12

Teleservice Provision via a Frame Relay Network - A User Case Study

R. Hunt
Senior Lecturer (Communications)
Department of Computer Science
University of Canterbury, New Zealand
E-mail ray@cosc.canterbury.ac.nz
Tel +64-3-3642347 Fax +64-3-3642999
WWW homepage http://www.cosc.canterbury.ac.nz/~ray

Abstract

Frame Relay is a modern teleservicing wide area architecture which is designed to complement a variety of existing services and in many cases it offers enhanced services at a lower cost that some of the traditional teleservices. Frame Relay has a particularly important role to play in providing elastic bandwidth communication in support of distributed teleservicing applications, and modern client/server LAN-to-LAN systems in particular. It also forms an important step in the direction to full multimedia ATM networking. The bandwidth requirements for emerging teleservice applications such as client/server based transaction processing, image and graphics transmission as well as distributed database systems are very different from those used in earlier applications. These modern applications all process large volumes of data, transmit intermittent high speed bursts and are intolerant of long delays. To satisfy these requirements, the bandwidth management system in LAN-to-WAN and LAN-to-LAN networking must offer access to high bandwidth on demand, direct connectivity to all other points in the network and consumption only of bandwidth actually needed. This paper examines users' experience resulting from three years use of a Frame Relay teleservicing network in New Zealand and discusses a range of key performance, operational and design issues.

Keywords

Teleservice architecture, Frame Relay networking, application trials, intelligent network, virtual circuit, committed information rate, elastic bandwidth, forward/backward congestion notification, open/closed loop congestion control

F. A. Aagesen et al. (eds.), *Information Network and Data Communication*

1 INTRODUCTION

Tuianet is a New Zealand Frame Relay virtual teleservice network designed to promote networking for research, education and information sharing. It comprises three separate physical networks. The first, CRINet (Crown Research Institute Network) provides networking facilities for the Government's ten Crown Research Institutes and consists of seven nodes. The second, AgNet is an Agricultural Research and Development Network and consists of five nodes. Both of these networks support the new networking requirements resulting from Government department restructuring in recent years. The third network provides interconnection for the eight main universities. Additional connections are also provided to information service providers such as the National Bibliographical System, the Antarctic Centre, and to the Internet.

Each of these networks are managed by the respective group according to their own criteria for management, charging, etc. However, through Tuianet the three individual groups agree on issues of performance, security, interfacing as well as matters relating to day-to-day operation. This paper, commences with a background to Frame Relay teleservice provision and is followed by an analysis of the implementation experience and resulting performance characteristics.

1.1 What is Frame Relay?

Frame Relay is a wide area networking solution designed to provide flexible, high performance interconnection for regional, national and international networks. It has characteristics which make it an excellent choice for interconnecting distributed LANs.

Frame Relay is a good choice for handling *bursty* high data-rate traffic commonly generated by LAN applications (Lamont, 1989). In the past the predominant technologies used to solve LAN-to-LAN interconnection have been:

- Circuit Switching by way of Time Division Multiplexing (TDM)
- X.25 packet switching.

Providing Frame Relay is used in conjunction with an appropriate teleservice application environment, it offers a solution which is frequently superior to either of these other two technologies. However it is important to note that there are certain traffic characteristics such as high volume continuous data, combined voice and conventional data, or low volume message traffic which would still make either of the above two options applicable in certain circumstances. However, in general, Frame Relay is more flexible and cost-effective than TDM and provides better throughput than X.25 packet switching.

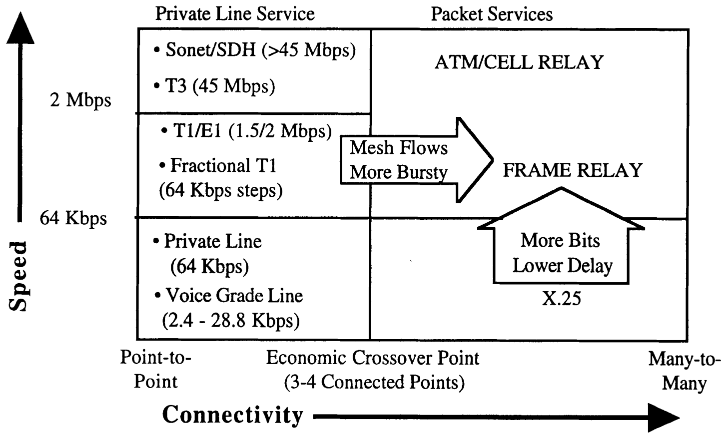

Figure 1 Speed/connectivity relationship for leased line and packet based teleservices (AT&T, 1992)

While Frame Relay provides performance improvement for X.25 users, it is unlikely to replace it as a teleservice network for a number of existing applications. Thus X.25 will retain a significant place as a networking teleservice. Applications which require high data rates or low delay benefit from Frame Relay as can be seen in Figure 1. Frame Relay networks support access speeds up to E1 (2.048 Mbps) in comparison with a maximum of 64 Kbps for most X.25 packet switched networks. Higher data rates of 34 Mbps (E3) and 45 Mbps (T3) have been demonstrated by some vendors and Frame Relay networks also offer lower delays than those found in conventional packet based networks.

