Impact of VC Merging on Buffer Requirements in ATM Networks

*A.L. Schmid**
Swiss Federal Institute of Technology, ETHZ
8092 Zurich, Switzerland
e-mail: Andreas.Schmid@switzerland.org
*Work performed at IBM Research Division, Zurich Research Laboratory.

I. Iliadis and P. Droz
IBM Research Division, Zurich Research Laboratory
8803 Rüschlikon, Switzerland
Tel. +41-1-724-86-46, Fax +41-1-724-89-55
e-mail: ili,dro@zurich.ibm.com

Abstract

For the implementation of multipoint-to-point connections in ATM, various approaches exist, each with its own advantages and disadvantages. VP-based methods require unique sender identification but they do not require reassembly in merging points. In contrast, VC-based methods do not require unique sender identification but they do require reassembly in merging points. It is likely that VC merging will be the method of choice as it is scalable and yet relatively simple to implement. One of its drawbacks is the increased output buffer space required at the switches because of packet reassembly at the merging points. This paper investigates the impact of the switch architecture and characteristics on the output buffer space by means of simulation. The results obtained demonstrate that for typical switch architectures, VC merging does not require significant additional buffering compared to VP merging.

Keywords

IP over ATM, VC merging, VP merging

1 INTRODUCTION

In current ATM networks, there exist only *point-to-point* (pt-to-pt) and *point-to-multipoint* (pt-to-mpt) connections. For the interconnection of routers across an ATM network as well as for many other information-gathering applications, *multipoint-to-point* (mpt-to-pt) connections appear to be more appropriate. Interconnection of N routers requires order N^2 labels for the order of N^2

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35388-3_42

H. R. van As (ed.), *High Performance Networking*

pt-to-pt connections as described by Calvinac *et al.* (1997). With mpt-to-pt connections only N labels for the N associated connections are necessary. This significantly reduces the required label space and thus makes the method more scalable. The same new ATM connection type could also be used in the context of merged connections for MPLS (Callon *et al.* 1997).

To implement mpt-to-pt connections, different solutions are possible. We focus on the two most important methods: *VP merging* and *VC merging*.

VP merging: Each sender is assigned a globally unique identifier having the format of a VCI. The identifier is carried in the VCI field of the ATM cell. The ATM switch translates incoming VPIs for the same destination to the same outgoing VPI. The receiver distinguishes amongst the different sources by the different VCIs. The key advantage of this scheme is that no VCC resources are required in the switching nodes as only VP switching is performed. This implies no change of hardware but only a change of the connection establishment protocol. Some of the disadvantages of VP merging are the lack of scalability caused by the VPI address space limitation of 4096 entries and the need for a "global VCI uniqueness" protocol. There are proposals to circumvent the nonscalability by enlarging the VPI address space at the expense of VCI address space. This is not desirable, however, because it requires changes in the switching hardware. *VC merging*: This method avoids the requirement for globally unique sender identifiers, and it consumes only one VCI per traversed link. These characteristics make this approach scalable. Each source participating in a mpt-to-pt connection has a unique VCI per link. The ATM switch translates incoming VCIs belonging to the same connection to a single outgoing VCI. This means that cells of packets belonging to different senders could be interleaved. As the receiver is not able to distinguish cells from different senders, packet reassembly has to be performed at the merging points, and all cells from a given packet must be sent contiguously so that reassembly at subsequent merging points and at the receiver will be possible. AAL 3/4 would solve the problem by introducing the Message Identifier (MID) field for sender identification in every cell. The use of AAL 3/4, however, has other drawbacks such as the limited space of the MID field, the inefficient encapsulation method, and the less powerful CRC capability. In this paper we consider the employment of AAL 5 because it is widely available and supported in ATM switches, especially in data networks. Packet reassembly at the merging points introduces additional buffer requirements on the switching architecture because all of the cells of a packet sent by a sender belonging to a mpt-to-pt connection have to be stored and must wait for the last cell of the packet identified by the "End Of Packet" (EOP) marker used by AAL 5 to arrive. Figure 1 depicts the cell interleaving problem. Packet reassembly also introduces additional delay for packets transported over a merged connection and adds burstiness to the traffic. This is because all the cells of a packet have to wait at every merging point. They appear afterwards as a burst of a whole packet at the output link. This burstiness becomes even worse

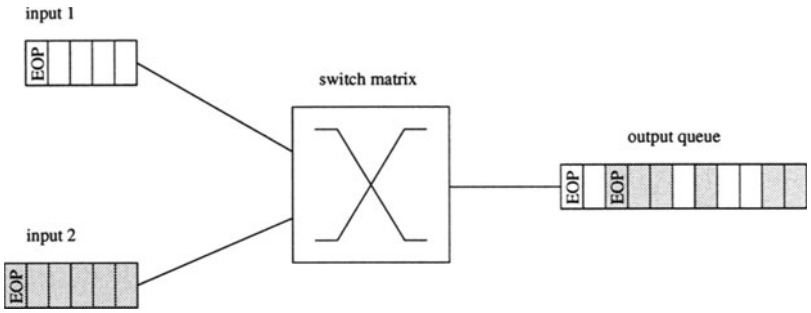

Figure 1 Cell interleaving problem.

as it is often cascaded and thus accumulated over numerous merging points. Heinanen (1997) gives some hints about how to solve the problems involved in mpt-to-pt VC merging.

