1998 International Computer Symposium
Workshop on Computer Networks, Internet, and Multimedia
December 17-19, 1998, N.C.K.U., Tainan, Taiwan, R.O.C.

Unknown Frame Forwarding in LANE (LAN Emulation)

I. L- J Thng and Y. S. Gan
Department of Electrical and Electronic Engineering,
National University of Singapore, 10 Kent Ridge Crescent, Singapore 119260, Singapore
Email: eletlj@nus.edu.sg

Keywords: LAN Emulation, LANE, TCP, unknown data frame forwarding, simulation, address resolution process, broadcast forwarding, unicast forwarding.

ABSTRACT

This paper investigates the merits of unknown frame forwarding in an ELAN (Emulated LAN) network supporting higher layer TCP applications. Our simulation results show that if Data Direct VCC setup time is significant in large networks, forwarding of a single unknown data frame is desirable only if the unicast approach is employed. If the broadcast approach is used for forwarding the unknown data frame in a large network, then our simulations show that a simpler ELAN network that does not support unknown frame forwarding has better throughput performance. In the case of forwarding multiple unknown data frames, our studies also show the same conclusions as that of the single unknown frame forwarding case.

1. INTRODUCTION

On March 1995, ATM Forum released LANE specifications version 1.0 [1]. The specifications provided the necessary framework to interconnect Legacy LAN clients onto a high speed ATM network backbone. The resulting network, which is also known as an Emulated LAN, has the ability to emulate the broadcast nature of Legacy LAN and at the same time offer point to point connectivity. All these features are achieved without requiring a change in the Legacy LAN hardware or a re-write of existing Legacy LAN protocol layers operating at or above the LLC (Logical Link Control) layer. The only change that is required is a mere software change at the MAC layer. This is perhaps one of the more attractive features of LANE.

By combining the connectionless nature of the Legacy LAN onto a connection oriented ATM network backbone, an initial latency due to MAC-ATM address mapping is incurred when a LEC (LAN Emulation Client) is connecting to another LEC for the first time. To reduce this latency, the LANE specifications mention the use of the unknown data frame forwarding mechanism. Several specific implementations of unknown frame forwarding has been discussed in the literature [2] but have not been examined through simulation especially in scenarios where the top most protocol layer is an ubiquitous TCP application. In this paper, we contribute such results for two different unknown frame forwarding mechanisms, the first of which is the broadcast approach and the second, the more complex unicast approach. The performance of ELAN networks employing the unknown frame forwarding mechanism are also compared with the performance of similar ELAN networks which does not employ the unknown frame forwarding mechanism. The comparison is important as it provides a measure of practicality in the implementation of the unknown frame forwarding mechanism. This paper is organised as follows: Section 2 provides a brief description of LANE and its essential components. Section 3 discusses the three unknown frame forwarding approaches. A brief discussion of LANE timing latencies is also provided in this section. Section 4 presents the ELAN network that we use to produce our simulation results. Section 5 presents the relevant numerical results and observations, and section 6 concludes the paper.

2. LANE OVER ATM REVISITED

A brief revisit of LANE is presented in this section to enable us to define several important components of the network. Figure 1 shows a schematic illustration of a typical ELAN where several LECs (LANE Client), viz. LEC1, LEC2... etc, are served by the LES (LANE Server) and the BUS (Broadcast and Unknown Server) over an ATM network backbone. There is another component called the LECS (LAN Emulation Configuration Server) and as its name suggests, it is only used to assist configuration of LANE components at initial start-up. Consequently, it has no relevance in this paper. We now provide a brief description of each of the components shown in Figure 1.

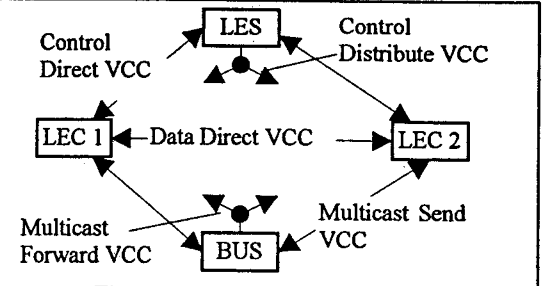

Figure 1. Typical ELAN network

As opposed to the LEC notation used in Figure 1, strictly speaking, the LEC is a protocol layer that has been specified to operate just below the LLC layer (where the MAC layer resided originally) of a Legacy host protocol stack as shown on the left of Figure 2. The LEC interface maps the MAC address of the data from the upper layer into the corresponding ATM address before sending the data to the destination LEC through a Data Direct VCC. An internal address resolution table is maintained in the LEC to facilitate this address mapping. Now, the LEC