The designers of Frame Relay recognised the need for higher performance packet based networks and sought to avoid the duplication in the functionality between X.25 networks and end-system higher layer protocols. In addition they took account of the low error rate associated with modern digital telecommunications systems. The designers therefore created a protocol that permitted switches to be built with speed and not error-free transmission as the primary objective. The resulting architecture can be seen in Figure 2.

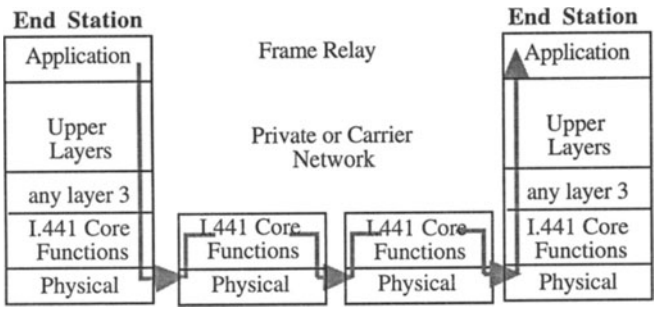

Figure 2 Architectural stack of a Frame Relay Network (Minoli, 1995).

1.2 Frame Relay in the spectrum of alternative teleservices

Frame Relay has been acclaimed as an ideal for LAN-to-LAN interconnection teleservice architecture with resulting cost savings over the use of private leased DDS (Digital Data Service) teleservices, particularly as a consequence of its elastic bandwidth. In fact over 80% of Frame Relay's use is directed at LAN-to-LAN communication with the remainder focusing on host-to-terminal and host-to-host communication.

In some ways Frame Relay is of more significance to the network provider than the user since 64 Kbps - 2 Mbps point-to-point links have been widely available for many years. However, the provision of such bandwidth on a continuous basis is uneconomical for many applications and a more flexible or elastic ("rubber") bandwidth system was needed.

There are many LAN-to-LAN connections that have used DDS point-to-point links most satisfactorily for many years. At the lower end of the spectrum, 64 Kbps links have been used by some companies where reasonable data transfer volumes have justified the cost. Alternatively X.25 packet switching links at a similar speed often turned out to be expensive because of the significant connection and volume charges.

This is not to say that Frame Relay is superseding DDS or packet switching teleservices, but rather that a new option has emerged which for some applications falls neatly between the other options.

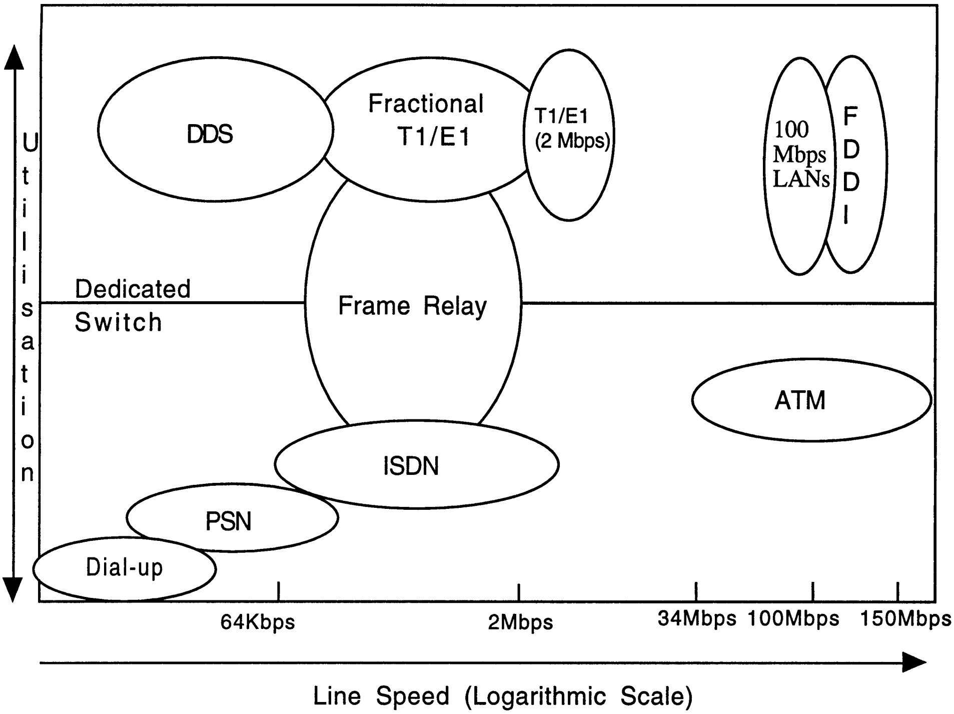

Figure 3 Positioning of telecommunication interconnection teleservices.

Figure 3 shows where Frame Relay lies in the bandwidth utilization/line speed spectrum covering both switched and dedicated teleservices. (Note that the line speed axis is a logarithmic one). It can be seen that overlap with other services does exist and in some situations it may be difficult to choose between some alternatives.