A third possibility is to handle a mpt-to-pt connection of N senders to one receiver like N pt-to-pt connections without applying any merging. This possibility again requires order N^2 labels for the order of N^2 pt-to-pt connections. Of the above possible solutions, VC merging appears to be the method of choice as it is relatively easy to implement and yet scalable. At the ATM Forum, VC merging has been almost fully accepted and will most likely be introduced in the PNNI v2.0 specification (expected to be finished in the spring of 1998). The only concern is with the reassembly required in the switches in terms of additional buffering and delay. The numerous simulations presented in the following sections are used to investigate the required additional buffer overhead for VC merging. It is also very likely that different methods of merging and nonmerging will exist simultaneously in an ATM network. Some interworking aspects of these methods are discussed by Widjaja *et al.* (1997).

Section 2 of this paper describes our switching architecture model for VC merging and the model of the arriving traffic. In Section 3 we show our simulation setup and discuss the results of the simulations. In Section 4 we give a summary and derive some conclusions.

2 SWITCH AND TRAFFIC MODEL

2.1 Switch Model

In this paper we consider the general class of single-stage, nonblocking $M \times M$ packet switches with both input and output queuing (Iliadis and Denzel 1993, Denzel *et al.* 1995). The shared output buffer is assumed to be sufficiently large so that the switch performance is close to optimal, corresponding to the pure output queuing. Cells are transferred from the head of the input queues to the shared buffer. The speed of the input and output switch ports is denoted R_S,

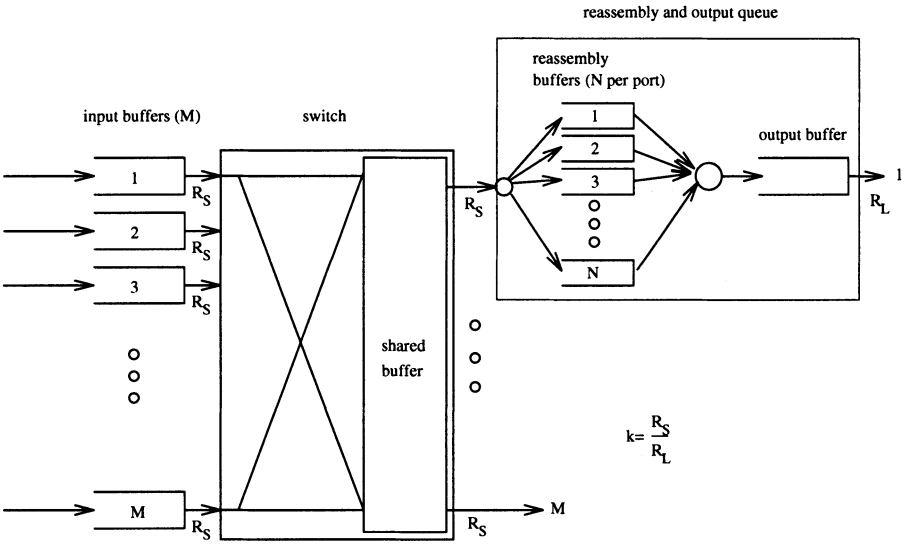

Figure 2 Switch architecture.

and the speed of the outgoing links is denoted R_L. Let k denote the speed ratio of the switch speed (per port) to the outgoing link speed, i.e. $k = R_S/R_L$. Typically k is greater than one, which implies that an output queue should be provided in order to cope with the speed mismatch.

As described above, VC merging requires an amount of additional output buffering due to the packet reassembly. We introduce a so-called reassembly buffer at each output port of the ATM switch. Figure 2 shows the concept of the reassembly buffer. A switch has M input ports and N sources of mpt-to-pt connections because it is likely that different connections will coexist. Hence N can be much larger than M. The model considered in this paper is valid for the case of $N \leq M$. The case of $N > M$ is not covered by the present switch model and is therefore a subject for further investigation. At every merging point, each of the sources participating in the corresponding mpt-to-pt connection is associated with a distinct reassembly buffer at the output queue. When the last cell of a packet with the EOP marker arrives at the reassembly buffer, all of the cells of a packet are instantly transferred into a single output buffer per output port. Physically the reassembly and the output buffers of one output port share a common memory pool. The transfer from the reassembly to the output buffer can easily be done by a pointer movement and will therefore not incur additional delay.