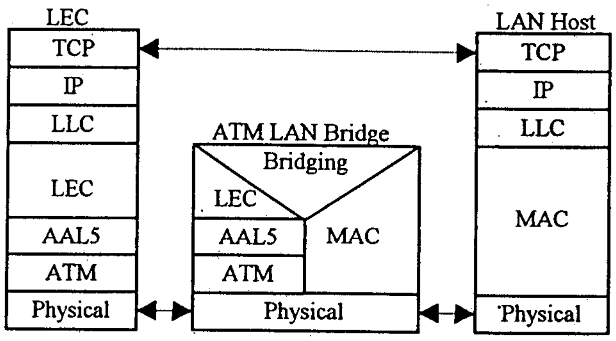

Figure 2. The LEC interface

interface can be said to disguise the ATM network as an Ethernet network whose services are accessed by the upper layers. On the other hand, the LEC interface can just as well be said to perform functions that converts the Legacy host to an ATM host in order to allow interconnection into the ATM backbone. The LEC interface also accommodates bridging functions if an ATM LAN Bridge is used to interconnect an LEC in an ELAN with a Legacy LAN as illustrated in Figure 2. The bridging aspect of LANE is beyond the scope of this paper and thus will not be considered further.

One of the functions of the **LES** is to record the MAC/ATM address pair of each LEC before assigning a unique LEC identifier (LECID) to it. This LECID is embedded in each packet sent in order to identify the sender easily. When the source LEC cannot resolve the destination MAC address into the ATM address successfully, it will send an address resolution protocol (LE_ARP) Request through a Control Direct VCC to the LES. Since all LECs in the ELAN network have registered their MAC/ATM address pairs with the LES on start up, the LES should be able to resolve the unknown destination address successfully. The resolved ATM address is sent back to the source LEC in a LE_ARP Reply. The LEC can then proceed to establish a Data Direct VCC to the destination LEC before sending data to it.

The **BUS** emulates the broadcast services in the Legacy network. Every LEC in the network is connected to the BUS through the Multicast Forward VCC and a Multicast Send VCC before it can send any data. When a LEC needs to broadcast a message, it is sent to the BUS that forwards it to all LECs in the network. Only the LEC whose ATM address is specified as the destination address will retain the received message. Other LECs discard the message when it is received, including the source LEC.

3. UNKNOWN DATA FRAME FORWARDING

Continuing from the previous section, it is noted that there ought to be some latency associated with transmitting messages for the first time. This is because the source LEC's internal address resolution table does not contain the corresponding ATM address mapping of the destination MAC address. Thus, the LEC must first query the LES for the relevant MAC-ATM address mapping before it is able to set up a Data Direct VCC to the destination. The querying of the LES for a MAC-ATM address mapping by an LEC is referred to as an LE_ARP (LAN Emulation Address Resolution Protocol) process in this paper. Thus, the first few data frames that are produced at the TCP layer before the MAC-ATM address mapping information is known are often referred to as unknown data frames. To reduce the latency, the LEC can forward the unknown data frames to the BUS, which will then re-direct the unknown data frames to the destination. This permits the destination LEC to receive part of the message while the MAC address is being resolved. When the mapping information has arrived and the Data Direct VCC established, the forwarding path via the BUS must be flushed of data using the LANE flush protocol. The flush protocol is necessary to prevent out of order delivery of ATM cells in the ATM backbone. Thereafter, the source LEC can use the Data Direct VCC path to send the rest of the message.

Now, the main interest of this paper is to examine the relative merits of three approaches for handling the unknown data frames during the LE_ARP process. These approaches have been discussed in [2] but we shall re-visit these approaches once again in order to introduce a number of new notations.

3.1 The No-Forward Approach

The first approach is what we will refer to as the **No-Forward** approach as illustrated in Figure 3. As the name suggests, the **No-Forward** approach does not forward the unknown data frames to the BUS. It merely waits for the LE_ARP process to complete, and then sets up the Data Direct VCC to send the queued data frames. In this approach, no flushing is required.

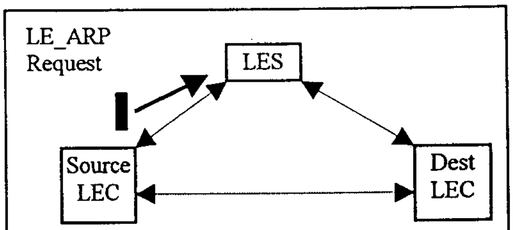

Figure 3. No Forward Approach

The time scale in Figure 4 illustrates the various delay components of the No-Forward approach where the notations used are explained as follows:

- LE_ARP time = Latency due to the LE-ARP process.
- $\delta_D(n)$ = Reception delay of the nth data frame using the Data Direct VCC connection.
- FF_delay (First Frame Delay) = Total time for the first data frame to reach the destination LEC.