It is generally accepted that DDS (including the various T1/E1 teleservices) is suitable for dedicated interconnection between sites but at a lower performance level. ISDN is appropriate for non-continuous services where moderate bandwidth is required for short holding times. However ISDN may not offer adequate bandwidth for certain teleservice applications. Frame Relay is appropriate for major sites and networks requiring bursty traffic flow between interconnected LAN applications. It is relatively cheap to implement on existing hardware - sometimes it merely requires a software upgrade to the existing X.21 interface on a router.

2 VIRTUAL CIRCUITS AND COMMITTED INFORMATION RATE

The Frame Relay standard specifies use of either PVCs or SVCs. Virtual circuits in the Frame Relay context are used for connectivity and addressing and are able to provide the equivalent of a private leased circuit but with substantially increased flexibility. All traffic for a particular virtual circuit uses the same path through the network. However, virtual circuits consume no resources when they are not active, but can (in principle) instantaneously carry up to 2 Mbps of traffic. Many PVCs can be established at a single access point between the end-station (usually a router) and the network. Each of these PVCs can link different end-stations and have different throughput capacity.

The protocols for dynamically setting up SVCs have only recently been finalised and are more complicated than those used for PVCs although SVCs will inevitably be required for international Frame Relay services. The PVCs in a Frame Relay network allow any single device to communicate with hundreds of other devices by way of a single access port.

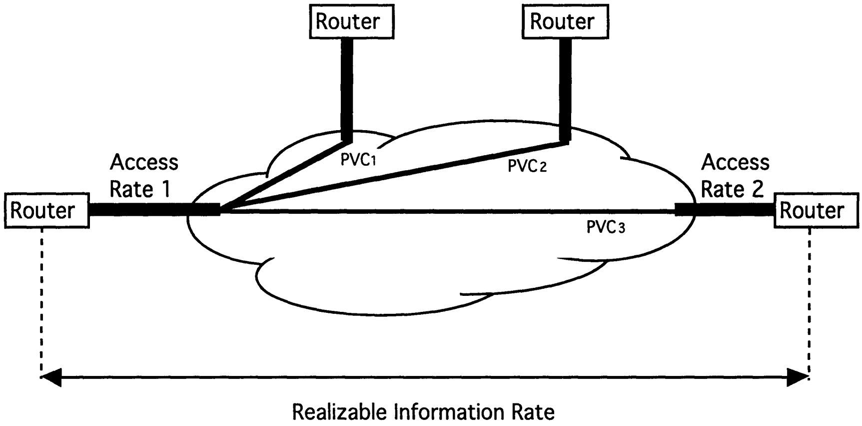

Figure 4 Frame Relay Interconnection Using PVCs.

Once frames enter the Frame Relay network they are forwarded along the PVC according to the connection identifier specified by the Data Link Connection Identifier (DLCI) in the frame. Data transmitted on a particular PVC is statistically multiplexed onto network trunks along with data from other users. Since this statistical multiplexing process involves variable store-and-forward delays, the throughput for individual transactions will be variable and usually less than the access rate. (See Figure 4).

Nodes in a Frame Relay network are linked by PVCs each with a Committed Information Rate (CIR) which defines the bandwidth guaranteed to be available to the connection although the maximum amount of bandwidth can be much higher if network capacity is available. Closely associated with the CIR is the Committed Burst Size (Bc) which specifies the maximum size of a traffic burst which the network commits to deliver during a specified time interval (T) and for a specified virtual circuit, ie

CIR = Bc/T

This does not mean that excess bandwidth is always available as the network can only supply bandwidth up to the access rate of the connection to the network. If the network is congested, the bursts may not make it through the network and data can be discarded. Commonly the CIR can be set to a minimum value of 8 Kbps and then increased in blocks of 4 Kbps. However this depends upon the equipment supplier and network operator.

CIR can be implemented on a per interface or per PVC basis, and is computed over a small sliding window interval. If the incoming data over the interface (or PVC) exceeds its pre-specified CIR, the excess data frames are marked with the Discard Eligible (DE) bit. Data frames marked with this DE bit are delivered on a best effort basis and are the ones discarded first in the event of network congestion. Although the CIR of the interface cannot exceed the access speed, the total CIR of all the PVCs through an interface can exceed the physical speed of the interface.

The Excess Burst Size (Be) specifies how much data above the Committed Burst Size (Bc), a user can transmit. While CIR deals with the long term traffic load of an interface (or PVC), Excess Burst Size determines the initial size of a transaction that the network can support. It too is computed over a small sliding window time interval. If the incoming data over an interface (or PVC) exceeds its pre-specified Excess Burst Size, then the network will discard the excess data frames. For a PVC, the CIR and Excess Burst Size can be different for the two directions.

For example Figure 5 shows that frames 1 and 2 fall within the CIR (or Bc). Frame 3 exceeds the CIR (or Bc) and is counted in the Excess Burst Size (Be) and is therefore marked with the DE bit. Finally, frame 4 causes the data rate to exceed the agreed CIR + Excess Burst Size (Bc + Be) and is therefore discarded by the network thus preventing one user from overloading the network at the expense of other users.

