

U.S. DEPARTMENT OF COMMERCE
National Technical Information Service

AD-A036 370

BEHAVIOR OF A LINK IN A PVC NETWORK

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON

7 DECEMBER 1976

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-76-338	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Behavior of a Link in a PVC Network		5. TYPE OF REPORT & PERIOD COVERED Technical Note
		6. PERFORMING ORG. REPORT NUMBER Technical Note 1976-45
7. AUTHOR(s) Alan G. Nemeth		8. CONTRACT OR GRANT NUMBER(s) F19628-76-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P. O. Box 73 Lexington, MA 02173		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element No. 33126K
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Agency 8th Street & So. Courthouse Road Arlington, VA 22204		12. REPORT DATE 7 December 1976
		13. NUMBER OF PAGES 24 23
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB Bedford, MA 01731		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) PVC network voice traffic queueing systems speech terminals data traffic		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The purpose of this note is to discuss the properties of a single link in a Packetized Virtual Circuit (PVC) network under a load consisting of a combination of voice and data traffic. Such properties as the mean waiting time for data packets, overall line utilization, and required amount of buffering are investigated as a function of the percentage of the channel dedicated to voice. A combination of simulation and analytic techniques have been used to obtain the results described here. These results indicate the possibility of achieving line utilizations of up to 98%, within quite practical values of buffering, and packet delay.		

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

BEHAVIOR OF A LINK IN A PVC NETWORK

A. G. NEMETH

Group 24



TECHNICAL NOTE 1976-45

7 DECEMBER 1976

Approved for public release; distribution unlimited.

LEXINGTON

ii

MASSACHUSETTS

ABSTRACT

The purpose of this note is to discuss the properties of a single link in a Packetized Virtual Circuit (PVC) network under a load consisting of a combination of voice and data traffic. Such properties as the mean waiting time for data packets, overall line utilization, and required amount of buffering are investigated as a function of the percentage of the channel dedicated to voice. A combination of simulation and analytic techniques have been used to obtain the results described here. These results indicate the possibility of achieving line utilizations of up to 98%, within quite practical values of buffering, and packet delay.

ADMISSION for ☒ White Section ☐
☐ Suit Section ☐
 NAME _____
 NO. _____
 EXAMINATION NO. _____
 JUSTIFICATION _____
 BY _____
 EXPIRATION / AVAILABILITY DATE _____
 DIS. _____
 401L GROUP SERIAL _____
 A

TABLE OF CONTENTS

ABSTRACT	iii
I. INTRODUCTION	1
II. ANALYTIC MODEL	3
III. SIMULATION RESULTS	3
IV. COMPARISON WITH OTHER TYPES OF NETWORKS	9
V. CRITIQUE	10
ACKNOWLEDGEMENTS	15
REFERENCES	16

I. INTRODUCTION

The Packetized Virtual Circuit (PVC) concept is a scheme for designing an integrated data/voice digital network, which has desirable properties of high line utilization, and low delay through the network. This behavior is obtained by the application of several design principles:

1. Two users A and B which are transmitting traffic from A to B are likely to continue to do so; therefore, a connection based scheme is employed. A and B must request the establishment of a connection before the transmission of any traffic. This permits the use of an abbreviated header on each packet. For example, typical header information in the network can be encoded in 16 bits including check bits.

2. Flow control is handled in a statistical fashion. No flow control operations are performed on a packet by packet basis anywhere in the network. This reduces the amount of protocol traffic, but requires us to adopt the attitude that buffer overflow at an internal node in the network is not catastrophic; rather, the loss of data may be recovered by end-to-end protocols which must already exist to cope with line errors. Flow control may then be provided by restricting expected line utilization to reduce the probability of buffer overflow to any desired value. Then, the routing scheme used for establishing connections also provides the flow control in the network. Further, voice packets may be occasionally lost without any need for retransmission, for the ear will not notice gaps in the speech with durations of tens of milliseconds if an appropriate reconstruction algorithm is used, and the gaps do not happen very frequently.

A computer simulation has been constructed which can model the behavior of a single link in a PVC network. The computer models a population

of speakers providing a voice load on the system, but does not model any data traffic. In order to investigate the link behavior under a variety of data loads, an analytic model of the behavior of the data traffic has been generated, which uses values for its parameters obtained from the computer simulation. The model about to be described has a number of assumptions built into it about the behavior of data traffic, and the relative priority of voice and data traffic. We shall later return to a discussion of relaxing these restrictions.

