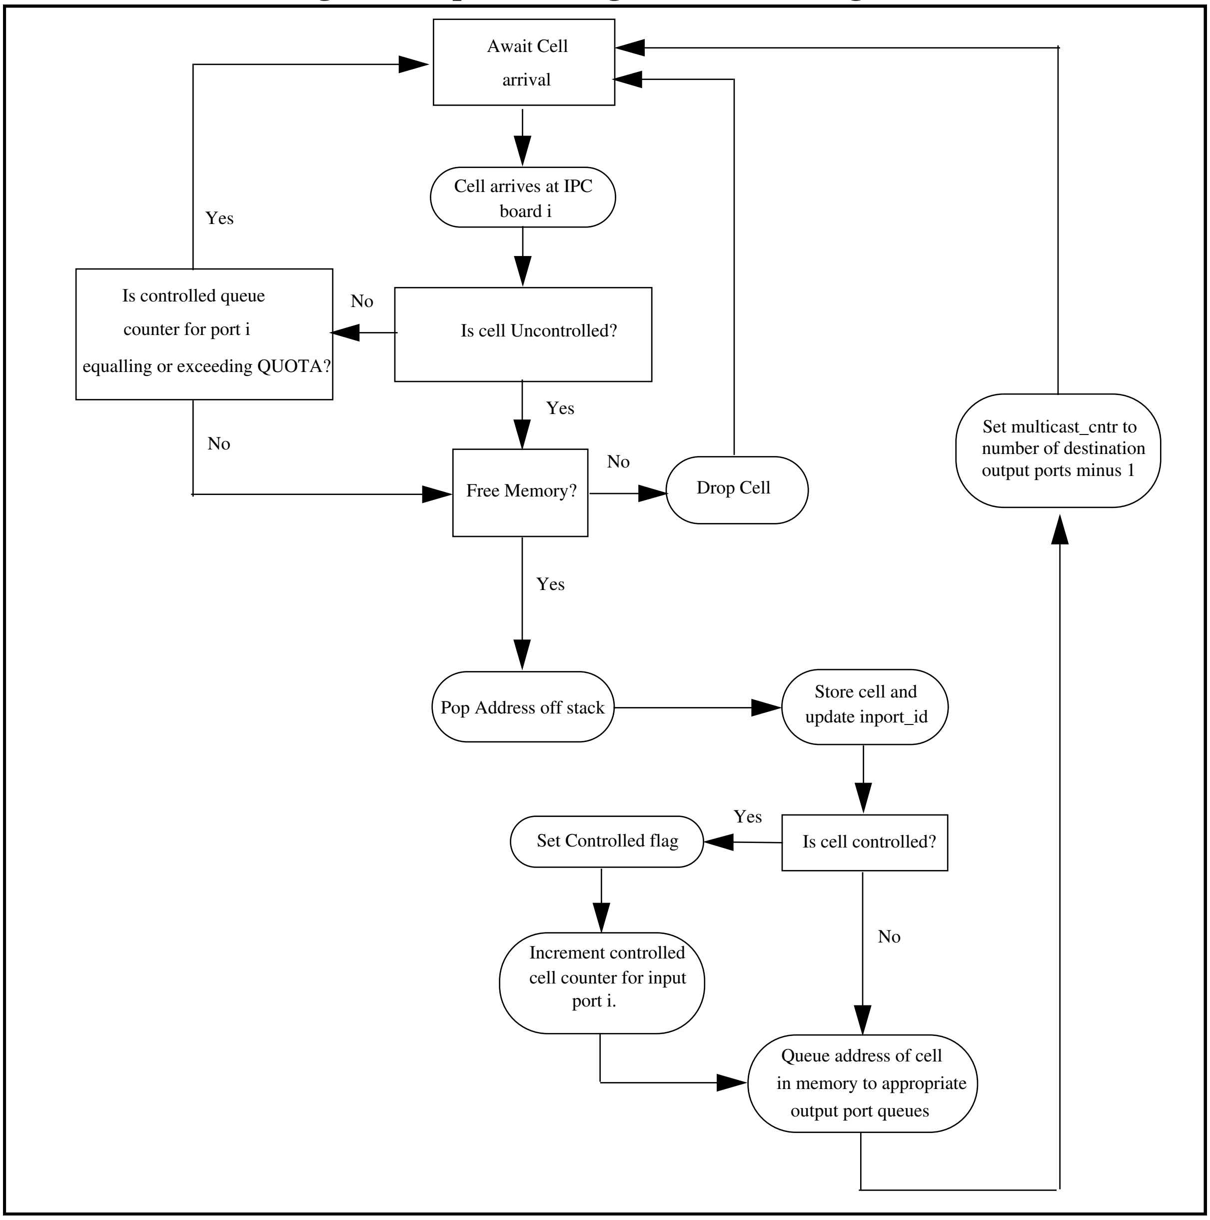

Figure 6: Output Port Algorithm Flow Diagram

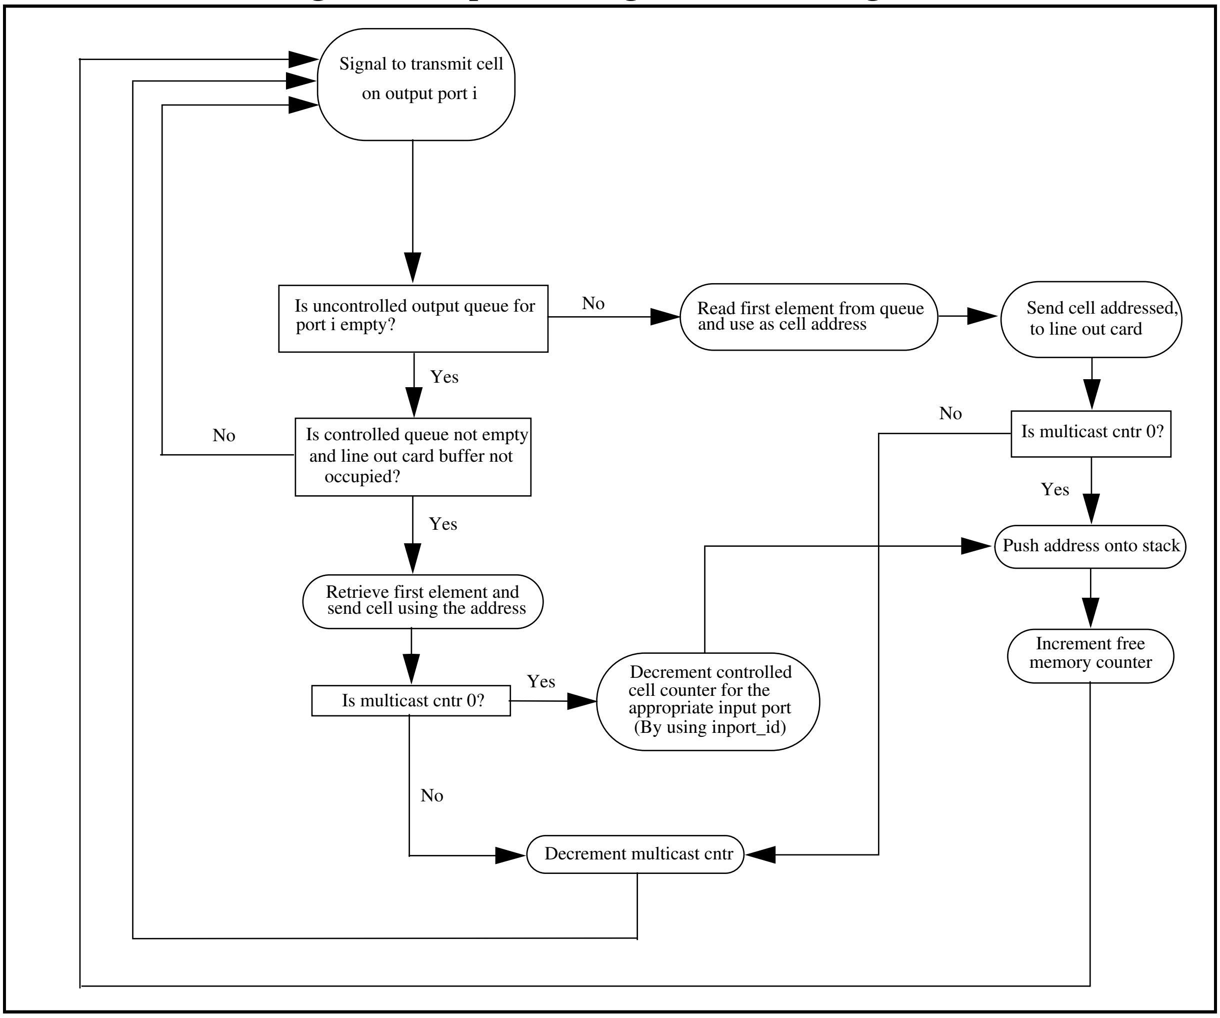

Figure 7: The Switch Architecture

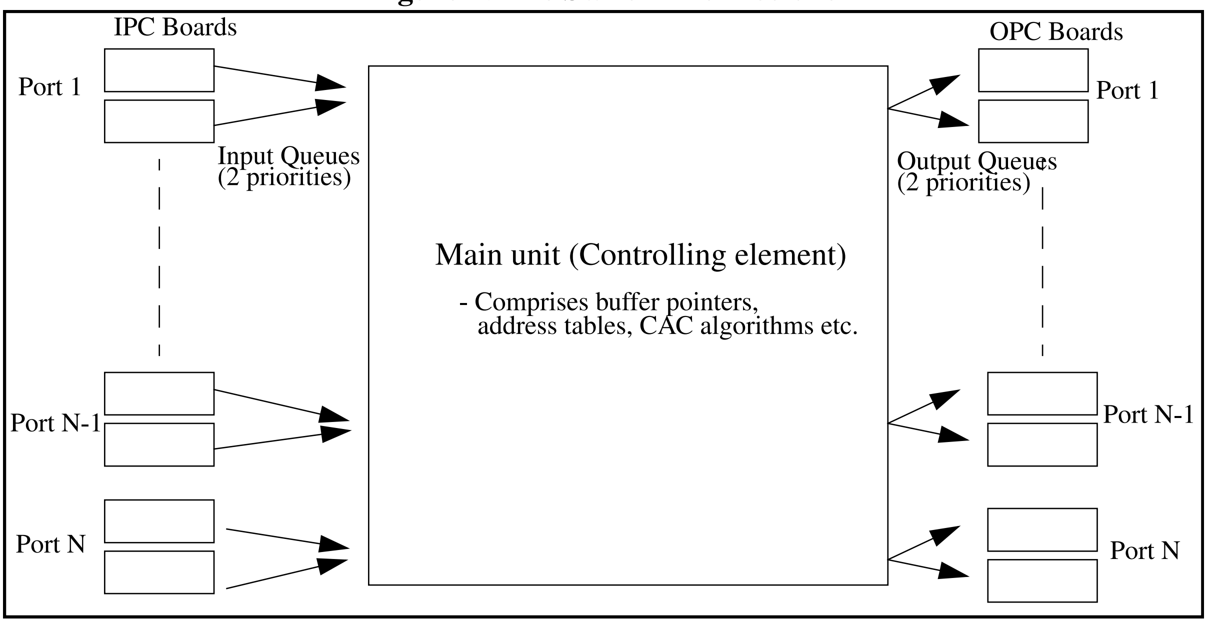

6.2 Simulation Configurations

Figures 8 through 10 show the configurations simulated for each of the three capabilities. The configurations shown represent is section within the network where there is a single link which is a bottleneck. The portion of network over which a source reaches the bottleneck may have a number of hops, but is represented by a single short hop. Similarly, the number of hops over which the traffic reaches the end user is also represented by a single short hop. The large balance of delay is assumed to be in the bottleneck. We are primarily interested in determining the smallest end-to-end delays provided by a transfer capability to bursty data traffic, of which bottlenecked elements contribute the largest balance. Therefore, if a transfer capability can provide the lowest delay through a bottlenecked section, then inevitably, this transfer capability will also provide the lowest end-to-end delay for bursty data.