CIR and Excess Burst Size is a form of admission control used to regulate the traffic load at the network interface (or PVC). Selecting the CIR and Be permits a user to optimize each PVC for the traffic to be carried. Information in the bursts is transmitted if possible but has a lower probability of getting through. Other admission control mechanisms include sliding windows and the leaky bucket credit manager algorithm.

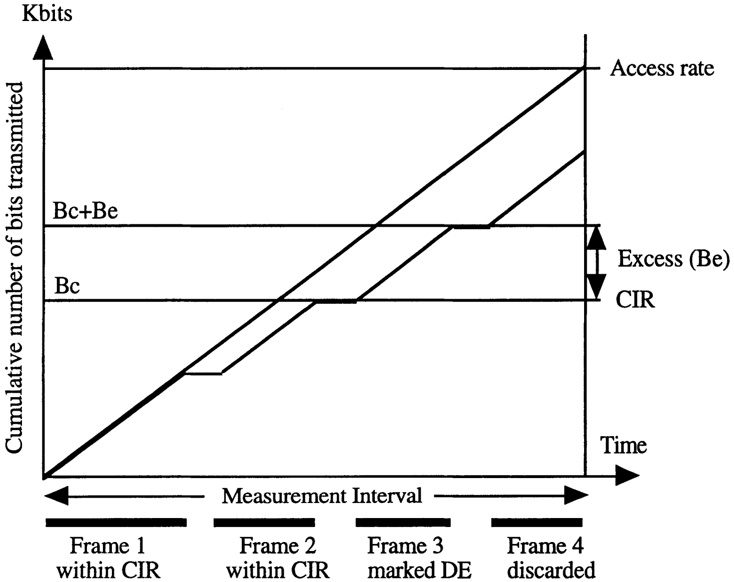

Figure 5 Frame Relay capacity management (Santoso, 1993).

3 ADMISSION CONTROL

The function of admission control is to regulate the volume of traffic entering the network in order that data from a single data source will not degrade the network performance as seen by other users (Chen, 1989). These schemes apply at the network interface and two of the most widely used are the sliding window and leaky bucket credit manager schemes.

Network nodes can still be overloaded even with good admission polices. For example, many end-stations might want to communicate with the same end-station and the combined traffic load to this end-station might overload part of the network. The procedures which manage these overload conditions are known as congestion control procedures.

Specification of the CIR is a relatively course measure and recent developments allow for the *initial*, *minimum* and *maximum* data rates on each PVC to be identified. For example some PVCs can be configured to provide fixed bandwidth while others can be configured to maximise the cost savings resulting from shared bandwidth.

In the sliding window algorithm admission control scheme, each interface (or PVC) is allocated a certain number of credits. As frames are transmitted the credit allocation is decremented while the receiver returns credits via acknowledgements once the frames have been successfully received. In a well tuned system the transmitter receives credits back at a rate such that the window is never exhausted.

In the credit manager scheme these credits represent units of data. They are allocated on a time basis to any unused PVC up to a defined maximum. When a frame arrives for transmission, the cells so generated decrement the current credit but are transmitted at the access speed (one credit per cell).

A PVC at the access interface can accumulate credit at a regular rate, up to the limit. When the end-station transmits the frames over the PVC, the credit allocation is decremented by the amount of data transmitted. The initial maximum credit limit regulates the initial maximum burst size that the network is willing to accept. When the credits for a PVC are exhausted no further frames will be transmitted until a new credit arrives. Under constant load, PVC credits arrive at a rate which permits data to be transmitted at the CIR.

Applications with low utilization and short transaction bursts such as database record processing, terminal/host transactions, electronic mail and perhaps short file transfers all operate well with this mechanism. Variations and enhancements to these two algorithms have been implemented by various Frame Relay vendors. An example of the credit manager algorithm in action is shown in Figure 6.

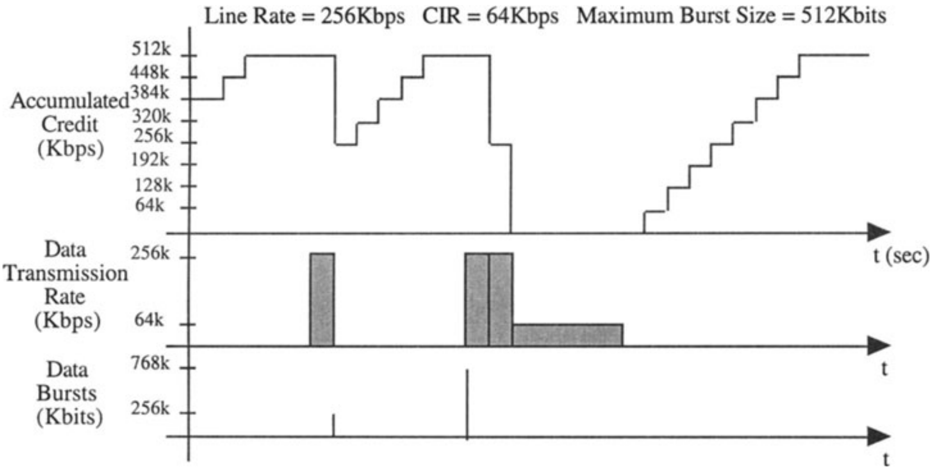

Figure 6 Burst transmission under control of the credit manager algorithm.

