Simulation of an ATM–FDDI Gateway*

Milind M. Buddhikot
milind@dworkin.wustl.edu
(314) 935-4203

Sanjay Kapoor
kapoor@dworkin.wustl.edu
(314) 935 4203

Gurudatta M. Parulkar
guru@flora.wustl.edu
(314) 935-4621

Computer and Communications Research Center
Department of Computer Science
Washington University
St. Louis, MO 63130

Abstract

Fiber Distributed Data Interface (FDDI) networks have become viable high-speed local-area networks (LANs). The research in Asynchronous Transfer Mode (ATM) switching indicates that the local-area ATM networks will soon be deployed. The problem of designing a high-performance gateway architecture for internetworking between these two important classes of networks was addressed in [6]. The design partitions the internet functions into critical and non-critical paths. The critical path consists of per packet processing, and is implemented in hardware. The non-critical path consists of connection, resource, and route management, and is implemented in software. An extensive simulation was carried out to establish the feasibility of this approach to design future high-speed gateways. The primary focus of this paper is on the description of the simulation model, and its use in characterizing the performance of the proposed gateway architecture.

1 Introduction

The future communication environment will continue to be an internetwork and will include high-speed networks such as ATM and FDDI. The future applications such as multimedia conferencing,video distribution and interactive remote visualization require high bandwidth and performance guarantees. To meet the demands of these exciting applications new protocols and architectures need to be developed. As an initial step in this direction, we have proposed a novel internetwork abstraction called Very High Speed Internetwork (VHSI) [8].

The two important components of the VHSI abstraction are the Multipoint Congram-oriented High performance Internetworking Protocol (MCHIP) and the high performance gateway architecture. The higher level protocols use MCHIP to communicate over the internetwork and obtain desired performance guarantees. The gateways in VHSI implement the MCHIP functionality in addition to the traditional functions of interconnecting diverse networks and internetwork routing.

*This work was supported in part by the National Science Foundation, and an industrial consortium of Ascom Timeplex, Bellcore, BNR, DEC, Italtel SIT, NEC, NTT, and SynOptics.

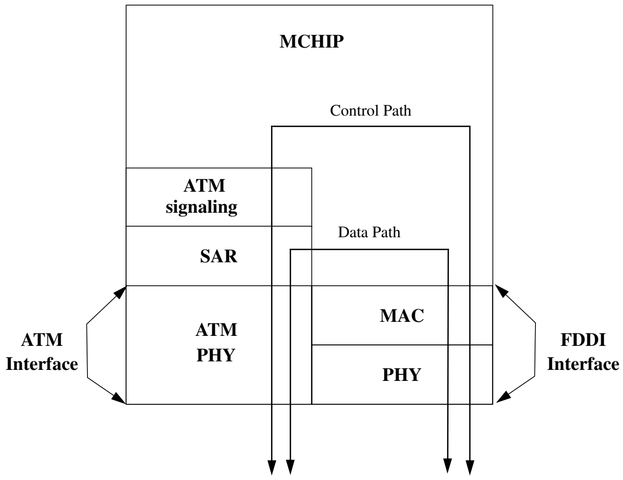

Figure 1: Protocol structure in the gateway

In order to understand design tradeoffs for future high-speed gateways, we designed a two-port gateway between ATM and FDDI networks which have diverse characteristics and show the promise of wide spread use in future. The detailed design has been reported in [6].

Figure 1 shows a high-level view of the layering and protocol structure upon which the gateway design is based. In addition to the ATM and FDDI physical interfaces, the figure also shows the ATM segmentation and reassembly layer (which implements the SAR protocol [3]), the ATM signalling layer, the FDDI media access (MAC) sublayer, and the internet layer protocol (in this case, MCHIP).

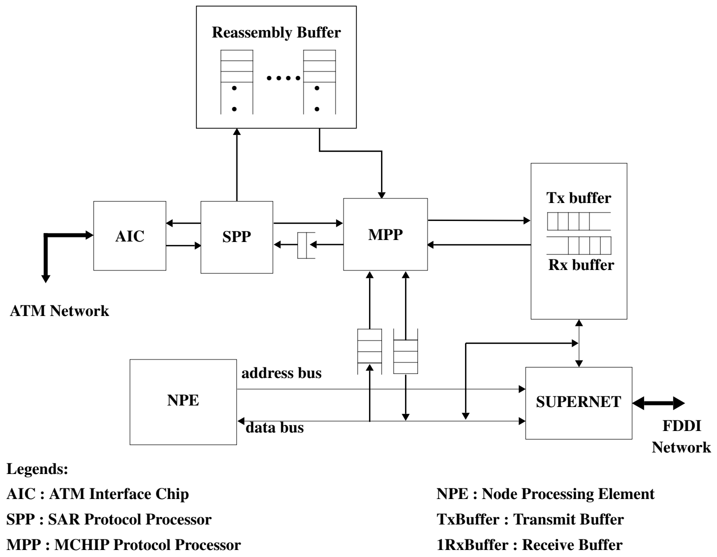

Figure 2: A two port ATM–FDDI gateway

A simplified block diagram of the gateway hardware architecture appears in Figure 2. Note that the non-critical (control) path functions are implemented in software run on a processor called the Node Processing Element (NPE). All the critical (per cell or per frame) functions are implemented in hardware using two custom VLSI chips and two commercially available chips.