For the DBR capability (Figure 8), the simulation comprises 1 bursty source transmitting to an end user, but with only a 1% partition of the available bandwidth. Because the source is bandwidth limited, it can be considered to be a DBR connection. The configuration for the DBR capability is illustrated in Figure 8. The available bandwidth that we speak of here is a 50 Mbps slice.

For the CCT capability (Figure 9), the network simulated comprised of 100 bursty sources transmitting to a common end terminal, but sharing the total available bandwidth (i.e. 50 Mbps). In the case of the ABT-DT capability (Figure 10), we again simulated a network of 100 bursty sources each vying for a share of the 50 Mbps allocated. From the workings of the ABT-DT capability, we were also required to determine and simulate different bandwidth allocation strategies. Of those chosen, we present only two, namely full rate and quarter rate. By bandwidth allocation, we mean the amount of bandwidth allocated to an ABT connection when it requests resources. The ABT-DT capability operates by sending out an RM_Request cell each time it has an ATM block (Group of ATM cells) to send. In the request cell, the connection requests the amount of resources it wants reserved for it. So by full rate, we mean that the RM_Request cell requests the whole 50 Mbps for each connection and by quarter rate, the connection only requests 12.5 Mbps. In the full rate case, only one connection can exist on the link at any one time, whereas in the quarter rate case, four connections can exist at the same time. An ABT connection is therefore similar to a DBR connection but doesn't permanently keep the bandwidth. A more detailed explanation of ABT-DT can be found in [5].

In every case that we speak of bursty source, we mean the following: The source generates exponentially distributed 'Chunks' (Or higher layer packets), with a mean of 20 cells. If a 'chunk' exceeds 105 cells, then it is segmented into two or more packets and transmitted with an inter-packet delay of 5 cells. The maximum packet length of 105 cells corresponds to a maximum packet size of 4.5 Kbytes, assuming a 44 byte payload. Interleaving of packets occurs if there is one or more packets still waiting to be transmitted when another arrives.