4 CONGESTION CONTROL

Congestion control is necessary since there is no mandatory link layer flow control mechanism between the end-station and the network as is found in LAP-B and LAP-D. The standard defines implicit and explicit mechanisms to alleviate network congestion. With implicit congestion control the intelligent end-stations reduce their load into the network once they realize that congestion is occurring. With explicit congestion control, signals are sent from the network to the end-stations to inform them of performance problems.

Explicit congestion notification can be achieved by setting the Forward Error Congestion Notification (FECN) and Backward Error Congestion Notification (BECN) bits in the control field of the frame. In a congestion situation, the Frame Relay network may be forced to discard low priority frames, ie those with the DE flag set.

Both the FECN and BECN indicators are set by the network when congestion occurs. Figure 7 shows how a network might behave as the traffic increases. Initially the throughput increases in step with the offered load. Beyond a certain point (x) congestion commences and some frames are discarded. Both the FECN and BECN indicators would be set in this mid congestion region which enables the network to recover from congestion by moving the operating point back into the non-congested region. Forward notification is useful when the receiver can control the transmitter, usually by means of a higher layer windowing mechanism. Backward notification relies upon traffic flowing in the reverse direction and in this way the transmitting device is informed that congestion is occurring. For example, BECN is used when interworking with X.25 traffic. The rate at which X.25 frames are encapsulated is controlled by reducing the X.25 window. Beyond point (y) the network congestion is severe.

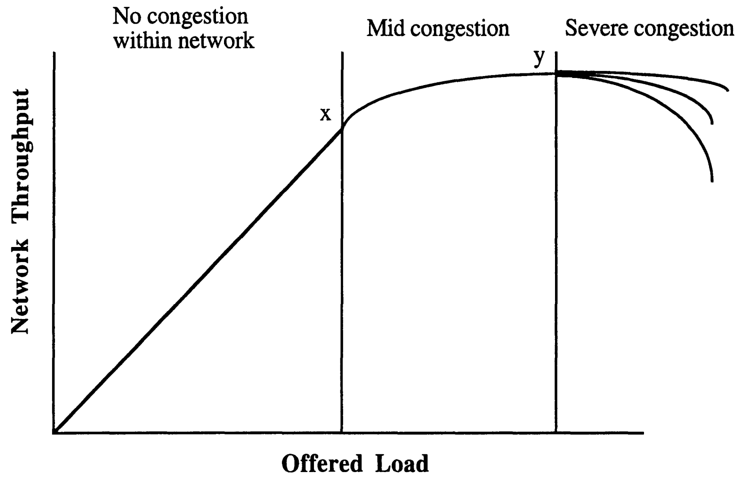

Figure 7 Network throughput versus offered load with various levels of congestion.

The access rate control mechanism has a good deal to do with the size and length of bursts possible above the CIR. Both *Open and Closed Loop Congestion Control Mechanism* are in use and have been implemented by vendors and are described below. There are a variety of congestion algorithms implemented by different vendors as well as many different views on how to set network parameters. The performance comparison of these offerings is difficult making the selection of equipment based upon this performance issue very tricky.

4.1 Open Loop Congestion Control

The simplest (but most inefficient) flow control mechanism is to limit data access to the network by providing explicit window flow control on each PVC with a transmission rate up to the CIR only. Another option is to provide storage capacity for burst data of up to about 5 KBytes. Although such mechanisms might be adequate for a lightly loaded network, better techniques and required in order to provide fine tuning and greater efficiency under conditions of heavy load. Other open loop congestion control mechanisms such as discard eligibility react to congestion by discarding data. This implies that there is no feedback mechanism to indicate congestion and/or no admission control mechanism to prevent congestion. Hence no action is taken until shared truck buffers are full, at which point data must be discarded. Full transmission queues create long delays for all virtual circuits. End-user application performance is reduced by these long queuing delays, and even more so by protocol time-outs caused by discarded frames.

Frames are admitted to the network at the access port speed without buffering. Once the delay experienced by users increases to a critical level, the FECN/BECN bits are set by the network in order to notify end devices that congestion is occurring. Many such devices ignore this notification and, if congestion continues to grow, the Discard Eligibility bit is used in order to decide which frames to discard. The setting of this DE bit is the most common method of handling open loop congestion control. It is the lack of foresight in this open loop system which makes it potentially inefficient as it cannot be assumed that the network will always carry spare capacity.

Although the FECN/BECN bits in the frame are used to signal congestion problems on the network, few access devices or LAN protocols act on this information. Even if they do, the time taken for this congestion information to reach the source end is often outside the time bound caused by the burst of data in the first place.