ATM Interface Chip (AIC): The AIC implements the ATM physical layer protocol and provides an interface to the ATM network.

SAR Protocol Processor (SPP): The SPP implements the SAR protocol, which is responsible for segmenting data and control frames from higher level protocols into ATM cells, and for reassembling these frames upon the receipt of all the constituent cells. Part of the job delegated to the SPP involves reassembly buffer and timer management; for this purpose, it interfaces to the AIC, MPP, and reassembly buffer memory.

MCHIP Protocol Processor (MPP): The job of the MPP is to route reassembled FDDI frames or ATM cell fragments to the appropriate destination based upon the type of the cell/frame and information contained in the MCHIP routing tables. The MPP is also responsible for detecting incoming control frames and handing them over to the NPE, as well as accepting and forwarding control frames from the NPE.

The MPP also reads the FDDI frames in the receive buffer of the FDDI interface and forwards them to the SPP with appropriate state information for subsequent fragmentation into ATM cells.

Node Processing Element (NPE): The NPE is a general purpose microprocessor running software implementations of the ATM signaling protocol, FDDI connection and station management, and MCHIP congram management.

SUPERNET [1]: The SUPERNET chip set implements the physical and media access FDDI protocols in VLSI [1].

This paper presents an extensive simulation model for the design discussed. The primary objective in undertaking this simulation study was to characterize the gateway performance under various input traffic conditions and design parameters. Another objective of the simulation was to perform an extensive functional verification of our design.

The outline of the rest of the paper is as follows: Section 2 focuses on the specific performance metrics that we were interested in, and the reasons for selecting these metrics. Section 3 presents the simulation model, while Section 4 describes the various experiments conducted and the results obtained from these experiments. Finally, Section 5 provides conclusions.

2 Motivation

In this section, we present the motivation behind carrying out a detailed simulation study. Enumerated below are descriptions of the performance metrics that can be used to characterize the gateway performance:

Load vs Throughput performance of the Gateway: The load-throughput behavior of the gateway characterizes what fraction of the offered load can be successfully carried by the gateway.

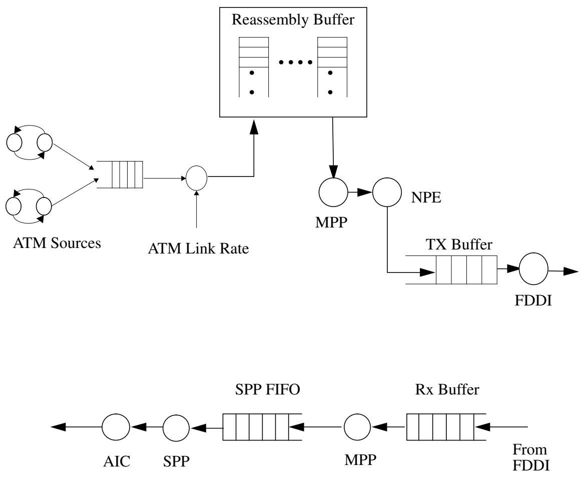

Figure 3: Queueing model for the architecture

Cell loss probability in SPP: The SPP drops cells if the reassembly buffer is full, or if the arrival rate is higher than the service rate (that is, the NPE is too slow). The cell loss probability in SPP is measured as the fraction of the cells processed by SPP that are dropped.

Frame loss probability in MPP: The frame loss probability in MPP is defined as the fraction of the total number of frames processed by the MPP that are dropped. It is assumed that the partially reassembled frames are of no use to higher level protocols or the end user, and are therefore dropped in the MPP.

Latency of completely assembled frames (through the gateway): This is the time spent in the gateway by the first constituent cell of a frame that is successfully reassembled within the gateway; it is a measure of the delay performance of the gateway architecture.

The performance metrics described above strongly depend on design parameters such as the sizes of the transmit, receive, and reassembly buffers, the processing speed of the NPE, and the maximum number of active connections supported by the gateway. In addition, they are influenced by changes in input traffic patterns. A careful analysis of the effects of these parameters on the metrics can cast light on design tradeoffs, and enable us to (over) engineer the gateway hardware so as to obtain acceptable levels of performance.

Figure 3 shows a simplified queueing model for the ATM-TO-FDDI and FDDI-TO-ATM pipelines in the gateway hardware. As shown, the reassembly buffer is serviced by a tandem server formed by MPP and NPE. This service mechanism depends on the state of the transmit buffer, such as the availability of the buffer space. The transmit buffer in the FDDI interface is serviced by the timed-token-rotation mechanism of FDDI. The receive buffer is serviced by the MPP, and the SPP FIFO is serviced by FDDI-ATM pipeline of the SPP.