II. ANALYTIC MODEL

We begin then by assuming that voice traffic has an absolute priority over data traffic. This assumption permits the data traffic to be modeled by a single server queue, where the service time for data packets is determined by the empty slots left in the channel after voice traffic has been transmitted. Thus, given the distribution of gaps between empty slots for the voice channel observed alone, we can then model the data channel without further regard for the details of voice channel behavior. In the succeeding work, the simulator has been used to collect the mean and variance of the time between empty slots in the voice traffic to obtain the service time distribution for the data channel.

The next assumption we make is that the data traffic arrives according to a Poisson model -- successive interarrival times are independent, and drawn from an exponential distribution. These assumptions then lead to an M/G/1 queue for which a number of results may be found in the literature [1]. In particular, the mean waiting time, $\bar{W}$ is given by the following expression:

$$\bar{W} = \frac{\lambda \overline{X^2}}{2(1-\rho_D)}$$

where

λ is the rate of arriving data traffic in packets/sec.

$\overline{X^2}$ is the variance of the service time,

ρ_D is the fractional utilization of the data channel.

If $\bar{X}$ is the mean time to service a data packet, then we have the following relation:

$$\rho_D = \lambda \bar{X} \quad .$$

This then gives

$$\bar{W} = \frac{\overline{X^2}}{2\bar{X}} \frac{\rho_D}{1-\rho_D} \quad .$$

The maximum capacity of the data channel R_D is given by

$$R_D = \frac{(\# \text{ of data bits/packet})}{\bar{X}} \quad .$$

We will assume that each packet consists of 8 16-bit words (128 bits) of which 32 are overhead. Because of this assumption of short packets, we will assume further that the packets are fully utilized, i.e., each packet contains 96 bits of useful data. With these assumptions, we may then plot $\bar{W}$ (the mean waiting time) as a function of $\bar{R}_D$, the rate of data traffic in bits/sec, given values for $\bar{X}$ and $\overline{X^2}$.

III. SIMULATION RESULTS

We have obtained values for $\bar{X}$ and $\overline{X^2}$ by simulation of the voice load on the channel. The data presented in Table I are based on a number of

TABLE I

	$\bar{X}$ μsec	σ_X μsec	ρ_V	MAX R_D kbits/s
130 CVSD	834	15,033	90.0%	115.066
120 CVSD	457	1,677	81.8%	210.159
110 CVSD	371	1,023	77.5%	258.813
100 CVSD	284	440	70.6%	338.537
75 CVSD	172	189	51.4%	559.302
100 CVSD, 135 LPC	865	6,820	90.4%	110.990
75 CVSD, 110 LPC	256	379	67.5%	374.613

assumptions:

(1) The link transmits 12000 packets/sec. Each packet contains 96 data bits, and 32 overhead bits, for a total of 128 bits/packet, or 1.536 MB/sec.

(2) Each speaker may independently be in silence, or active speech. If in silence, no packets are transmitted; in active speech, packets are transmitted at a rate determined by the vocoding technique. The decision for each speaker is made independently on the basis of curves representing the distribution of the duration of talkspurts and pauses. These distributions have been measured extensively by Brady [2]. Two representative techniques are used for the results here: 16 KB CVSD generating a packet every 6 msec or 3.6 KB LPC generating a data packet every 27.5 msec.

(3) The data represent 2 minutes of simulated link time.

(4) ρ_V is the fraction of the channel utilized for voice traffic.

(5) If the queue for voice traffic reaches 70 packets (5.83 msec delay), then half the packets are arbitrarily removed. In fact, this loss of data occurred only in the cases with $\rho_V > 80\%$. Only in the worst case (130 CVSD), was there an appreciable period of time (3.5 seconds) during which speakers would lose a significant ($\sim 13\%$) fraction of the packets transmitted.

(6) These data are the result of one simulation run for each case. The standard deviation for the service time, in cases where $\rho_V > 80\%$, has itself a large deviation, which is believed to account for the relatively low value of

σ_x in the 120 CVSD case. More extensive computer runs would undoubtedly modify these results, but the overall conclusions, and the amount of modification would be relatively small.