It would be hoped that in the future access devices will be configured to respond to the FECN/BECN notification bits. For example, such is now the case with the IBM 3745 FEP. It is important to note that even this explicit form of congestion notification cannot always be relied upon as a primary congestion avoidance mechanism since reference is required to higher level protocols at end-stations. Even assuming that voluntary end-user action is correctly implemented, in many instances the reaction time is too slow to avoid congestion. However most protocols operate with a window mechanism and since PVC buffers can often accept up to 64 KBytes, this turns out to be ample storage in order to accept a full window of data in many situations.

4.2 Closed Loop Congestion Control

Unlike other congestion control schemes, closed loop control does not rely on the access devices to reduce their transmission rate. This is important because the Frame Relay flow control signals are presently not implemented in many access devices and in most LAN protocols. Even if implemented, end-to-user device response to explicit congestion notification would typically be too slow to prevent short term congestion.

This lack of foresight on PVC buffering has the potential to cause congestion where many PVCs share the same access trunk. Once the shared buffer is full, frames will be discarded and all PVCs on the trunk will be affected, even if the congestion was primarily caused by just one user.

Thus a more satisfactory approach is to adjust the rate at which data is allowed to enter the network at an *individual* PVC's buffer as a result of information derived from the utilization of the interconnecting network trunks. Data received from the user's node at a rate in excess of the allocated rate enters the appropriate PVCs buffer at the network boundary for a limited period of time. Separate buffers for each virtual circuit at the boundary of the network ensure fairness and improve the efficiency of the higher level LAN protocols. The trunk loading parameter information is piggy-backed with normal frames where data is flowing. In the absence of such data, special control frames are transferred thus minimising any additional load on the network at these critical times.

In this closed loop control schemes (such as that implement by StrataCom by way of their ForeSight *congestion avoidance algorithm* [7]) (Kraemer, 1994), the rate at which data is allowed to enter the network on each virtual circuit is adjusted in response to the utilisation of the network trunks along that data path at that moment. By looking across the network to determine the level of bandwidth contention on each data path, the busiest network segment traversed by a virtual circuit will determine the amount of bandwidth allocated to that virtual circuit at the entrance to the network. A virtual circuit which traverses a downstream trunk that is momentarily fully utilised will have its bandwidth allocation at the access node *decreased*, thus avoiding delays and possible data loss at the downstream trunk. At the same time, another virtual circuit originating at the same node and routed over the same trunk at that node but over less busy downstream trunks may have its bandwidth allocation *increased*. This intelligent bandwidth allocation maximises the performance of all virtual circuits and is an important advantage over open loop congestion control mechanisms.

It is important that these closed loop increments and decrements occur within the expected timings for typical LAN transactions. In one vendor's package these changes can occur 30 - 40 times per second on each PVC. This type of dynamic control is ideal for applications in which transactions last for several seconds - file, document, bit map and image transfers. An ideal solution is one which combines an open loop credit allocation system (to control the rate at which frames are accepted at the network boundary) together with a closed loop foresight control system (which monitors the utilization of interconnecting trunks). This data can then be used to control initial, minimum and maximum bandwidth on each PVC and is an ideal solution to minimize congestion and provide adaptive bandwidth allocation.

Further, if properly used, explicit congestion notification by way of the FECN/BECN notification bits will assist in preventing data loss due to an excessive data arrival rate. If data can be buffered at the network boundary the efficiency of upper layer flow control is maximised. Thus data is delayed rather than incurring the problems associated with frame discarding and application time-out delays for the retransmission which inevitably result.

5 FRAME RELAY PERFORMANCE AND OPERATIONAL ISSUES

Early experience with Frame Relay networks has indicated that properly designed networks do perform well but that there are a number of pitfalls. Carefully analysis must be given to the choice of access data rate, the higher level protocols that are to be transported across the network, the characteristics of the traffic pattern (average and peak transfer rates) as well as the performance characteristics of different vendors' equipment.

Some users indicated that the performance of Frame Relay over DDS offers very little but that the real advantages lie in improved connectivity, ie a "single hole in the wall" with the PVCs providing a replacement for multiple leased DDS circuits thus offering more connectivity per node at a lower cost. Frame Relay has the potential to offer savings with this single physical connection philosophy even though customers have to specify the CIR between every pair of nodes on the network. In almost every case the cost of a Frame Relay configuration will be lower than a partially meshed point-to-point DDS network and always lower than the cost of a fully meshed DDS network.

Congestion management is not well understood, resulting largely from very limited practical experience. Neither is it always clear how the various vendors' switches handle congestion. Although congestion can radically reduce network performance, the answer is to design the network to avoid congestion. However this is often easier said than done as it is difficult to stop a subscriber from flooding the network with substantial volumes of data if their equipment configuration chooses to ignore congestion control advise from the network.

Although Frame Relay has congestion management control built in to the protocol by way of the FECN/BECN bits, a number of Frame Relay routers on the market explicitly ignore these bits and/or are not intelligent enough to respond to congestion notifications from around the network. The matter is further complicated by the fact that users with routers who do obey the requirements of these congestion bits may in fact be opening up bandwidth for users with routers which do not respond to these control measures!