The queueing and service mechanisms in the gateway architecture are very complex and an analytical performance model to characterize performance metrics seems mathematically intractable. On the other hand, a detailed function level simulation study can be done with ease. Such a study can also be used to do a extensive functional verification of the design. A graphical general purpose simulation tool called BONES[2], allows typical architectural components like multiplexors, memories, buffers, timers and processors to be modeled easily and hence is useful in developing functional model for the gateway architecture.

3 Simulation Overview

This section presents the simulation model of the gateway architecture. The model is described using the BONeS blocks developed in the course of this simulation.

Figure 4 shows a high level view of the system. The corresponding BONeS system block diagram appears in Figure 5 (A). Note that the gateway block is pictured between a model of an ATM network and a module that simulates FDDI network traffic. The latter module is part of the BONeS standard distribution. The ATM network is simulated as a set of bursty sources (representing independent connections) feeding an infinite length queue which is serviced at the ATM link rate of 155 Mbps. The particular service discipline

[2]BONeS (Block Oriented Network Simulator) is a trademark of Comdisco Systems, Inc.

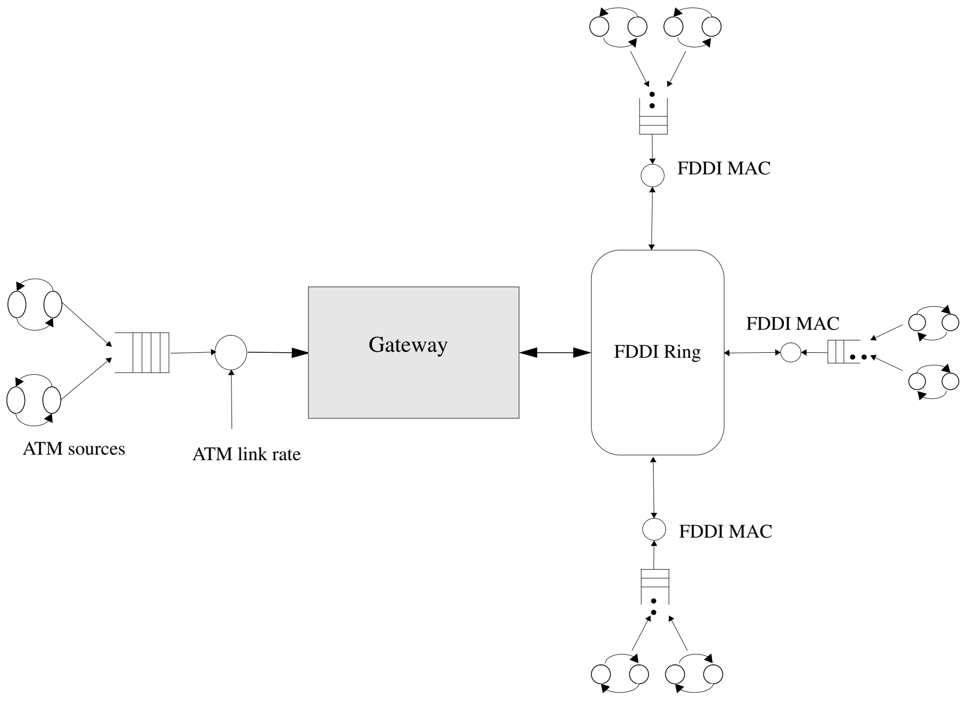

Figure 4: Simulation definition

used is First In First Out (FIFO). Each bursty source represents the network traffic on an active ATM connection. A two state Markov chain is used to statistically describe each source, as shown in Figure 5 (B). The sources alternate between an active and idle state, with the time in each state being exponentially distributed. During the active state, sources generate Poisson traffic, and during the idle state, no traffic is generated. A three tuple $< \Lambda_p, \Lambda_a, B >$ can be used to fully characterize the bursty nature of a source. Here, Λ_p is the peak bandwidth, Λ_a is the average bandwidth, and B is a burst factor that provides a measure of the amount by which the average bandwidth is exceeded during an active period.

The gateway model itself is constructed as a hierarchy of seven layers of BONeS blocks. Descent through this hierarchy corresponds to an increase in the perceived degree of detail. An effort was made to maintain a close resemblance between BONeS blocks and the hardware components comprising the gateway.

Figure 5 (C) shows the next lower level of BONeS blocks comprising the gateway block in Figure 5 (A). The SPP and the MPP have been explicitly modeled as BONeS blocks, as have the ATM interface chip and the receive buffer memories. Note that no control path traffic has been modeled; it is assumed that control traffic does not interfere with data path processing. The NPE has been modeled as a processing delay, placed in SPP-MPP acknowledgment path.