Statistics recorded for each of the simulations are packet delays and chunk delays only. Delay statistics are calculated as being the difference between when the first bit of the first chunk is transmitted at the source and the last bit of that same chunk is received at the destination. Similarly for packet delays.

Figure 8: Network Configuration for Source Allocated a 1% Partition of Bandwidth

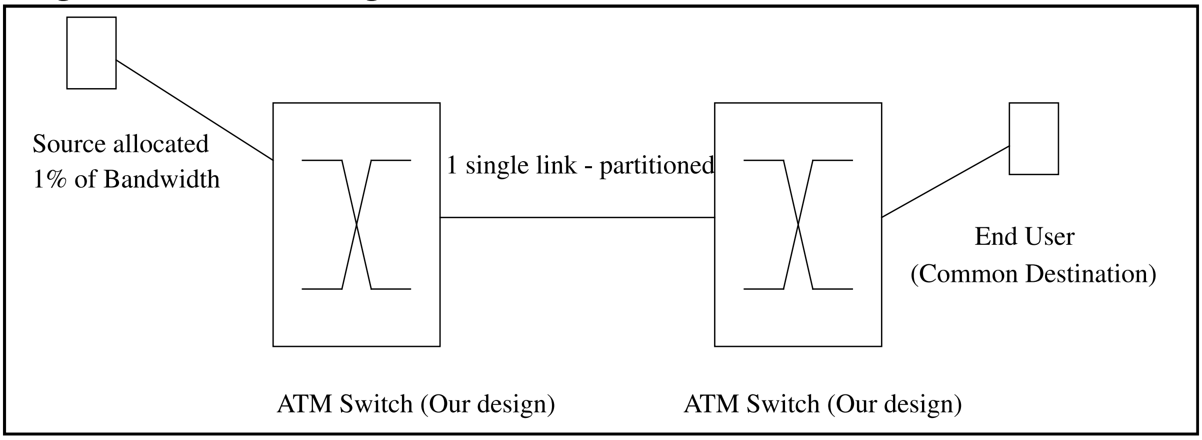

Figure 9: Network Configuration for Shared Link Simulations

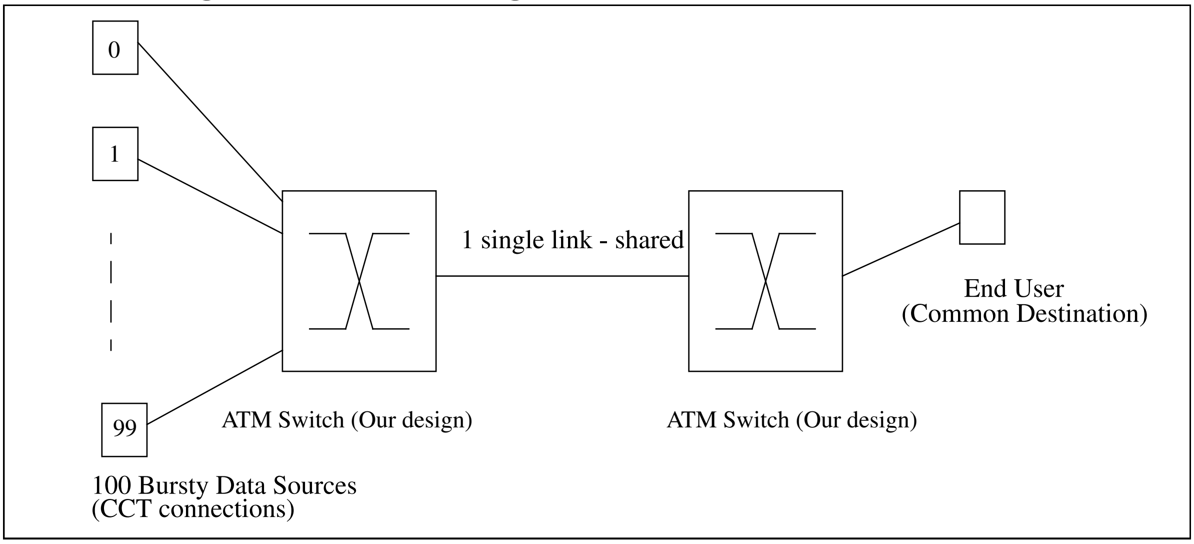

Figure 10: Network Configuration for ABT-DT Simulations

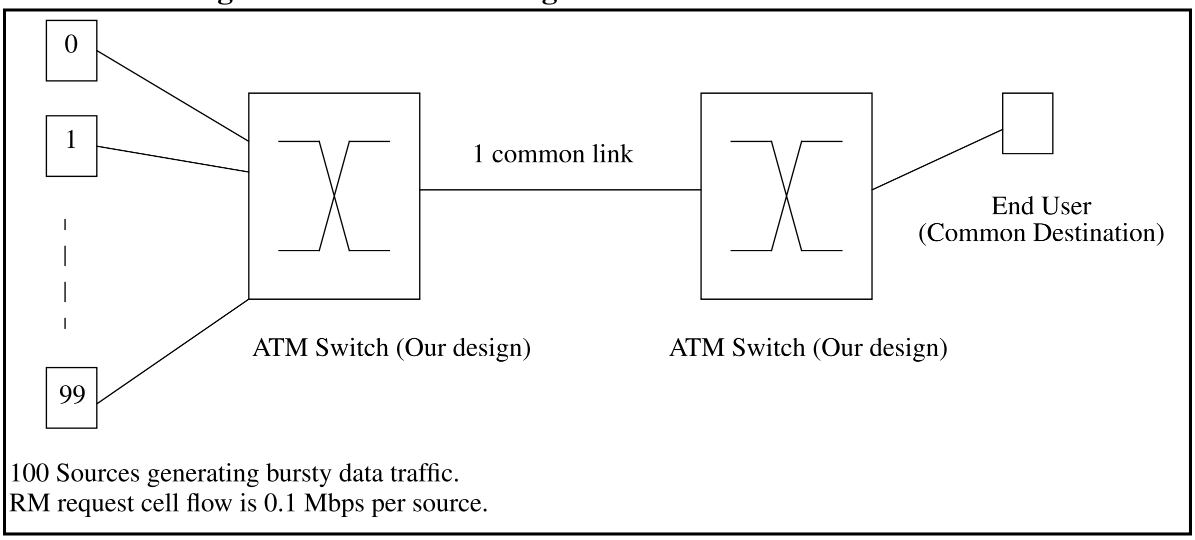

6.3 Numerical Results

The results of simulating the above configurations for 10 seconds are now discussed.

In Figure 11 below, results for the DBR case are presented. The results presented are cumulative histograms of 'chunk' delays.

Figure 11: Chunk Delays for 1% Partition - DBR Capability

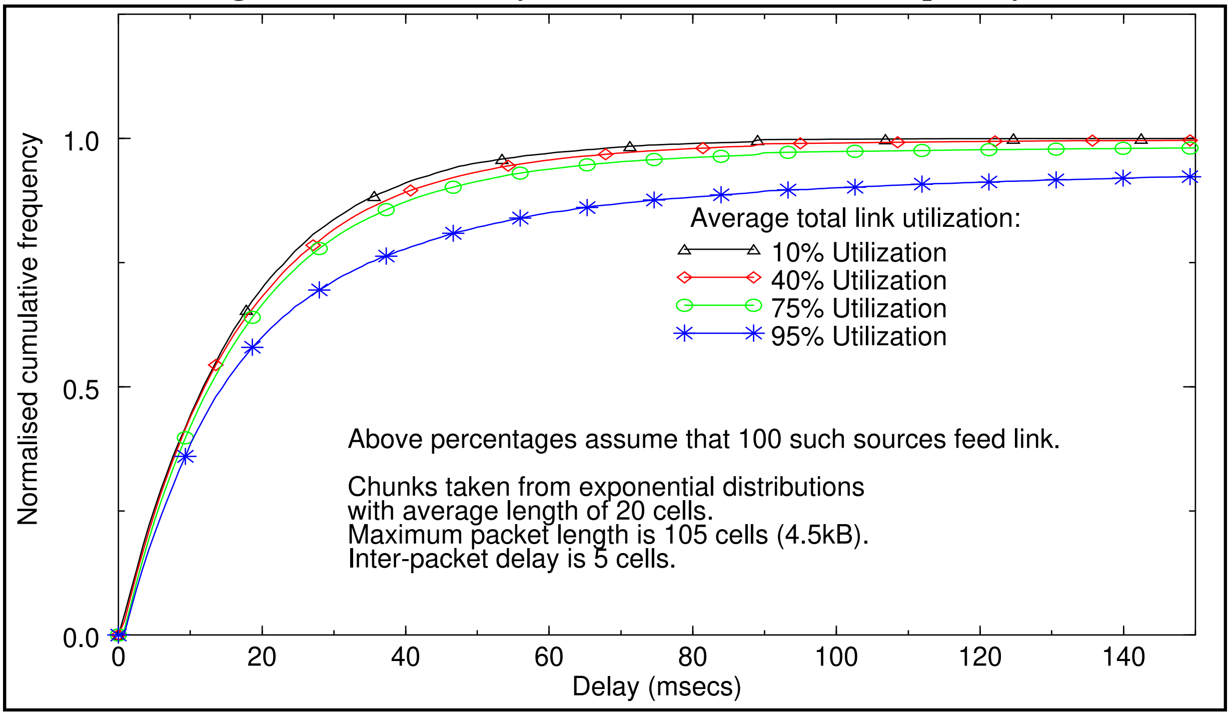

In this case, four simulation runs were undertaken, corresponding to average link utilisations of 10%, 40%, 75% and 95% respectively. From the above graph, it is apparent that the delay distributions were very similar in all but one of the four cases of link utilisation. The only case which was not similar was the 95% case. However, the only noticable difference between the 95% case and the other three cases is in the delay experienced. This is quite intuitive because you would expect that as the load increases, so too would the number of 'chunks' requiring transmission and therefore longer delays in fully receiving those 'chunks' would occur. One major item to note is that of the time scale used on the horizontal axis. Note that in the 10%, 40% and 75% cases, that the longest chunk delay is approximately 120 msecs. From the scale of the graph, it is impossible to determine the maximum delay experienced by 'chunks' under 95% loading.