The simulation models for VC and VP merging are shown in Figures 3 and 4, respectively. Cells belonging to the various VCs are transferred from the head of the switch input queues in the shared buffer and, subsequently, to the corresponding output queues. It is assumed that the traffic is uniform, i.e. the destination of an arbitrary packet can, with an equal probability, be any

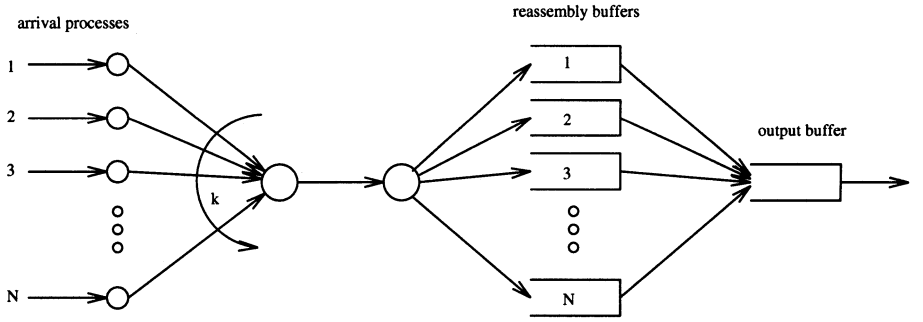

Figure 3 Our simulation model for VC merging.

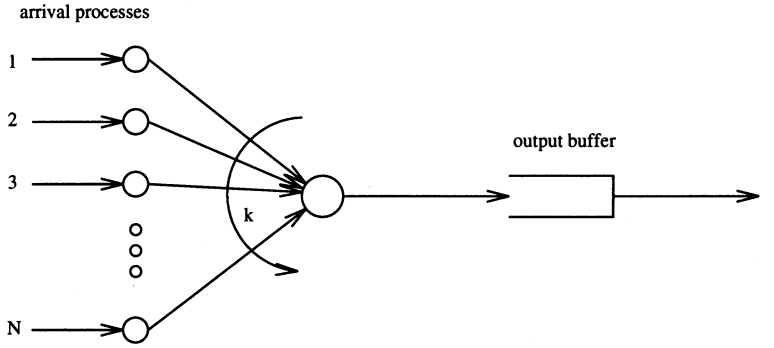

Figure 4 Our simulation model for VP merging.

of the output ports, and that successive packets are independent regarding their output port destinations. Owing to traffic symmetry, all of the output queues have identical behavior. Let us turn our attention to a particular output queue and study its behavior. The corresponding simulation model considers N sources feeding the output queue in a round-robin fashion governed by the factor k. This model is also appropriate for the case where the switch fabric is capable of transferring only a limited number of cells to any given output (Oie *et al.* 1989).

2.2 Traffic Model

The traffic and simulation model we use is shown in Figures 3 and 4. We use N arrival processes, which correspond to the traffic destined to the output queue. Packets are assumed to arrive according to either a Poisson process (nonbursty traffic with the mean arrival rate λ) or a hyperexponential process (bursty traffic). The hyperexponential process is generated by a two-stage hyperexponential distribution. The mean values corresponding to the two stages

are $0.51 * \lambda$ and $16.48 * \lambda$, respectively. The corresponding routing probabilities for the two stages are 0.97 and 0.03, respectively, so that the mean arrival rate is again equal to λ. Each packet is assumed to contain a number of cells geometrically distributed with a mean of E cells (Chao and Smith 1992, Widjaja and Elwalid 1997). We used $E = 10$, $E = 30$, or $E = 180$ cells (10 cells correspond to 472 bytes, 30 cells to 1432 bytes and 180 cells to 8632 bytes).

It is shown by Widjaja and Elwalid (1997) that the mean packet size in a core network where ATM is likely to be applied is about 289 bytes. This yields 6.2 ATM cells of data using AAL 5 with the null encapsulation method as described by Heinanen (1993) (additional overhead of AAL 5 is 8 bytes per packet). The dominant packet sizes in an Internet backbone are 40 or 44 bytes at about 36% of the traffic (TCP acknowledgment packets, TCP control segments such as SYN, FIN, ..., and Telnet packets carrying single characters), 552 or 576 bytes at about 25% (512 and 536 bytes of TCP implementations without path MTU discovery as the default maximum segment size (MSS) for nonlocal IP destinations, yielding a 552 or 576-byte packet size), 185 bytes at about 2.7%, and 1500 bytes at about 1.5% (Ethernet traffic). These statistics were collected on Feb 10, 1996, in FIX-West network as a sample wide-area network, and are given on the NLANR homepage (1996). A more recent study of traffic characteristics in an Internet backbone was conducted in August of 1997 (Thompson *et al.* 1997). It is shown that almost 50% of the traffic is 40 or 44 bytes in packet length. More prominent packet sizes are 532, 576, and 1500 bytes, each representing 15% of the traffic. Comparing the two studies we observe a shift to smaller packets of size 40 or 44 bytes and larger packets of size 1500 bytes.