Figure 6 shows the BONeS block diagrams for two pipelines that comprise the system's view of the SPP. One of these pipelines models data flowing from the ATM interface to the FDDI transmit buffers, while the other pipeline models data flow in the reverse direction (that is, from the FDDI receive buffers to the AIC). Note that the function of these pipelines is the same as that of the two hardware pipelines required in the VLSI implementation of the SPP.

Just as the SPP, the MPP is also modeled by two pipelines. The ATM-FDDI pipeline of MPP models the NPE and the handshake between SPP and MPP in detail. The FDDI-ATM pipeline is simpler and models the table translations and header processing done by MPP on each frame. The BONeS block diagrams for these pipelines are not shown in the paper, due to lack of space.

The FDDI network block used in the simulation belongs to a library of local area network modules provided by BONeS [4]. This block models a set of nodes on an FDDI ring, with one of the nodes serving as a gateway. Figure 5 (D) shows the BONeS block diagram of a typical node in this layout. The traffic generator block shown in this figure is capable of generating synchronous traffic as well as asynchronous prioritized traffic. The FDDI MAC block implements the FDDI media access protocol. For the purposes of our simulation, we modified this block on the gateway node to reflect the handshaking between the SPP and the MPP. We have assumed that all the traffic is asynchronous [7]. The implication of this assumption is that performance results are optimistic.

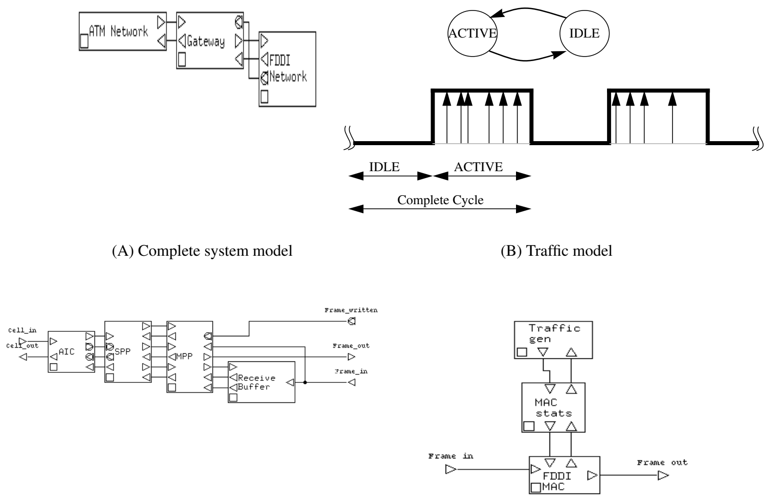

(A) Complete system model

(B) Traffic model

(C) Gateway model

(D) FDDI node model

Figure 5: Model in BONeS

4 Simulation Experiments

A number of experiments were conducted using the simulation; these experiments fall into two basic categories: functional verification experiments and performance evaluation experiments. This paper presents results of experiments in the performance evaluation category. However, we note that the functional verification experiments were immensely useful in verifying the correctness of the design as well as correctness of different parts of the simulation. The performance evaluation experiments can be classified as follows:

- Effect of input ATM traffic patterns.
- Effect of buffer size constraints.
- Effect of the NPE delay.
- Effect of FDDI parameters.
- Effect of the number of active connections.

Table 1 lists 7 parameters that remained constant throughout these experiments and their values. The behavior of the gateway hardware is largely independent of the number of stations on the FDDI network and therefore, all the experiments use a fixed four node FDDI network. The results of the experiments in the first three categories are discussed below.

4.1 Effect of Burst Factor

Recall that the burst factor is a measure of how much extra traffic the source generates over and above its average bandwidth when it goes into the active state. A higher burst factor makes the active and idle periods of a source longer. Longer active periods demand larger buffers on a transient basis. Therefore, at very high burst factors, the gateway is expected to experience a high cell/frame loss.

Table 1: Constant parameters

Parameter	**Value**
Number of active connections	8
Clock Cycle	25 ns
Maximum FDDI frame size	4096 bytes
ATM link capacity	155 Mbps
FDDI link capacity	100 Mbps
FDDI node priority	7
TTRT	50 ms

Table 2 lists parameter settings for the experiment. There are eight active connections, each of which is assumed to have a frame size (negotiated at set up) equal to the maximum size of a FDDI frame. The offered load from the ATM network is increased in steps, while all other hosts generate 1 Mbps background traffic.