It is noted that for the No-Forward approach,

$$FF\_delay = LE\_ARP\ time + Data\ Direct\ VCC\ Setup\ Time + \delta_D(1) \quad (1)$$

3.2 The Broadcast Approach

The second approach is what we will refer to as the Broadcast approach as illustrated in Figure 5. In this approach, the source LEC will forward a certain number of unknown data frames, say, m number of unknown data frames. The BUS, in turn, shall broadcast all m forwarded frames to all LECs via its Multicast Forward VCC connections. When the m number of unknown frames have been forwarded, the flush protocol is then initiated to clear the forwarding path. Now, depending on value of m, two different situations may arise on completion of the flush protocol: the under-forwarding case as illustrated in the time

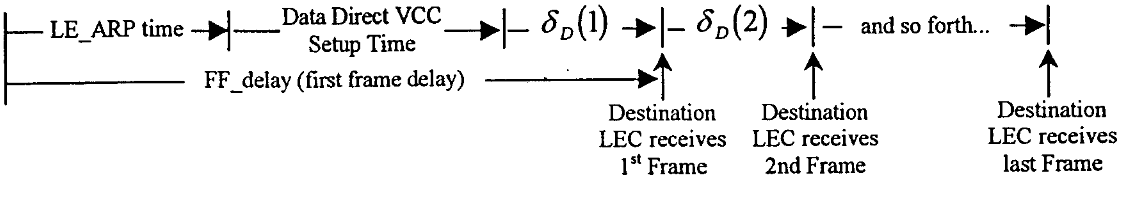

Figure 4. Delay Components of the No-Forward Approach

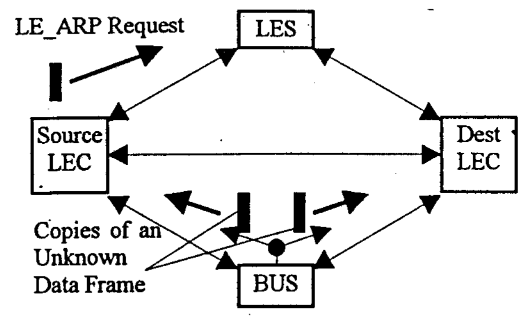

Figure 5. Broadcast Approach

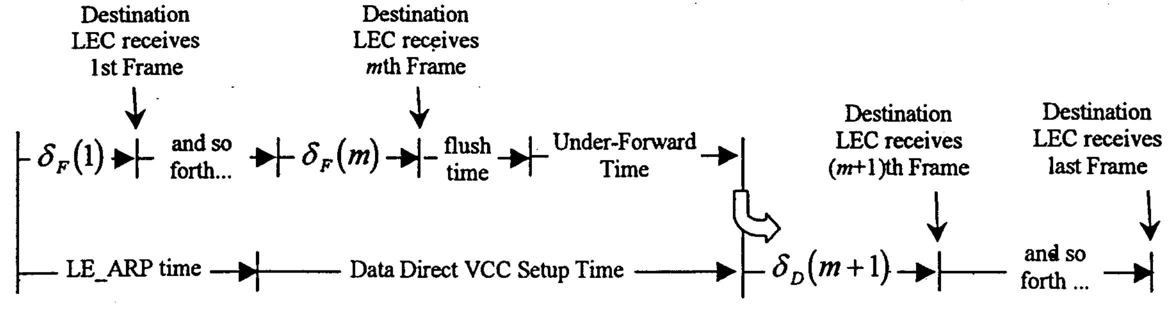

Figure 6. Under-forwarding of Unknown Data Frames

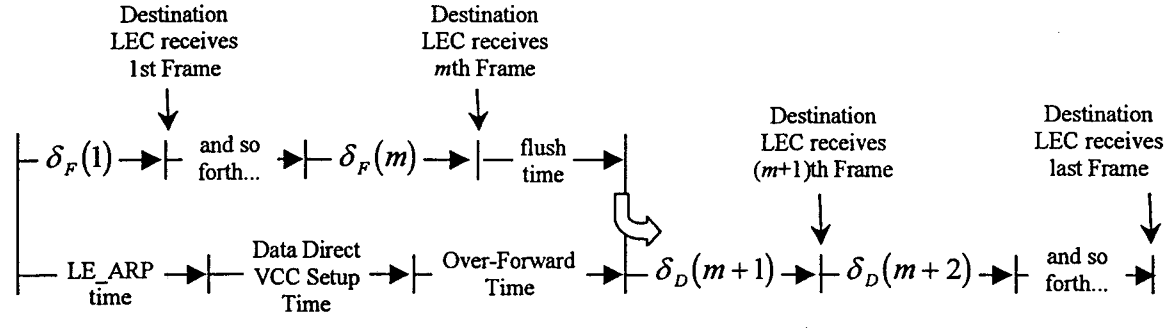

Figure 7. Over-forwarding of Unknown Data Frames

scale of Figure 6 or the over-forwarding case as illustrated in the time scale of Figure 7. The new notations introduced in Figure 6 and Figure 7 are explained as follows:

- $\delta_F(n)$ = Reception delay of the *n*th data frame using the Forwarding path via the BUS.
- Flush time = Time taken to complete the flush protocol.
- Under-Forward Time = Delay component associated with the resumption of data frame transmission (on the Data Direct VCC) due to under- forwarding of unknown data frames.
- Over-Forward Time = Delay in the usage of the Data Direct VCC due to the over-forwarding of unknown data frames.