For the future development of packet sizes, the spreading of the use of path MTU (PMTU) discovery will have a significant impact. PMTU will affect MTUs in IPv4 as proposed by Mogul and Deering (1990) and even more MTUs in IPv6 over faster LANs. There will be numerous packets with possible sizes up to 64 kilobytes (max. packet format for AAL 5 is 64 kilobytes (Laubach 1994)). A single packet of this size involved in reassembly could alone fill the entire reassembly buffer in a switch output queue. Atkinson (1994) gives an overview of other typical frame sizes being applied on AAL 5. These are 8 kilobytes used by the Network File System (NFS) and the 9180 bytes of IP MTU over SMDS (Piscitello and Lawrence 1991) that became the default value for IP MTU over ATM AAL 5 (Laubach 1994). These big packet sizes in conjunction with VC merging could induce present problems that VP merging would not encounter. On the other hand there will also be much more real-time traffic (e.g. voice) in the Internet. Real-time traffic typically produces a large amount of very small packets.

3 SIMULATION SETUP AND RESULTS

This study concentrates on the additional buffer space required for reassembly. It is conducted under loads $l = 30\%, 70\%, 90\%$, with different traffic characteristics (bursty, nonbursty), factors $k = [1, .., 16]$, $N = 16, 64, 128$ sources, M ports ($M \geq N$), and with a mean packet size of $E = 10, 30, 180$ cells. The default values for the simulations are $N = 16$ sources, $E = 10$ cells, $l = 90\%$, and $k = 16$ unless specified otherwise. Figures 5-13 show the VC merging buffer size (solid line) and the corresponding VP merging buffer size (dashed line). The results serve to compare VP and VC merging. They cannot, however, be used directly to show the required output buffer space in an ATM switch because no flow control has been taken into account. The simulations were carried out for an extremely large number of events such that 95% confidence intervals were very small.

Figure 5 shows the results for $l = 30\%, 70\%, 90\%$. The difference between the solid and the dashed lines (VC and VP merging for a specific load) is about 19 to 21 cells for $l = 90\%$, about 21 to 23 cells for $l = 70\%$, and about 30 to 37 cells for $l = 30\%$ over some magnitudes of overflow probability. At high loads the output queue contains a large number of cells, which translates to long delays. Therefore, by the time the first cell of a packet is ready for transmission at the output link, the corresponding last cell has most likely arrived and, consequently, the packet reassembly has been completed. In this case, therefore, the additional overhead due to reassembly is almost negligible. It is also important to note that the workload of today's switches normally lies at high levels of around 70% or 80%. In contrast, at low loads, the first cell may be ready for transmission while the reassembly is in progress. In this case it has to be delayed until the reassembly process has been completed. How-

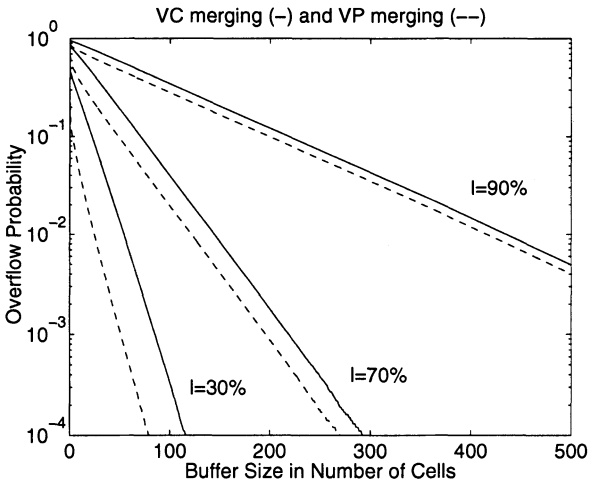

Figure 5 Simulations for $l = 30\%$, 70%, 90%, nonbursty arrival process.

ever, owing to the low load, the number of packets under reassembly is small and, therefore, the additional buffer requirement of VC merging is minimal. The results obtained are in agreement with those presented by Widjaja and Elwalid (1997). Furthermore, the packet delay corresponding to VC merging approaches that corresponding to VP merging as the load increases.

We then made the same simulations with bursty arrival processes. We model the bursty arrival process by a hyperexponential packet arrival process as described in the previous section. The results for $l = 30\%, 70\%, 90\%$ are shown in Figure 6. We see that the buffer requirements for both VC and VP merging grow significantly for high loads. Of course flow control would alleviate this problem to some extent due to the overall load reduction. The additional buffer requirements for VC merging compared to VP merging are minimal even for the case of bursty traffic. In particular, for high loads they become negligible for the reasons given above.