Most Frame Relay switch manufacturers work on the principle that if extra bandwidth is not available when a substantial burst of traffic is received by the switch, frames are forwarded only up to the CIR with the remainder being buffered. Once the buffer has filled to a critical point the FECN/BECN bits are set to alert other nodes on the network that this particular PVC is experiencing congestion problems. If the other nodes do not respond by the time that the switch's buffer is full, frames are discarded requiring the user's end-to-end protocols to arrange for retransmission.

The net result is a slow down in the overall performance and the end-to-end protocols which normally implement a window driven protocol do not continue to send frames indefinitely without some form of returning acknowledgement. Hence the dominating factor becomes *increased delay* and not an excess of *lost data.* Note however that this only affects the individual PVC rather than the other users on the rest of the network. Congestion is most undesirable since it defeats the ideals of a Frame Relay network.

Excess capacity exists in many of the new Telecom Frame Relay networks. This invites customers not to be concerned with internal mechanisms of operation or the idea that data bursts greater than the CIR might not be accepted. However in a more heavily loaded network this can become a major problem.

The overall performance objective is to provide maximum throughput on all PVCs while avoiding congestion that can be caused by momentary traffic bursts. The open loop system simply accepts traffic bursts and discards data if sufficient bandwidth is not available while the closed loop system derives continuous feed-back on trunk utilization from across the network and then correspondingly adjusts the rate of data acceptance on each PVC as previously discussed.

5.1 Network configuration for optimum throughput

In order to determine the access rate and CIR it is essential to analyse traffic volumes and determine to which nodes the traffic is flowing. Getting the CIR correct is a most important issue. If it is *too low* it will cause congestion and poor response times while one that is *too high* will cost the organization financially. Regrettable at the present time there are no magic rules for determining the correct CIR for a particular application. One would normally start with a CIR equivalent to that used for a private point-to-point DDS circuit. For example if a connection utilizes 25% of a 128 Kbps DDS circuit then the CIR could be around 32 Kbps although the utilization of a Frame Relay network is often higher than for a private DDS network. Further, the time period over which this utilization is measured is also a critical factor.

Another way to assist in determining the appropriate CIR is to observe the number of DE (Discard Eligibility) bits set in a traffic sample. If many frames have this bit set it indicates congestion and the CIR should be set higher while few frames with the DE bit set indicates that the CIR may be too high thus costing the organization unnecessarily. However, one major Frame Relay supplier does not normally set the DE bit as it attempts to control congestion at the network boundary in order to avoid discarding frames.

Although bursts can exceed the allocated traffic capacity, in the long run the network cannot carry significantly more traffic than the negotiated CIR as Frame Relay networks are vulnerable to overloading but also the cost of providing excess CIR is high. Another problem occurs when there are too many PVCs per node with low traffic levels per PVC as this results in poor throughput. As the number of PVCs increases, the speed of the access port must be increased accordingly which in turn increases the overall network cost.

Once the CIR is reached, enforcement is up to the network provider who must determine whether additional frames are accepted or rejected and over what time period this might be done. Over-subscription (excess capacity in the network) as discussed earlier is a way of offering some measure of protection for customers who do not have a precise idea about their traffic flows - which turns out to frequently be the case with LAN-to-LAN interconnections. The benefit of over-subscription is that it reduces access port costs for the user. On the other hand if the customer under-estimates the CIR, congestion delays are likely to be encountered.

In configuring connections to a Frame Relay network there are some key issues to consider:

- Determine network traffic requirements - for example the types of traffic to be transported (TCP/IP, IPX/Netware, Appletalk, etc.), to which locations this traffic flows, and what peak volumes are likely to be encountered (eg. file sizes).

- Choice of access data rate appropriate for the applications. It may be necessary to progressively increment the access rate, eg. 64 Kbps → 128 Kbps → 256 Kbps → 1 Mbps → 2 Mbps, in single steps.

- Choice of CIR appropriate for the applications. A realistic CIR will need to be selected. For example a 2.4 Kbps CIR in conjunction with a 128 Kbps access rate is unrealistic. In general

$$\sum_{i=1}^{i=n} CIR_i <= .5\ \textit{Access Rate}$$

  where: CIR_i = Committed Information Rate for the ith channel and
  n = number of PVCs at an individual port operated at the designated *Access Rate* .

- It may be necessary to reconfigure PVCs once traffic statistics are available from the traffic management software.

- Choice of interface equipment (routers etc.) for access to the network.

- Ensuring that the carriers network has been designed to address problems such as congestion in a fair manner.

- Full meshing of the network does not necessarily make for an appropriate configuration. It is better to commence with lightly meshed network topologies and work towards more complex meshing only as the delay requirements dictate. However it should be noted that some routers do not allow transit routing.

- The granularity of the CIR is quite critical if the nodes have low average traffic flows.

- Not all vendors define CIR in the same way. The definition of CIR is determined in part by the standards, in part by the capabilities of the switch and in part by how much the users know about their traffic profiles.