In the Under-forwarding case illustrated in Figure 6, it is noted that the number of unknown frames forwarded via the unknown frame forwarding mechanism is inadequate. As a result, there exists some delay in the resumption of data transmission (on the Data Direct VCC) after the last unknown frame has been received. This resumption delay, Resume_Delay, is given in general by

$$\text{Resume\_Delay} = \text{flush time} + \text{Under-forward time}, \qquad (2)$$

where the Under-forward time is some value ≥ 0.

In the Over-forwarding case illustrated in Figure 7, it is noted that due to the excessive forwarding of unknown data frames, there exists some delay (i.e. the Over-Forward time) in the usage of the Data Direct VCC. It is also noted that in this case, the Resume_Delay is just given by:

$$\text{Resume\_Delay} = \text{flush time}. \qquad (3)$$

Finally, it is noted that the first frame delay, i.e. FF_delay, for both the under-forwarding and over-forwarding case, is given by

$$\text{FF\_delay} = \delta_F(1). \qquad (4)$$

3.3 The Unicast Approach

The third and final approach is what we will refer to as the **Unicast** approach as illustrated in Figure 8. The **Unicast** approach exploits the co-location of the LES and the BUS to forward unknown data frames to the destination LEC. When the BUS receives the unknown data frames, it maps the data frames to the appropriate destination ATM address by looking into the mapping tables of the LES. With the ATM address, the BUS can then unicast the frames to the destination LEC through the appropriate Multicast Send VCC. No broadcasting is done and the network traffic is expected to be much less than the traffic generated in the **Broadcast** approach. Similar to the **Broadcast** approach, the forwarding path must be flushed before the source

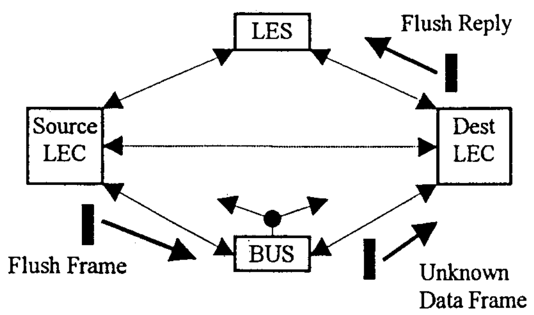

Figure 8. The Unicast Approach

LEC can switch its transmission path to the Data Direct VCC. The delay components of the **Unicast** approach are similar to that of the **Broadcast** approach. Therefore, depending on *m* (the number of unknown frames forwarded), either the Under-forwarding scenario illustrated in Figure 6 or the Over-forwarding scenario illustrated in Figure 7 can arise. The relations in (2), (3) and (4) apply to the **Unicast** approach as well.

4. THE TEST NETWORK

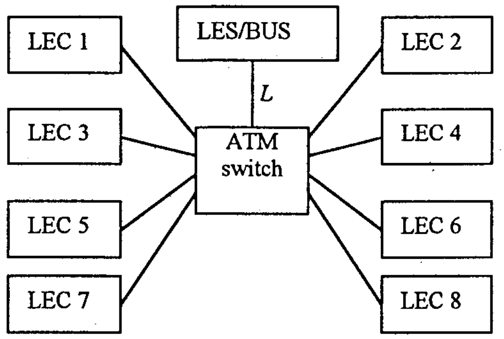

Figure 9. 8-host network

Figure 9 illustrates the general configuration of the 8-host network used in our simulations. We have also considered a 16-host network that is very similar to the 8-host network whereby there are 8 more LECs (hosts) connected to the ATM switch. The simulation modules have been constructed in the NIST ATM Network Simulator [3] environment. The links and the switch are operating at 155 Mbps. The transceiver in the physical layer of each LEC has been specified to operate at a maximum 10 Mbps (Ethernet standard). The capacity of the LES/BUS and the link *L* are scaled up whenever the number of hosts in the network increases. Thus, in the 8-host network, the LES/BUS and link *L* will always operate above 80 Mbps the capacity will always be above 160 Mbps in the 16-host network. This arrangement was made to investigate the bottleneck effect of the ATM switch and not that of the LES/BUS. In addition, all buffer sizes in the LECs, ATM switch and the LES/BUS were set to infinity to prevent any cell loss from occurring that may affect the network performance. All network modules were constructed to focus specifically on the performance of the three unknown frame forwarding mechanisms. More details of the LANE modules that were constructed can be obtained in [4]. We will now describe how data was generated at the top most TCP layer for each LEC.

4.1 The TCP Source

For each LEC, raw messages of size 10500 bytes destined to some other LEC are continuously generated as long as the previous message has been verified by TCP to have been received without errors. The raw message is segmented to TCP payload sizes of 1500 bytes and appended with TCP headers to form TCP data frames. At the ATM layer, TCP data frames are further segmented into 48 bytes ATM payload and appended with ATM cell header. The cells are then transmitted to the destination LEC. Once all data frames have been acknowledged (control frames for acknowledgement have size 0 bytes) via the TCP acknowledgement protocol, the same cycle of events are initiated again where the next batch of raw message of size 10500 bytes is transmitted off to another destination LEC. For the transmission of data frames, all TCP layers will transmit using the TCP slow- start protocol. For an *N*-host network, destination LECs are chosen so that each transmitting LEC will have to initiate LE_ARP processes consecutively for their first *N*-1 messages generated. In addition, we also made sure that none of the individual LEC receiver is a source of bottleneck when multiple source LECs are sending data to it.