Simulation results were obtained for different values of k and different loads l. By varying k we expected to see an influence on the additional buffer requirement. Surprisingly, only the extreme value $k = 1$ resulted in a big additional buffer requirement for VC merging. It is obvious that VP merging requires almost no output buffer with $k = 1$ as the speed of the switch output port (R_S) is equal to the speed of the output link (R_L). We then tried to determine the critical k for every load factor l considered. The critical value of k is defined as follows: For all values of k larger than the critical value, there is practically no distinction between VC curves and VP curves, whereas for all smaller values of k the curves start becoming distinguishable. We found that the critical k lies close to the extreme value $k = 1$. The range of the critical k is between 1.1 and 1.3 for $l = 90\%$ and $l = 70\%$, respectively. This means that the critical k becomes larger with lower loads, but it is still far away from

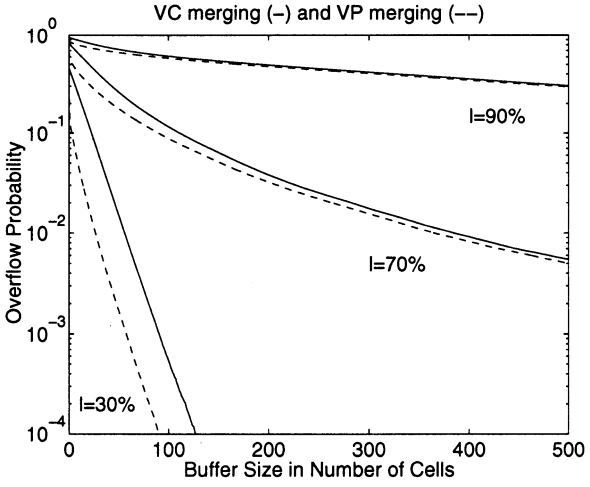

Figure 6 Simulations for $l = 30\%, 70\%, 90\%$, bursty arrival process.

the values implemented in today's switches (greater than 2). To substantiate these observations we investigated the critical k for $l = 30\%$, too. In this case the critical value for k is approximately 1.5, which is still much smaller than 2 and thus confirms our theory.

Figure 7 shows the results of our simulations for $k = 1.1, 1.2, 16$ at $l = 90\%$. We observe that all of the curves for the output buffer of VC merging at different values of k lie close together. The value of $k = 1.1$ is the critical one because the corresponding curve starts to show a deviation. The same applies to the curves for VP merging. For values of k greater than the critical one, the additional buffer requirement for VC merging at low overflow probabilities is minimal. However, the difference between VC and VP merging becomes noticeable for values of k less than the critical value.

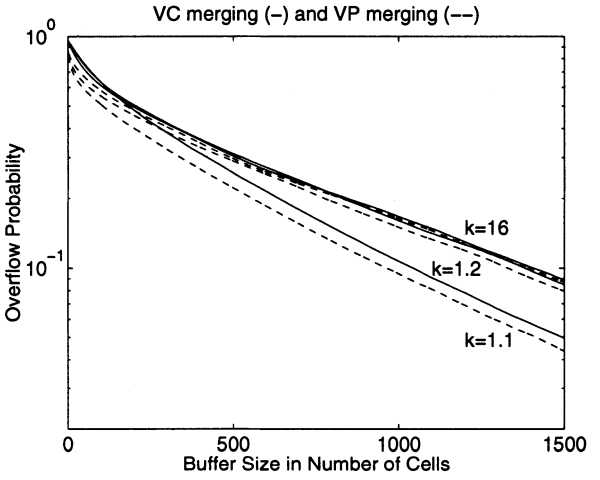

Figure 7 Simulations for $l = 90\%$, $k = 1.1$, 1.2, 16, bursty arrival process.

Figure 8 shows the results of our simulations for $k = 1.2, 1.3, 16$ at a lower load of $l = 70\%$. In this case, we observe that the different curves for VC and VP merging lie fairly close together with the exception of the curves for $k = 1.2$. This shows that the critical k is slightly larger for $l = 70\%$ (about $k = 1.2$) than for $l = 90\%$ (about $k = 1.1$). Here again, the additional buffer requirements for VC merging at low overflow probabilities become noticeable for values of k less than the critical value.

Figure 9 shows the results of our simulations for $k = 1.5, 2, 16$ at a low load of $l = 30\%$. We observe again the similarity of the curves for VC merging over the entire range of k. The curves for VP merging vary slightly so that the additional buffer space becomes smaller for a larger k, with a critical k at about $k = 1.5$. There is a noticeable additional buffer requirement for

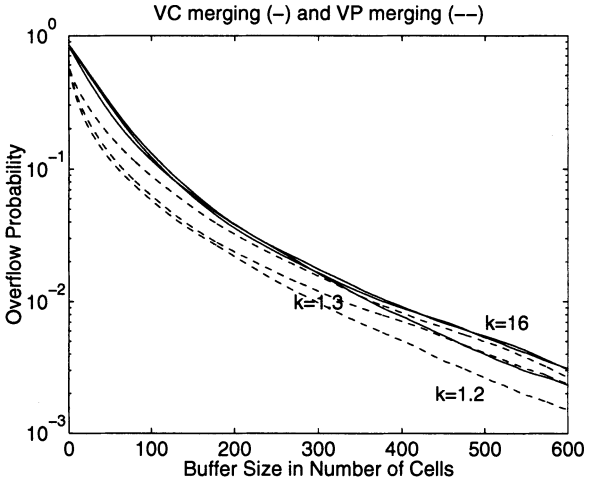

Figure 8 Simulations with $l = 70\%$, $k = 1.2$, 1.3, 16, bursty arrival process.