For the CCT capability, the results obtained are presented in Figure 12 below.

Figure 12: Chunk Delays for 100 Bursty sources on Shared Link - CCT Capability

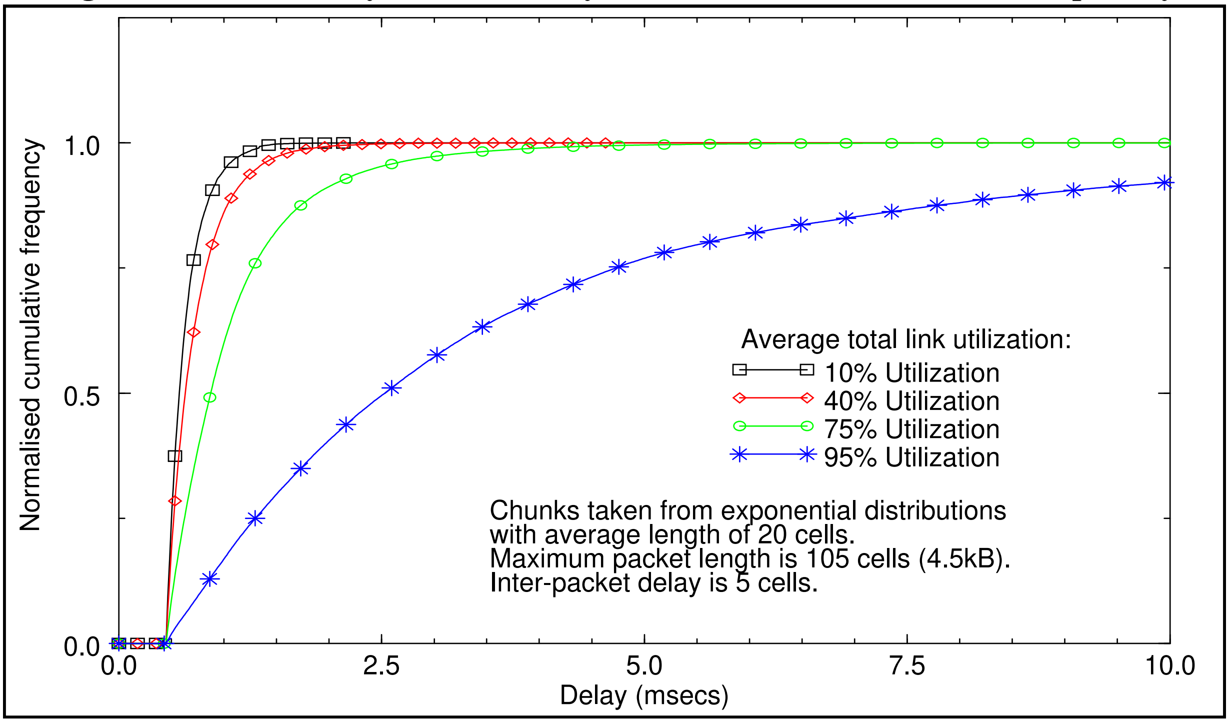

Again we simulated four different traffic intensities and recorded the delays experienced. As was the case for the DBR capability, the 95% utilisation curves are a lot lower than the other three, while the other three are very similar, differing only slightly at the beginning. Reasons for this difference is again due to the loading on the link. As the loading increases, so too do the number of ‘chunks’ requiring transmission. As there is only finite space on the link, inevitably, more and more ‘chunks’ get delayed whilst waiting to be transmitted. The effect of this is to increase the total delay in receiving the ‘chunk.’ In the case of 95% utilisation, the loading is probably at its limit and so the curve effectively represents a worst case scenario.

If we now note the scale in this particular figure, we can see that for the 10%, 40% and 75% utilisation curves, the maximum delay is approximately 5 msecs. If we now compare this figure to the one obtained previously, we can see that DBR gives an overall delay that is 24 times larger than the delay given under CCT.

We now present the ABT results for two different bandwidth allocation strategies.