5. NUMERICAL RESULTS AND OBSERVATIONS

5.1 Simulation studies on forwarding 1 unknown Frame (m = 1)

The computer simulations presented in this section considered the forwarding of only one unknown data frame, i.e. *m* = 1, for the **Broadcast** and **Unicast** approaches.

The first two plots in Figure 10 illustrates the average raw message throughput measure (i.e. sum of all LEC raw message throughputs divided by the number of LECs) of the respective unknown frame forwarding approaches for the case when the Data Direct VCC set-up time is zero. Such a situation is likely to be seen in small ATM networks catering specifically to non-delay sensitive data where all required Data Direct VCC routes have been established as UBR-type (Unspecified Bit Rate) connections beforehand on a permanent basis. This arrangement

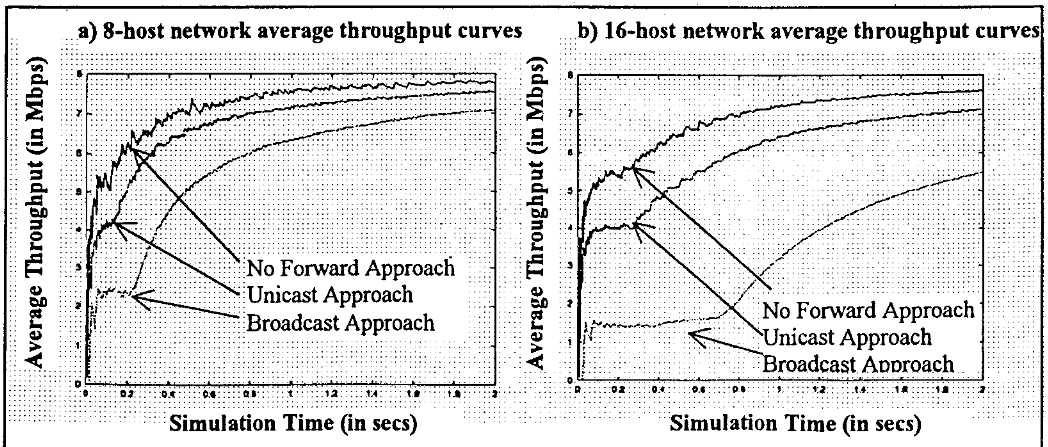

Figure 10. Average throughput curves without Data Direct VCC Setup Time

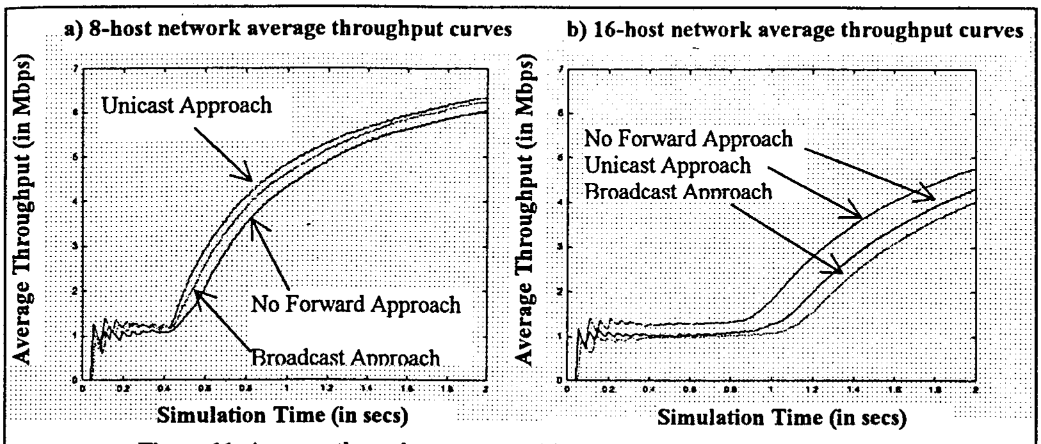

Figure 11. Average throughput curves with Data Direct VCC Setup Time

is also suitable since the problem of frame loss due to cell discards in the UBR connections can be easily overcome by the retransmission protocol of higher TCP layers. We will provide more explanations on the peculiarities of the throughput curves later.