Figure 7 (A) shows the variations in gateway throughput as the offered load (average ATM bandwidth) changes. This figure indicates that up to a threshold, gateway throughput increases linearly with the load and is independent of the burst factor of the source. Beyond this threshold however, the throughput begins to level off as it is restricted to the FDDI link speed of 100 Mbps. At higher loads, the throughput is lower for higher burst factors because of increased cell and frame loss at high loads. An ATM cell which

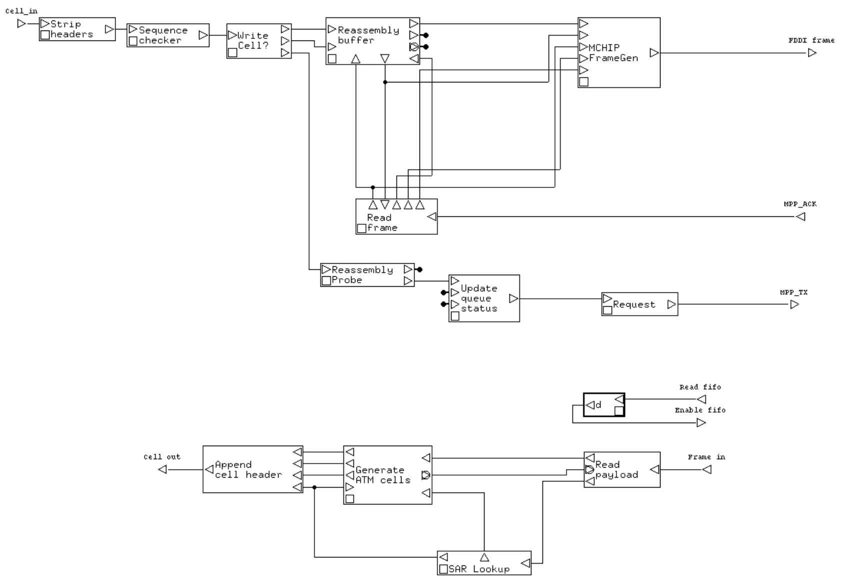

Figure 6: SPP model

Table 2: Input parameters : vary burst factor

Parameter	**Value**
Transmit buffer size	16 maximum size FDDI frames
Reassembly buffer size	2 maximum size FDDI frames per connection
Frame size assembled	85 ATM cells
Peak-to-Average ratio	2
Burst factor	1000,500,250

has an overall size of 53 bytes, contains 8 bytes of overhead information and 45 bytes of data. Gateway hardware strips off the 8 bytes of overhead information and uses it to reassemble the 45 bytes of data into a MCHIP segment. This explains why there is approximately a 15 percent difference between the offered load and the carried load in the region below the threshold bandwidth, although there is no cell/frame loss in this region.

Figure 7 (B) shows the variation of cell loss probability in the SPP with the offered load, at various burst factors. The cutoff bandwidth beyond which the SPP begins to drop cells decreases for smaller burst factors. Moreover, the cutoff bandwidth appears to be sensitive to the burst factor. At loads close to the cutoff bandwidth, the loss probability is small. However at very high loads, loss probability can be as high as 5 percent.

Figure 7 (C) shows how frame loss probability varies with a changing load. Note that the cutoff bandwidth for the frame loss probability is the same as that obtained from the cell loss probability plots. Even a small cell loss probability in the SPP translates to a large frame loss probability in the MPP. Intuitively, this is correct, since even dropping a single cell in the segment being assembled results in the rest of the segment being discarded. This phenomenon becomes especially relevant, when large segments are reassembled under heavy load.

Figure 7 (D) shows the variation in gateway latency for completely assembled frames with changes in the burst factor and the load. At low loads, the gateway latency is dominated by reassembly delay, which decreases exponentially with a fixed number of active connections and increased offered load. Observe that at very high loads, the latency experienced by a frame is independent of the burst factor. One of the reasons for this is that at high loads, the transmit buffer occupancy reaches its maximum, and incoming frames always see the transmit buffer as being full. Another reason is that the service mechanism for the transmit buffer queue is the periodic timed-token-rotation mechanism in FDDI, and so the average delay experienced by the frame when the transmit buffer is almost full is independent of the input traffic characteristics.

4.2 Effect of Peak-to-Average Ratio of the Source

The peak-to-average bandwidth ratio of a source is another measure of the burstiness of a source, since it affects the rate at which the source generates traffic when in the active state. A large peak-to-average ratio would make the source generate a large burst of traffic

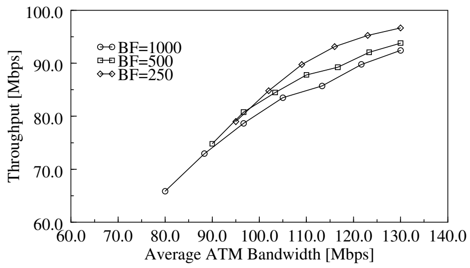

(A) Gateway throughput vs Offered load (for various burst fatcors)

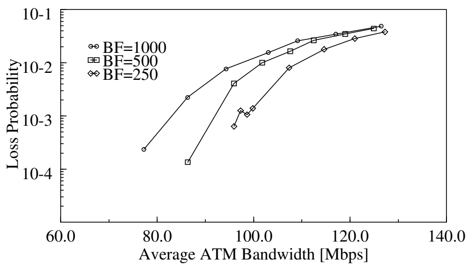

(B) Cell loss probability vs Offered load (for various burst fatcors)