Figure 13: Chunk Delays for ABT Capability - Full Rate Bandwidth Allocation

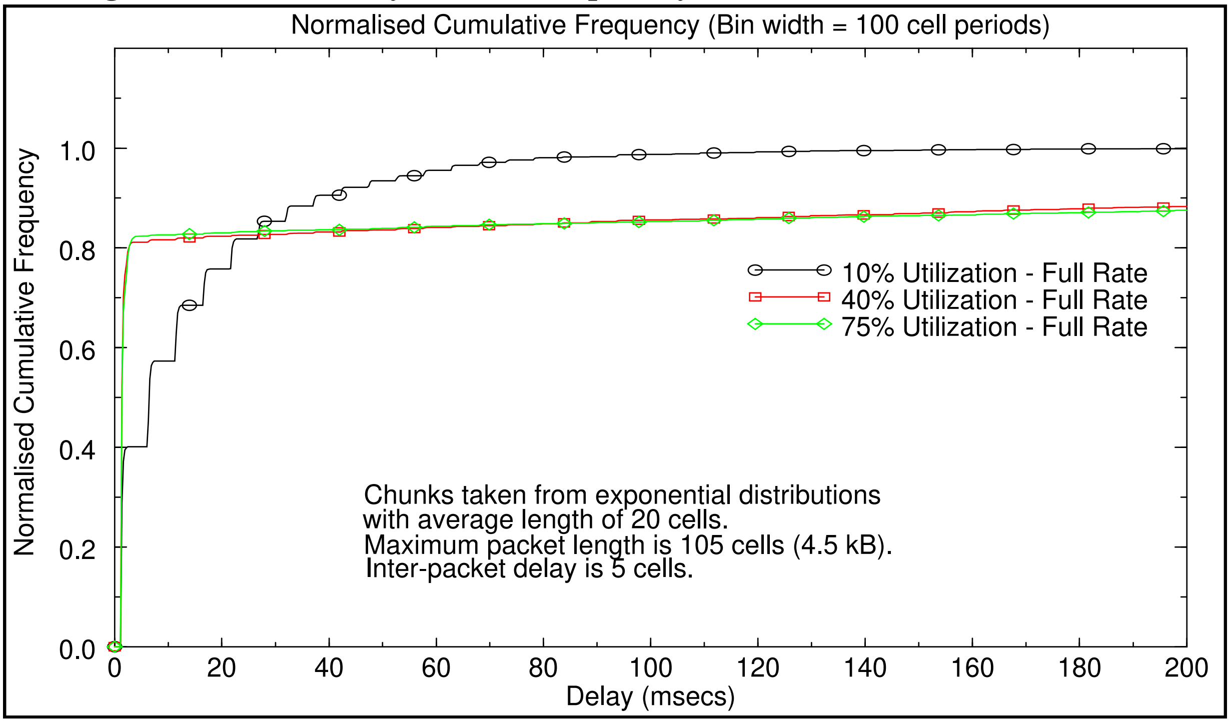

Figure 14: Chunk Delays for ABT Capability - Quarter Rate Bandwidth Allocation

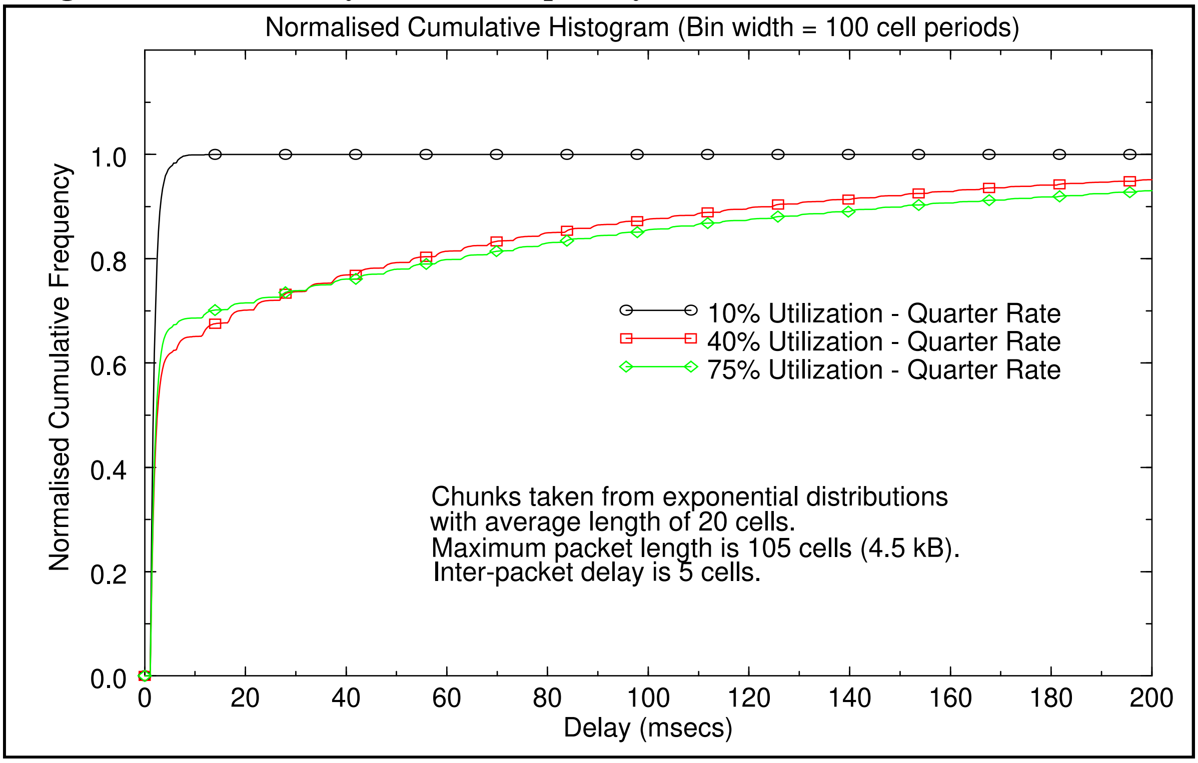

In Figures 13 and 14 above, it is clear that under the quarter rate bandwidth allocation, that the maximum delay experienced by ‘chunks’ is smaller than under the full rate allocation. This can be attributed to there being only one source using the link under the full rate allocation whereas four can simultaneously use the link under the quarter rate allocation. Hence, under the different traffic utilisations, the maximum delay will be reduced under the quarter rate as opposed to the full rate. The differences are far more noticable for traffic utilisations greater than 10%. This is because at the 10% rate, very low activity is occurring in each source. Hence, when a source does eventually have something to transmit, it can do so rather quickly regardless of whether or not the allocation is full rate or quarter rate. In general the link will not be heavily loaded and so the chance of finding it in such a manner is also low.

If we now note the scales on the graph, we see that the maximum delay is somewhat larger than 200 msecs. Therefore, the ABT capability will give maximum delays more than 40 times larger than the CCT capability. The CCT capability is therefore far superior under these conditions.

From the above results, we have shown that under bursty, variable length message sources, CCT delivers the messages with low delay. CCT therefore satisfies the service requirement of transmitting variable length messages with low latency.

7. An Application

One application for CCT is in support for multipoint Virtual LAN connections. For example, consider a multi-department business separated geographically from each other. Each department may wish to be connected to the other departments via the ATM network, so that they can share applications and other business information. The capability chosen to implement this requirement is crucial to the overall operation of the Virtual LAN (An appearance of a LAN).