The second set of plots in Figure 11 illustrates the average raw message throughput of the respective unknown frame forwarding approaches for the case when the Data Direct VCC set-up time is set to 0.036 seconds. Now, the Data Direct VCC set-up time varies according to the network components used. Such studies have been publicised in [5] which shows that for a small network (one switch with 8 to 16 nodes), a range of 0.036 seconds (for a network with highly optimised components) to 0.619 seconds (for a network with poorly optimised components) is typically required to set up a Data Direct VCC. For the purpose of our simulation studies in Figure 11, we have decided to assume that our network components are highly optimised since we specified a Data Direct VCC set-up time of 0.036 seconds. We now provide explanations on a number of peculiarities in the throughput measures shown in Figure 10 and Figure 11 as follows:

- ***Rising Throughput and Steady State*** - We first note the gradual rising characteristic of the average throughput curves before settling down to some steady state value just under 10 Mbps. The gradual rising characteristics is due to the TCP slow start mechanism, while the average steady state raw message throughput is expected to fall just below 10 Mbps since all LEC transceivers can operating at 10 Mbps and there are packetisation overheads incurred in the TCP and ATM layers.
- ***Plateau phase*** - We note that for each of the throughput curves, there is a plateau phase in the initial portion of the gradual rising throughput curves. From the figures, the size of the plateau depends on the number of LECs in the network, the data direct VCC set-up time and the type of unknown frame forwarding approach. The plateau phase can best be explained by first recalling that in our simulation scenario, we specified that the first N-1 messages of each LEC, where N is the total number of LEC hosts in the network, will be destined to different LECs. Since the address mapping of destination LECs are all unknown at start up, a source LEC sending its first N-1 messages will require an accompanying LE-ARP process for each message. Now, the LE_ARP process will introduce a break in the transmission. For example, in the No Forward approach, each LEC will go through an initial period where data transmission consists of N-1 stop-start cycles as shown in Figure 12. This initial stop-start phenomena also occurs for the Broadcast and Unicast approach as well. In these cases, the stop periods are due to at least the flush time (or flush time + under-forward time) as illustrated in Figure 6 and Figure 7. By taking the average of all LEC's stop-start throughput phase, we obtain an average throughput that is almost flat for a couple of milliseconds. After all the Data Direct VCC has been set-up and the MAC-ATM mappings have been completed for each LEC, subsequent messages are transmitted on Data Direct VCCs with no interruption. Therefore, the throughput is expected to gradually rise again according to conventional TCP slow-start characteristic and this is confirmed in Figure 10 and Figure 11.
- ***The effect of increasing the number of LECs*** - Both Figure 10 and Figure 11 show that by increasing the number of hosts by two times, the plateau phase of the throughput curves increases by at least twice its previous duration. This is expected since each host will now have to go through twice the number of stop-start cycles initially. The increased traffic brought on by having more hosts should also lengthen the stop-start cycles of each LEC.
- ***Throughput curves of Figure 10*** – Figure 10 illustrates the significantly better throughput performance of the No Forward approach in a network where Data Direct VCC set-up time is zero. In the No Forward approach of Figure 10, only the LE_ARP time contributes to the stop periods during the plateau phase. By comparing the LE_ARP time with the flush time + under-forward time, which contributes to stop periods during the

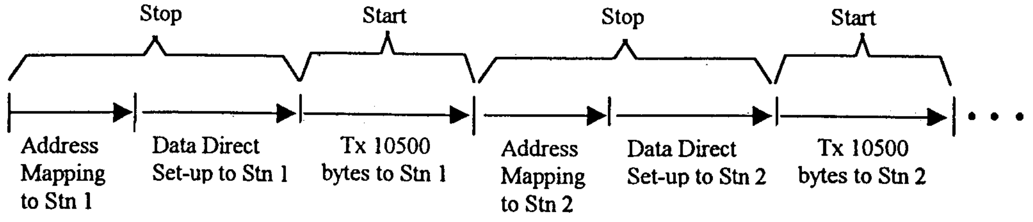

Figure 12: Initial Stop-Start Characteristics of an LEC utilizing the No Forward Approach

plateau phase of the Unicast and Broadcast approach, we found that the LE_ARP time was significantly less. Thus if the No Forward approach is used, then LECs will experience stop periods which are significantly of shorter duration than the stop periods associated with the Unicast and Broadcast approach during initialisation. We note as well that the throughput of the Broadcast approach is the poorest of all. This can be envisaged by considering the number of unnecessary traffic that would be generated during a forwarding process. In this scenario, one data frame (since $m = 1$) is forwarded for each unknown LEC mapping event in the Broadcast Approach. In an N-host network, a total of $N(N-1)$ data frames is forwarded to the BUS which will create N identical frames for broadcasting for each frame received. Thus the total traffic generated by the BUS during the forwarding process is a hefty $N^2(N-1)$ frames. If the network size increases by 2 times, the amount of traffic generated during the forwarding process would increase by 8 times and thus we note the significantly poorer performance of the Broadcast approach when the network size is increased from 8 to 16 in Figure 10a and 10b. In the Unicast approach, the BUS creates only $N(N-1)$ frames during the forwarding process. When the number of hosts increases by 2 times, the number of frames created for forwarding purposes increases 4 times. This explains why in Figure 10a and 10b, the throughput performance of the Unicast approach is better than the Broadcast approach and is also not so badly degraded as that of the Broadcast approach when the number of hosts doubles.