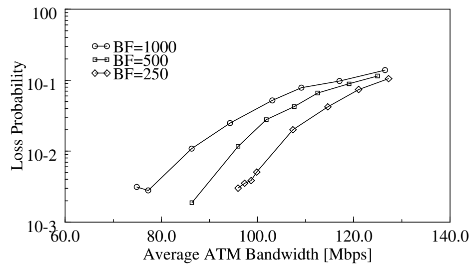

(C)Frame loss probability vs Offered load (for various burst factors)

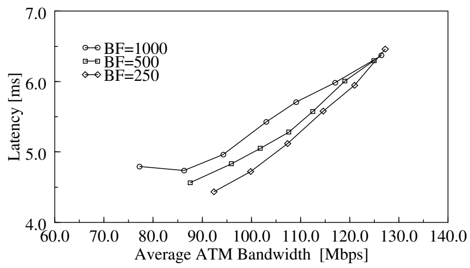

(D) Gateway latency vs Offered load (for various burst fatcors)

Figure 7: Gateway performance for various burst factors

in a short time interval; this would mean that large buffers could be requested at a very short notice.

In our experiments, the burst factor of the source was 1000. The transmit buffer size, the reassembly buffer size and the segment size are the same as in Table 2. All the experiments measure various performance metrics as the offered load is varied, for peak-to-average ratios of 2, 4 and 8. Note that at a peak-to-average ratio of 8, when the offered average load is 100 Mbps, the peak offered load is 800 Mbps, which is well over the FDDI link bandwidth of 100 Mbps. For a transient period this can strain the gateway hardware significantly. Such high peak-to-average ratios will be be common in high bandwidth applications of the future [2].

The gateway throughput performance for different peak-to-average ratios is plotted in Figure 8 (A). Observe that for high loads, the throughput is lower for higher peak-to-average ratios due to the increased frame and cell loss. The dependence of cell loss probability on the offered load for different peak-to-average ratios is shown in Figure 8 (B). At high loads, doubling the peak-to-average ratio increases the cell loss probability dramatically. For example, with an offered load of 120 Mbps and a peak-to-average ratio of 2, the source generates traffic at a peak rate of 240 Mbps. If the peak-to-average ratio is doubled, the peak rate is increased to 480 Mbps, which is around five times the FDDI link capacity. Therefore, with a high peak-to-average ratio, the gateway hardware is forced to assemble a continuous train of back-to-back cells. This causes buffer overflow resulting in significant cell loss.

4.3 Effect of the Node Processing Delay

Earlier, we presented the queueing model for the ATM-FDDI path in the gateway hardware. This model shows that the reassembly buffer is serviced by a complex service mechanism that involves the MPP and NPE, and depends on the transmit buffer occupancy. If the reassembly buffer is not drained at a fast enough rate, the SPP will begin to drop cells. An example of a scenario that could result in such a situation is when the NPE is slow and the gateway has to assemble small frames to meet the demands of a high bandwidth application. In order to study the behavior of the gateway under these conditions, two sets of exper-

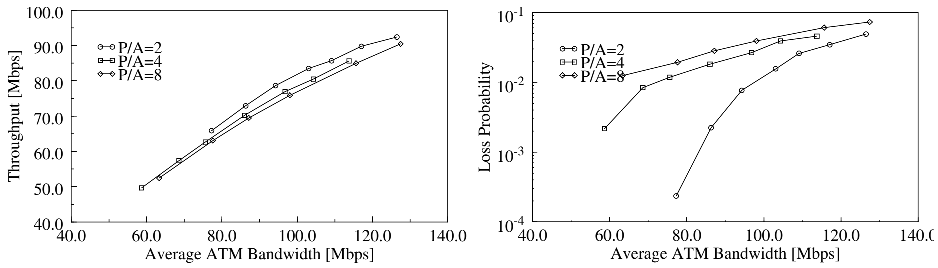

(A) Gateway throughput vs Offered load
(for various Peak/Avg ratios)

(B) Cell loss probability vs Offered load
(for various Peak/Avg ratios)

Figure 8: Gateway performance for various Peak-to-Average bandwidth ratios

iments were conducted. In the first set of experiments, a small burst factor (= 4) was used, while the second set used a moderate burst factor (= 43). In either case, the size of the frame being assembled was small. Both of these experiments were run with the transmit buffer and reassembly buffer size the same as in Table 2. The segment size was fixed at 8 cells, and a peak-to-average ratio of 2 was used. The NPE was configured so that requests were processed at a rate R, varied from 25 to 100 requests per microsecond.