What is needed for the support of such Virtual LANs is that the network should provide bandwidth that can be dynamically shared, the connection should have a guarantee on the aggregate bandwidth. Of the current capabilities, DBR and SBR cannot efficiently share bandwidth. The reason for this is that if we have a multipoint-to-multipoint connection, how would we control individual client access to the connection? With DBR and SBR, we can only assign certain rates for each client access, but then if only a few clients are transmitting, the connection may be only half used. This would not justify the cost in keeping such a connection alive. ABR on the other hand, does not give us a guarantee on the bandwidth and the policing aspect is a concern. For instance, if one server is dominating transmissions onto the connection, then it will subject all other traffic flowing, to the ABR constraints and control. It therefore does not give control over individual links. CCT however, does allow efficient sharing of network resources and can guarantee bandwidth.

8. Conclusion

We have discussed the service requirements needed to provide LAN like service for ATM, along with reasons for looking at this area. The CCT capability was presented and discussed as a capability aimed solely at supporting LAN like services.

It was also shown through intuition and simulation results that the CCT capability satisfied all of the service requirements specified. In particular, the simulation results showed that CCT outperforms both DBR and ABT capabilities in terms of the delay experienced by ‘chunks’, when all three are placed under the same conditions. This is not surprising considering that the DBR and

ABT capabilities were not designed to handle just bursty traffic. CCT was however, designed for just that purpose and the results are evidence of this.

CCT is therefore an efficient capability, capable of supporting LAN like services over the wide area.

9. References

[1] Z. L. Budrikis, *"Proposal of Controlled Cell Transfer Capability,"* ITU COM13-Q8 Working document, Perth, November 1995.

[2] Z. L. Budrikis, *"Proposal for Two Classes of ATM Bearer Services,"* ATNAC document, Sydney, December 1995.

[3] S. Van Luinen, *"Controlled Cell Transfer Capability,"* ATRI document, December 1995.

[4] V. W. Wittorff, *"Controlled Cell Transfer Protocol,"* ATRI Seminar, September 1995.

[5] S. Van Luinen, Z. Budrikis, A. Cantoni, *"Performance of the Controlled Cell Transfer Capability for Bursty Data Applications,"* Document currently under review, June 1996.