- ***Throughput curves of Figure 11*** – The tables are turned when Data Direct VCC set-up time is considered. In the No Forward approach, the stop periods of the LECs would now consist of the LE_ARP time + Data Direct VCC set-up time (0.036 sec). Adding the Data Direct VCC time does not affect the Unicast and Broadcast approaches that badly as that of the No Forward approach. This is because in the forwarding approaches, the stop periods are due to the flush time + under-forward time and if we refer to the phase diagrams of Figure 6 and Figure 7, only a percentage of the Data Direct VCC set-up time would affect the duration of these two components. We now note that in this scenario, the Unicast approach is desirable in terms of maximising network throughput. As for the Broadcast approach, its throughput performance is slightly better than the No Forward approach in the 8-host network. However, by doubling the number of hosts to 16, the degradation in throughput performance due to increased traffic generated by the Broadcast approach makes it less desirable to implement compared to the No Forward approach.
- ***Statistics of Frame Forwarding time*** – We now present a number of ensemble average measures of the FF_delay (see (1) and (4)) that we obtained for one simulation run of each scenario described in Figure 10 and Figure 11. Table 1 illustrates these measures as follows:

Net. Size	Data Direct VCC Set-up Time (ms)	Ensemble FF_Delay Average (ms)		
		No Forward Approach	**Unicast Approach**	**Broadcast Approach**
8-host	0.0	6.8	6.3	10.3
	36.0	44.6	5.3	9.7
16-host	0.0	7.9	8.9	17.2
	36.0	46.9	6.9	17.6

Table 1. Ensemble FF_delay Average for various unknown data frame forwarding approaches

Note from Table 1 the significant increase in FF_Delay for the No Forward approach when Data Direct VCC set-up time was considered. If we refer to Figure 4, the 6.8 ms figure (row 1, column 3) is the ensemble LE_ARP time + $\delta_D(1)$ average for the No Forward simulation run where Data Direct VCC set-up time was zero. In the subsequent No Forward simulation run, we added a 36 ms component as the Data Direct VCC set-up time. Therefore according to (1), we would expect the ensemble FF_Delay for the subsequent No Forward simulation run to be around 42.8 ms (6.8+36 = 42.8 ms). Due to some minor statistical variations in the traffic generated, we measured 44 ms (row 2, column 3) which is quite close. In contrast, we note that for the Broadcast Approach and Unicast Approach, there are no significant differences in the ensemble FF_Delay average for either values of Data Direct VCC set-up time. This is expected according to (4). Nonetheless, the ensemble FF_Delay average for the Unicast Approach can be seen to be a touch shorter than that of the Broadcast Approach.

5.2 Simulation Studies on Forwarding Multiple Unknown Data Frames

Figure 13 and 14 illustrate the average raw message throughput of the Broadcast and Unicast approach when multiple unknown data frames (i.e. varying m) are forwarded. We have also included the throughput curve of the No Forward approach as a benchmark curve for comparison. All simulation plots were obtained in a scenario where Data Direct VCC set-up time was set to 36 ms. We note a number of interesting observations in regards to Figure 13 and 14 as follows:

- ***Throughput curves of Figure 13*** – The throughput measure of the 8-host network is slightly improved when we set $m = 2$. When $m = 5$, the throughput measure degrades to an extent that the No Forward approach becomes more desirable for implementation. In summary, we make the observation that for the Broadcast approach, implementing multiple unknown frame forwarding does not significantly improve the throughput performance of the network and may even degrade the throughput measure. Note that the total amount of data frames that would have been generated by the BUS during the plateau phase of the Broadcast approach for an 8-host network is of the order $m \times 8 \times 8 \times 7 = 448m$, where m is the number of unknown frames to be forwarded.
- ***Throughput curves of Figure 14*** – In contrast, the performance of the Unicast Approach can be improved by forwarding multiple unknown data frames. This is because selective unicast ensures that traffic load is kept to the minimum during LE_ARP process. A case of over-forwarding occurs when we set $m = 20$. When Over-forwarding occurs, the throughput curves will have a slower rising slope to steady state as observed

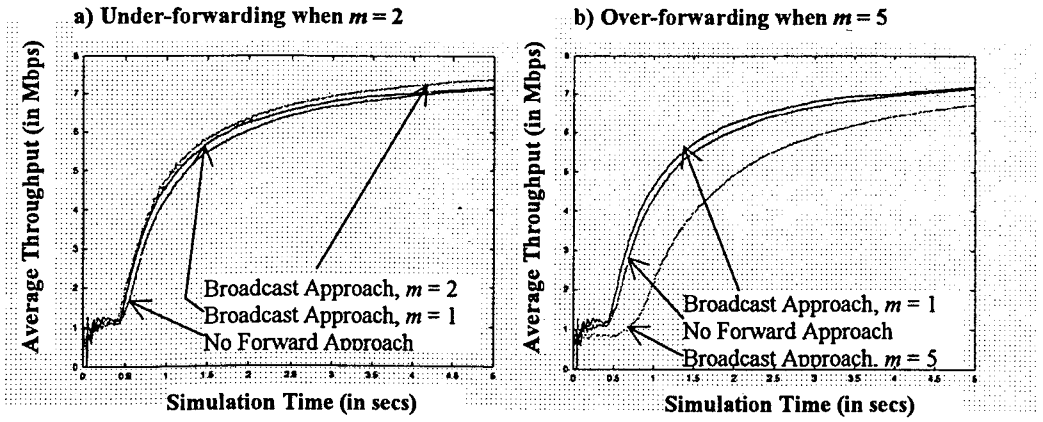

Figure 13. Multiple unknown frame forwarding in Broadcast Approach for 8-host network