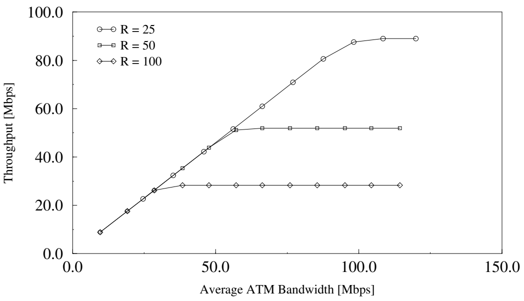

Figure 9: Throughput vs. offered load for various processing speeds

Figure 9 is a plot of the gateway throughput versus the offered load for different NPE processing rates. It is seen that for a slow NPE, the maximum achievable throughput is restricted by the NPE rate. Any offered load greater than NPE rate results in heavy cell and frame loss. In other words, given the NPE processing rate and the sizes for the receive, transmit, and reassembly buffers, there is a critical minimum segment size. If the segment size drops below this critical minimum, it could result in very high cell/frame loss. Observe also that at low loads, the cells are generated at a slower rate, resulting in a lower request rate to access the transmit buffer. The throughput linearly increases upto a "knee" point at which the request rate just matches the rate at which the NPE can process the requests. At all loads above this point, throughput remains constant and equals the maximum rate at which segment write requests from the MPP to the NPE can be handled. For loads above the knee, the NPE becomes a bottleneck and the SPP will drop cells. As expected, this will result in increased cell loss probability as the load increases. This can be verified from the Figure 10, which shows the cell loss probability variation for various processing rates and burst factors. It can be seen that the cell loss dramatically increases with the use of a slower node processor and increases exponentially with load, once the throughput saturation sets in.

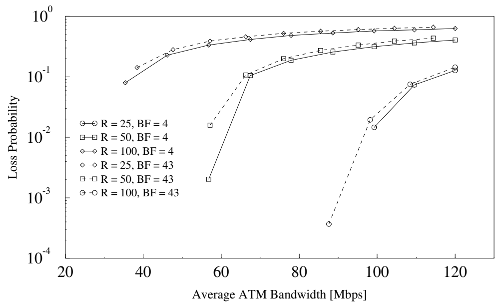

Figure 10: Cell loss probability vs. offered load for various processing speeds

4.4 Effect of Transmit Buffer Size

As long as an application uses large frames and a sufficiently large transmit buffer is available, the reassembly buffer and the NPE can be engineered so that

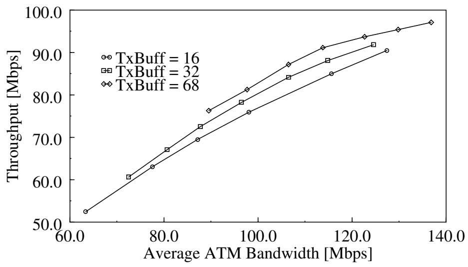

(A) Gateway throughput vs Offered load

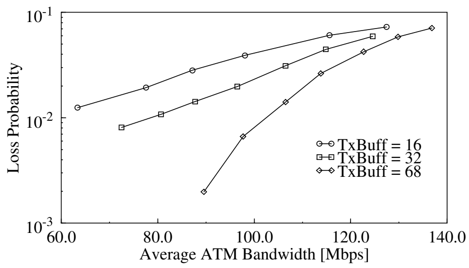

(B) Cell loss probability vs Offered load

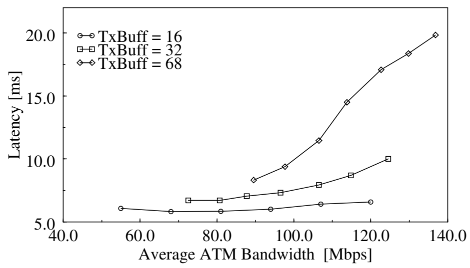

(C) Gateway latency vs Offered load

Figure 11: Performance of the gateway for different transmit buffer sizes

cell/frame loss is kept low. In what follows, we discuss the effects of transmit buffer size on cell/frame loss under very bursty traffic conditions. Note that when the traffic is very bursty and high load conditions exist, the transmit buffer becomes full more often. The MPP will place a hold on acknowledgments sent to the SPP until space is created in the transmit buffer by successful transmissions over the network. The longer the hold on the acknowledgements, higher the probability that the allocated reassembly buffer fills up and the SPP is forced to drop the incoming cells.

The transmit buffer size was set to 16, 32 and 68 full size FDDI frames during the course of the experiments, and the effect of each of these sizes on the performance metrics was observed. The sources used were extremely bursty, with a peak-to-average ratio of 8 and a burst factor of 1000. The reassembly buffer size was the same as in Table 2.

In Figure 11 (A), we see the effect of variations in transmit buffer size on the throughput versus offered load curve. As expected, a larger number of buffers improves the gateway throughput significantly.

Figure 11 (B) plots the cell loss probability as a function of offered load for different transmit buffer sizes. The cutoff bandwidth strongly depends on the transmit buffer size. It is seen from the figure that the maximum transmit buffer size allowed by AMD's SUPERNET chip set is inadequate if we wish to support a high load that exhibits a high degree of burstiness. Also, note that since the other hosts on the FDDI network are generating a minimal amount of traffic, the gateway FDDI interface does not have to wait for a token. If the FDDI network is heavily loaded and the gateway is receiving bursty traffic, the observed loss can be even higher.