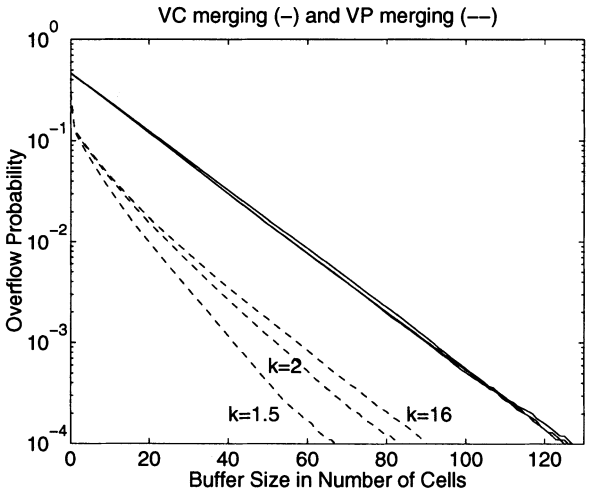

Figure 9 Simulations for $l = 30\%$, $k = 1.5$, 2, 16, bursty arrival process.

VC merging in the entire range of values of k. Furthermore, the additional requirement increases as k decreases.

We then tried to investigate the possible influence of more specific traffic characteristics such as larger packet sizes and increased numbers of sources in a mpt-to-pt connection on the additional requirements of VC merging compared to VP merging. First, we performed simulations for a larger mean packet size of the arrival process ($E = 30$). Figure 10 shows the curves for $l = 90\%$ and $k = 1.1, 1.2, 16$ with a mean packet size of $E = 30$. Compared to Figure 7 we observe a greater difference between the curves for $k = 1.1$ and for $k = 1.2$.

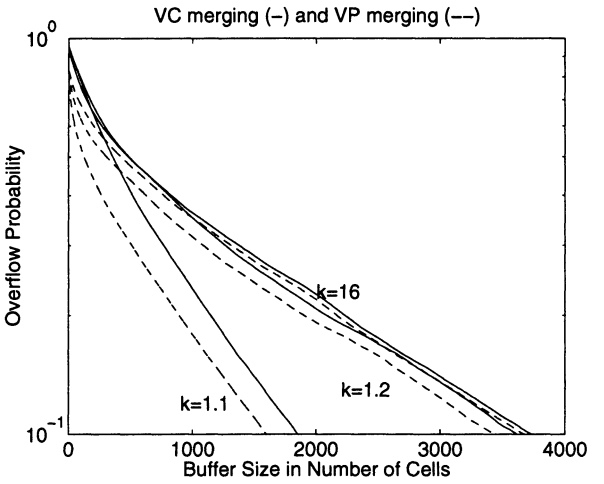

Figure 10 Simulations for $l = 90\%$, $k = 1.1$, 1.2, 16, bursty arrival process and $E = 30$.

It appears that the critical k is shifted to a value slightly larger than $k = 1.1$ (between $k = 1.1$ and $k = 1.2$). Furthermore we see that the mean packet size, which is three times larger, requires an output buffer size that is also three times larger. Moreover, the additional output buffers for VC merging are about three times larger for $E = 30$. Therefore the additional buffer requirement for VC merging appears to grow linearly with the mean packet size. This trend is also verified by our simulations for $E = 180$.

Figure 11 shows the results of the same simulation for $l = 70\%$, $k =$

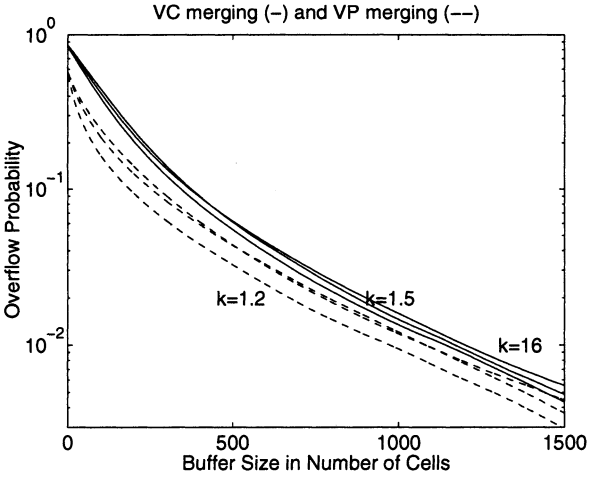

Figure 11 Simulations for $l = 70\%$, $k = 1.2$, 1.5, 16, bursty arrival process and $E = 30$.

1.2, 1.5, 16 and an increased mean packet size of $E = 30$. Compared to Figure 8 we again observe a shift of the critical k from a value of about $k = 1.2$ to a slightly larger value. Concerning the additional buffer requirement for VC merging, the same observations were made as in the case of load $l = 90\%$. This means that, also at this load, as the packet size increases, the additional buffer requirement increases by the same factor.