- Buffering must be allowed for in both the Frame Relay switch and the DTE (usually a router). The lower the CIR is chosen to be, the more buffering will be required.

- Other factors to be considered include priority queuing in the routers, what higher layer applications and supporting protocols are to be operated (for example this may require adjustments to time-out parameters), what type of congestion control is implemented, and what statistical and management reporting information is available.

5.2 Performance statistics

Access to a traffic analysis package is essential if traffic profiles are to be accurately determined and the network tuned for optimum performance. A variety of these software monitors are available. One widely used package (StratCom, 1994) reports on a range of variables including transmitted/received frames (peak and average), transmitted/received frames above CIR, marked DE and discarded frames. Packages such as these have gone a long way in improving the overall evaluation of the performance of the network and assists in planning for the connection of additional nodes as well as adjustments in access rate speeds, PVC allocation and CIRs for these PVCs.

6 FUTURE OF FRAME RELAY AND ATM

For the next three to five years the future of Frame Relay looks very positive. Although ATM could ultimately absorb Frame Relay it is more likely that Frame Relay networks will interface with ATM networks. Further, ATM's strength lies in high bit rate transmission (155 Mbps) and not at the bit rates currently being utilised by Frame Relay (up to 2 Mbps). As new applications demand bandwidth for voice and video within the multimedia spectrums, ATM will become the preferred option since Frame Relay is directed to conventional data transmission requirements (LAN-to-LAN in particular).

Higher speed options are increasingly in demand as a result of the enormous growth in the use of Internet. Frame Relay at 34 Mbps (E3) and 45 Mbps (T3) have been demonstrated although it is not clear yet where the Frame Relay/ATM boundary will be positioned. It is certainly questionable whether ATM will replace Frame Relay at speeds below 34 Mbps since ATM relies upon high speed switching and high bandwidth trunks to overcome congestion problems. Therefore it is appropriate to consider Frame Relay and ATM as complimentary networking technologies. It is likely that Frame Relay will be carried over ATM networks (ITU-TSS Recommendation I.555, 1992) (and various protocol structures for the interworking of Frame Relay and ATM have already been defined.

7 CONCLUSIONS

Frame Relay has been the subject of much technical media attention and often unreasonable expectations as to its merits. It has a place in the overall arena of interconnection teleservices. It overlaps in part with services such as DDS, Packet Switching, and some ISDN teleservices but nevertheless offers a very important option for LAN-to-LAN interconnection as well as a variety of other application teleservices. There are many unknown factors relating to issues such as access link speed, CIR, oversubscription, number of PVCs and the all to familiar unknowns about users' traffic profiles.

The Frame Relay interface is designed to be simple, fast and efficient and is optimised for reliable, digital communications circuits. With today's low error rates, it is unnecessary and inefficient to manage acknowledgements and retransmissions at each segment of the network. If a frame is corrupted or lost, it is not retransmitted within the network as acknowledgements and retransmissions are all handled by the end-systems.

Frame Relay has the potential to save network costs when designed appropriately and where applications in turn are tuned to run on these networks. The fact that there are scenarios where Frame Relay does not fit well is not unique as similar comments can be made about most networking systems. If use of Frame Relay is being considered, then care will have to be exercised to ensure that performance and cost improvements can be obtained. Finally, it is important to note that, as with most networking solutions, this service fills a gap in the spectrum and must not be automatically considered to be the solution for LAN or application interconnection in every case.

8 REFERENCES

AT&T Frame Relay Knowledge Base, Release 2.0, (1992), AT&T Data Communication Services, 55 Corporate Drive, Brightwater NJ 08807.

Chen, K.J. and Rege, K.M., (1989) A Comparative Performance Study of Various congestion controls for ISDN Frame Relay Networks, Proc. *IEEE INFOCOM'89*, Ottawa, Canada, pp674-675.

ITU-TSS Recommendation I.555, (1992) Interworking between Frame Mode Bearer Services and other Services.

Kraemer, H, (1994) *Evaluating Frame Relay Platforms*, StrataCom Inc, White Paper #107.

Lamont, J., Doad, J. and Hui, M, (1989) LAN Interconnection via Frame Relaying, *Proc. IEEE INFOCOM'89*, Ottawa, Canada, pp686-690.

Minoli, D. (1995) Technology Overview: Frame Relay, *Managing LANs Datapro Report #1225* McGraw-Hill.

Santoso, H. and Fdida, S., (1993) Frame Relay: A Solution for High Bandwidth Networking, *Computer Communications*, Vol 16 No. 7, pp432-439.

StrataView Plus Traffic Database, (1994) StrataCom Inc, #NMA.02.

Ray Hunt graduated in Electrical Engineering from the University of Canterbury and is now a Senior Lecturer in Communications. He has acted as a telecommunications consultant for a number of Government Departments and Corporations and as an adviser on aspects on local and wide area networking in both New Zealand and Asia. Currently he works in close conjunction with Telecom on aspects of network design, performance and operation and is the supervisor for a number of industry sponsored research projects. He has also worked with the airlines particularly in the area of inter-airline communication protocols.