Figure 11 (C) shows how the latency versus offered load curve behaves when subjected to different transmit buffer sizes. Clearly, larger transmit buffer sizes result in greater delay through the gateway. With the increase in the offered load, the gateway latency increases exponentially. When the transmit buffer size is set to its maximum allowable value, the gateway latency is observed to be remarkably higher at larger values of the offered load. This increased delay is caused by the fact that there is more queueing at the transmit buffer. Therefore, for very bursty traffic, there is a tradeoff between cell/frame loss and gateway latency; this tradeoff can be made to tilt one way or another by adjusting the size of the transmit buffer memory.

4.5 Experiments with the FDDI to ATM Path

In Figure 3 showed a simplified queueing model for the reverse data path. It can be seen that the service mechanisms in this queueing system are simpler than their counter parts in the queueing model for the forward pipeline. If the receive buffer has sufficient space, the reverse data path should be able to handle peak 100 Mbps traffic from FDDI. We carried out an experiment (which has not been reported here), which proves that a receive buffer equal to 16 full size FDDI frames is sufficient for very bursty traffic on FDDI destined to an ATM host.

5 Conclusions

This paper presented a detailed simulation model for the design of a high performance ATM-FDDI gateway based on a new connection-oriented internetwork abstraction called VHSI. The hierarchical simulation model, constructed using a graphical simulation tool called BONES, served the dual purpose of functional verification and performance evaluation. The conclusions of this simulation study are discussed below.

Under common case packet processing conditions represented by small transmit buffer size and less bursty traffic generated by sources with burst factors less than 1000 and peak-to-average ratio of 2, a reassembly buffer with buffer space equivalent to two maximum sized FDDI frames per connection is sufficient to obtain acceptable cell loss in the range of 10^{-4}. With full size transmit buffers, the cell loss behavior is significantly improved and gateway can deliver almost ideal throughput.

The experiments that study the performance for the input traffic characteristics indicate that high burst factors (≥ 1000) and high peak to average ratio (≥ 8) for bursty applications can result in intolerable cell loss ($\approx 10^{-2}$) and frame loss ($\approx 10^{-1}$), if the transmit buffer is not adequate. Also the cutoff bandwidth above which the gateway drops cells and frames is very sensitive to the input traffic.

The experiments that study the effect of the transmit buffer size indicate that under very bursty traffic conditions, the maximum allowable transmit buffer size of 256 KB (68 full size FDDI frames) is not sufficient for obtaining zero cell loss in the gateway. At high loads and with adequately fast NPE, restricted transmit buffer capacity can cause significant cell loss and loss in throughput.

It has also be seen that the architecture of the gateway is constrained by restrictions imposed by the design of AMD's FDDI interface chip set. It requires that the access to the transmit buffer should be obtained through the intervening Node Processing Element (NPE), which should be sufficiently fast, otherwise if the gateway is assembling small sized frames at high loads, the NPE can become a bottleneck. This fact was conclusively proved by the experiments describing the effect of the node processing elements.

Also, the maximum achievable gateway throughput is restricted by the maximum rate at which requests of MPP to access the transmit buffer can be handled. Note that no resource management and enforcement functions were modeled and hence no control traffic was simulated. It was assumed that the control packet always receive the prioritized service needed from the NPE. In light of this fact engineering the NPE adequately is very important.

It was also found that due to simpler processing needs of the FDDI to ATM path, there are no bottlenecks in this path and the gateway can successfully carry 100 Mbps load from FDDI to ATM side.

This work also highlights the utility of a high level graphical simulation tool such as BONES in developing complex simulation models.

References

[1] Advanced Micro Devices Inc., *The SUPERNET Family for FDDI*, 1990.

[2] Cox, J., and Turner, J., "Project Zeus: Design of a Broadband Network and its Application on a University Campus," Technical Report WUCS 91-45 , Department of Computer Science, Washington University in St. Louis, Missouri, 1991.

[3] Escobar, J., and Partridge, C., "A Proposed Segmentation and Reassembly Protocol for Use with Asynchronous Transfer Mode (ATM)," *Protocols for High Speed Networks*, II, (M. Johnson, ed), pp. 353-368, 1991.

[4] Comdisco Systems, Inc., *Block Oriented Network Simulator (BONeS) Additional LAN Models*, 1991.

[5] Kapoor, S., Buddhikot, M., and Parulkar, G., "Design of an ATM-FDDI Gateway," *Journal of Internetworking: Research and Experience*, Vol. 4, No. 1, pp. 21-45, March 1993.

[6] Kapoor, S., *Design and Analysis of a Two Port ATM-FDDI Gateway,* MS thesis, Department of Electrical Engineering, Washington University in St.Louis, December 1991.

[7] Katz, D., "Transmission of IP and ARP over FDDI Networks," *Internet Engineering Task Force, Internet Draft*, 1992.

[8] Parulkar, Gurudatta M., "The Next Generation of Internetworking," *ACM SIGCOMM Communication Review*, Vol. 20, No. 1, pp. 18-43, Jan. 1990.