Finally we performed simulations for a larger N ($N = 64, 128$) to assess the influence of a large number of sources associated with one mpt-to-pt connection on the additional buffer requirements of VC merging due to reassembly. An increased number of sources could translate to an increased degree of reassembly. This again would lead to a significantly larger required buffer space for reassembly than for nonreassembly. Figure 12 shows the results for the simulations for $N = 16, 64$ sources, factor $k = 16$ and nonbursty traffic. The results obtained also apply in the case of nonbursty traffic. This is explained by Palm–Khintchine's theorem (Heyman and Sobel 1982, p. 156), which states that summing up a large number of iid processes (for instance hyperexponential processes as used for our bursty traffic) results in a process of Poisson type (our nonbursty traffic). As our simulation has 64 sources, each with an iid process for the arrival traffic, we are able to apply this theorem and to simulate nonbursty arrival traffic. All of the corresponding curves in Figure 12 lie close together. Concerning VP merging, as N increases, the corresponding curves converge because the aggregated arrival process tends to a Poisson one. For VC merging, the buffer requirement does not increase with the number of sources. This is because increasing the number of sources translates to decreasing the arrival rate per source such that the load at the output link

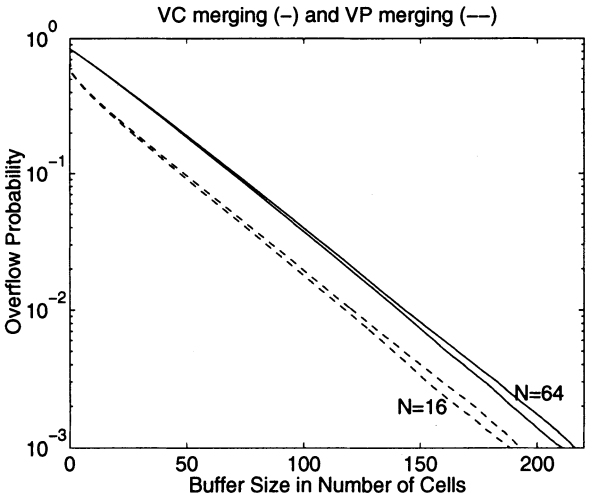

Figure 12 Simulations for $N = 16, 64$, $l = 70\%$, $k = 16$, nonbursty arrival process.

remains constant. This shows that our previous simulation results hold also for a larger scenario with a larger number of senders.

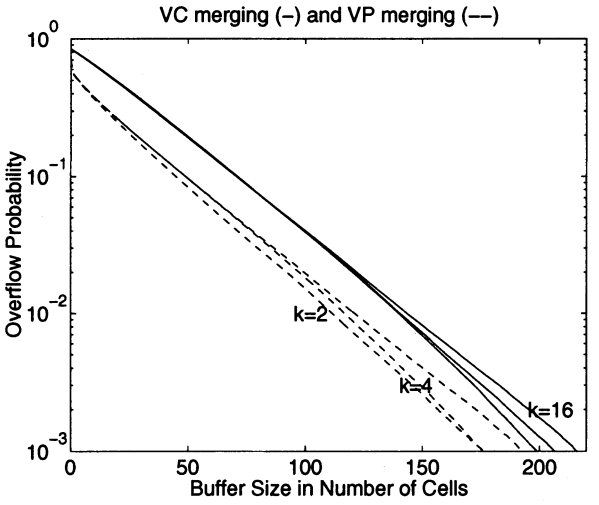

Figure 13 Simulations for $N = 64$, $l = 70\%$, $k = 2$, 4, 16, nonbursty arrival process.

We have investigated the impact of varying k given large values of N. Previously we found a critical k of about 1.1 to 1.3 at $N = 16$. Figure 13 shows the results of the simulations with $N = 64$ and $k = 2, 4, 16$ with nonbursty traffic. The different curves for VC and VP merging again lie close together and we see no significant difference between the curves belonging to the values $k = 2$ and $k = 16$. Consequently the critical value of k is smaller than 2. Once again, for values of k greater than the critical value, the additional buffer requirement for VC merging does not increase.

4 SUMMARY AND CONCLUSIONS

VC merging is likely to become the method of choice to implement mpt-to-pt connections in ATM networks. Because of the cell interleaving problem created by VC merging, reassembly has to be performed in the merging points. The effect of reassembly has been investigated assuming an output queue switch architecture. The results obtained demonstrate that, at high loads and for arbitrary arrival processes, the implementation of VC merging in the switches will not require much additional buffer at the output queues of the switches. In contrast, at low loads, additional buffer is required but this is minimal. Furthermore, it was found that the additional buffer requirement for VC merging is proportional to the average packet size. Consequently, large

packet sizes can result in large reassembly buffer requirements. We further investigated the effect of the speed ratio between switch output port and output link and came to the conclusion that for sufficiently large speed ratio values ($k > 2$) the output buffer requirement for VP and VC merging remain the same, respectively. We found a critical k which grows with decreasing utilization and also with growing mean packet sizes of the arrival traffic. But it always remains between 1.1 and 1.3 for high utilization of 70% and 90%.