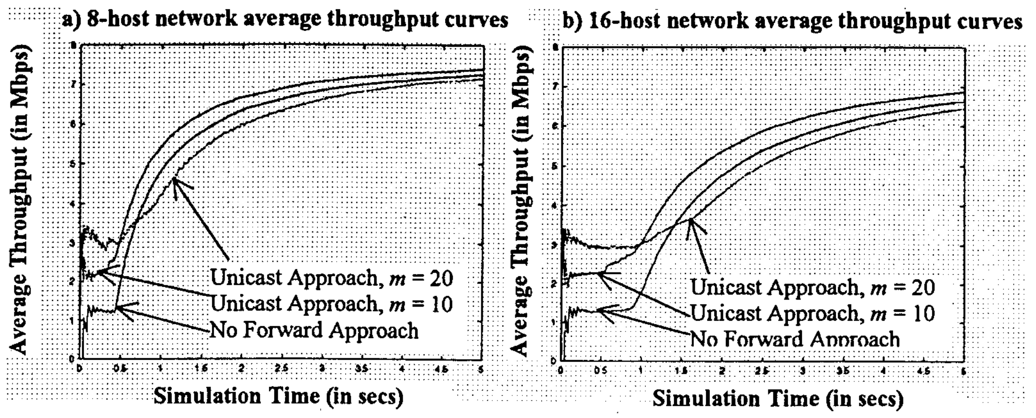

Figure 14. Multiple unknown frame forwarding in Unicast Approach

in Figure 14. This restriction imposes a limit on the number of unknown data frames that can be forwarded.

6. SUMMARY AND CONCLUSION

Several aspects of the unknown frame forwarding mechanism used in ELAN networks with higher layer TCP applications have been investigated closely in this paper. Our simulation studies show that not only does Data Direct VCC setup time determine the need for implementing such a mechanism, the methodology for implementing the unknown frame forwarding mechanism is also important.

When Data Direct VCC setup time is significant and the network is small (8-host network), the throughput performance of an ELAN network that broadcasts a single unknown data frame is found to be marginally better than a similar ELAN network that does not support unknown data frame forwarding. By increasing the network size (16-host network), the ELAN network which uses the broadcast method for forwarding a single unknown frame is now found to be slightly worse off in throughput performance compared to a similar ELAN network which does not support unknown data frame forwarding. The degradation in throughput can best be understood by noting that by increasing the network size linearly, the amount of traffic generated due to the broadcast approach increases by that same amount raised to the cubic power. In contrast, forwarding a single unknown frame by the unicast method is desirable for both small and large networks when Data Direct VCC setup time is significant.

This paper has also presented simulation results on forwarding multiple unknown frames in ELAN networks with significant Data Direct VCC setup times. Our results show that under-forwarding in broadcast approach only improves the throughput measure in the ELAN network marginally compared to a similar ELAN network that does not support the forwarding of unknown frames. Over-forwarding in the same network degrades the throughput measure severely. In the case of the unicast approach, the forwarding of multiple unknown frames is very desirable. However, it is necessary to determine the number of unknown data frames to be forwarded as excessive forwarding can affect the throughput of the network to the extent that it can even fall below the throughput measures of an ELAN network which does not support unknown frame forwarding. The optimal number of unknown frames to be forwarded ought to be that number of data frames to make up the LE_ARP process so that when the Data Direct VCC is established, data frames can be immediately transmitted on it.

In conclusion, unicast forwarding of unknown data frame is desirable in ELAN networks where Data Direct VCC setup time is significant and the amount of unknown frames that are forwarded is not too excessive. Our simulation results does not support the use of the simpler broadcast approach as it has been found that the even simpler ELAN network that does not support unknown frame forwarding can have better throughput in many scenarios. Nevertheless, the broadcast and the unicast approach permit the destination to receive part of the data while its address is being resolved. This should improve the initial response time of the network at the possible expense of a good network throughput performance.

7. REFERENCES

[1] "LAN Emulation over ATM version 1.0", ATM Forum af-lane-0021.000

[2] Hong Linh Truong, William W. Ellington Jr., Jean-Yves Le Boudec, Andreas X. Meier, and J. Wayne Pace, "LAN

Emulation on an ATM Network", IEEE Communications Magazine May 1995, pp. 70 to 85.

[3] "The NIST ATM Network Simulator: Operation and Programming version 1.0", National Institute of Standards and Technology.

[4] Gan Yung Sze, "Flush Protocol in LAN Emulation over ATM", National University of Singapore 1998, pp. 21 to 36.

[5] Robert Manville and Devil Shah, "Mix and Match ATM", Data Communications Tester's Choice July 1997, pp. 43 to 57.