REFERENCES

Atkinson, R. (1994) Default IP MTU for use over ATM AAL5. RFC 1626, May 1994.

Callon, R., Doolan, P., Feldman, N., Fredette, A., Swallow, G. and Viswanathan, A. (1997) A framework for multiprotocol label switching.
Internet Draft <draft-ietf-mpls-framework-02.txt> (November 1997).

Calvignac, J., Basso, C., Droz, P. and Dykeman, D. (1997) Dynamic IDentifier Assignment (DIDA) for merged ATM connections. ATM-Forum / 97-0504 (July 1997).

Chao, H.J. and Smith, D.E. (1992) A shared-memory virtual channel queue for ATM broadband terminal adaptors. *Int'l J. Digital and Analog Commun. Systems*, **5**, 29–37.

Denzel, W.E., Engbersen, A.P.J. and Iliadis, I. (1995) A flexible shared-buffer switch for ATM at Gb/s rates. *Comp. Networks & ISDN Systems*, **27**, 611–24.

Heinanen, J. (1993) Multiprotocol encapsulation over ATM adaption layer 5. RFC 1483, July 1993.

Heinanen, J. (1997) Multipoint-to-point VCs. ATM-Forum / 97-0261 (April–May 1997).

Heyman D.P. and Sobel, M.J. (1982) *Stochastic Models in Operations Research.* Vol. 1. McGraw-Hill, New York.

Iliadis I. and Denzel, W.E. (1993) Analysis of packet switches with input and output queuing. *IEEE Trans. Commun.*, **41**, 731–40.

Laubach, M. (1994) Classical IP and ARP over ATM. RFC 1577, January 1994.

Mogul J. and Deering, S. (1990) Path MTU discovery. RFC 1191, November 1990.

National Laboratory for Applied Network Research (NLANR)
http://www.nlanr.net/NA/Learn/packetsizes.html

Oie, Y., Murata, M., Kubota, K. and Miyahara, H. (1989) Effect of Speedup in Nonblocking Packet Switch, in: *Proc. ICC'89*, Boston, MA, pp. 410–4.

Piscitello D. and Lawrence, J. (1991) The transmission of IP datagrams over the SMDS service. RFC 1209, March 1991.

Thompson, K., Miller, G.J. and Wilder, R. (1997) Wide-area internet traf-

fic patterns and characteristics. *IEEE Network*, November/December issue.

Widjaja I. and Elwalid, A.I. (1997) Performance issues in VC-merge capable switches for IP over ATM. ATM-Forum / 97-0675 (July 1997).

Widjaja, I., Wright, S. and Chatterjee, A. (1997) Interworking of VP-merge, VC-merge, and non-merge ATM switches in a multipoint-to-point environment. ATM-Forum / 97-0748 (September 1997).

BIOGRAPHIES

Andreas L. Schmid received an M.S. degree in Electrical Engineering in April 1998 from the Swiss Federal Institute of Technology (ETH) in Zurich, Switzerland. From October 1996 to March 1997, he was affiliated with the University of Stuttgart, Germany, where he participated in the ERASMUS student exchange program. From November 1997 to March 1998, he was affiliated with the IBM Zurich Research Laboratory, where he did his Master's thesis.

Ilias Iliadis received a B.S. degree in Electrical Engineering in 1983 from the National Technical University of Athens, Greece, an M.S. degree in 1984 from Columbia University, New York, as a Fulbright Scholar, and a Ph.D. degree in Electrical Engineering in 1988, also from Columbia University. From 1986 to 1988, he was affiliated with the IBM Thomas J. Watson Research Center in Yorktown Heights, NY, as a work-study student. In 1988, he joined the IBM Zurich Research Laboratory as a member of the switching systems group working on broadband switching. He was responsible for the performance analysis and design of the IBM's PRIZMA switch chip. Currently, he works in the field of ATM-based customer premises networks. His research interests include analysis of distributed systems, performance evaluation of computer communication networks, switching architectures, development of protocols and congestion control schemes, and optimization and network design algorithms. Ilias Iliadis is a member of Sigma Xi, IEEE, and the Technical Chamber of Greece.

Patrick Droz studied computer science at the Swiss Federal Institute of Technology (ETH) in Zurich. He received an M.S. in 1992. He then joined the ATM Networking Group at the IBM Zurich Research Laboratory in Rüschlikon, Switzerland, as a PhD student. During this time, he also worked on the design and implementation of the ATM control point for the 8260 campus backbone hub. He received a Ph.D. in 1996 for his dissertation entitled "Traffic Estimation and Resource Allocation in ATM Networks." Since then, he has been working in the ATM Networking group as a research staff member. He is currently working in the area of IP/ATM integration, which involves participation in IETF as well as the ATM Forum.