These values may then be used to derive the curves shown in Figure 1 showing the relationship of mean waiting time to data rate for the data channel.

As an example, these curves tell us that on a 1.536 MB/s channel carrying 100 16 KB/s CVSD speakers, we may also carry about 313 KB/sec of data traffic with an average data packet delay of 6 milliseconds. This represents a long term channel utilization of 1084 KB/sec for voice, and 417 KB/sec for data for a total of 1501 KB/sec of traffic, or a line utilization of 97.7%. These figures for line utilization include the fact that each packet contains 32 bits of routing information that utilize the channel. If these bits are not counted then a line utilization of 73.3% for user data and voice traffic is obtained.

Note that in this example, because the voice load on the channel is about 71%, no voice packets were observed to be discarded. Further, the average space utilized for data packet buffers is about 160 16-bit words or 2560 bits.

The rather obvious hitch in all this good news comes from considering what the maximum delay, or alternatively, what the maximum amount of buffering will be for the data traffic. This figure is not so easily obtained but some estimates can be obtained.

One technique for estimating the amount of buffer space is based on an approximation to the waiting time distribution for a G/G/1 queue under heavy traffic [1]. By dropping the assumption that data traffic arrivals




Fig. 1. Mean waiting time vs. data rate; vertical bars are maximum data rate; various levels of voice load are shown.

are Poisson, we obtain for $\rho_D \approx 1$, that the waiting time is approximately exponentially distributed with mean wait given by

$$(\sigma_A^2 + \sigma_X^2) / 2(1 - \rho_D) \bar{t}$$

where $\bar{t}$ = mean interarrival time for data packets, and σ_A = standard deviation of the interarrival time. This approximation is valid when

$$2(1 - \rho) \bar{t} \ll \sqrt{\sigma_A^2 + \sigma_X^2}.$$

This approximation then permits us to estimate the waiting time achieved with certain low probabilities, i.e., we may find that, for some particular case, the probability of waiting longer than 100 millisec is 10^{-7} . We may then use this information to estimate further the amount of buffering space required to store this traffic, i.e., 100 millisec of traffic represents about 3000 words of transmission at the mean service time of the channel, or expressed differently there is a probability of 10^{-7} of overflowing a 3000 word buffer.

We may use this technique to estimate buffering requirements for the 100 CVSD speaker example discussed earlier. In order to do this, we must have a value for σ_A , the standard deviation of the interarrival time. In the earlier example, by assuming Poisson arrivals, we have used a standard deviation equal to the mean. Continuing with this assumption, we obtain that the waiting time is exponentially distributed with a mean of 6211 μ sec. Note the small discrepancy from the M/G/1 result of 6000 μ sec. This difference comes from the more general nature of the arrival time distribution,

leaving the possibility of more skewed arrival patterns (reflected in the values of third and higher moments of the interarrival time distribution).

We may then tabulate for specific probabilities of overload, P_{OVER} , the values of waiting time W_{OVER} which are achieved with that probability, and the amount of buffering represented. Table 2 shows some values for the 100 CVSD speakers, 313 Kb/s data traffic example.

TABLE 2

P_{OVER}	W_{OVER} (μsec)	Buffer space 16-bit words
10^{-3}	42,904	1209
10^{-4}	57,205	1611
10^{-5}	71,506	2014
10^{-6}	85,808	2417

Thus, to summarize the example, on a 1.536 MB/sec channel, we are able to place 100 16 KB/sec CVSD speakers, and 313 KB/sec of data traffic with good values of packet delay (voice < 6 msec, data-average 6 msec) and practical amounts of buffering (560 16-bit words for voice, ~ 2000 16-bit words for data), to obtain a channel utilization of 97.7%.

IV. COMPARISON WITH OTHER TYPES OF NETWORKS

The example presented above may be compared with performance on a pure circuit switched network as well as a pure packet switched network.

In a circuit switched network, to support 100 16 KB/sec CVSD speakers would require 1600 KB/sec data rate on the link without including any multiplexing overhead. This is already larger than the capacity of the 1.544 MB/sec link we have assumed. The only mechanism in such a network

for handling data traffic would be to establish circuits capable of handling the peak data rates for each data connection. If we assume a peak-to-average ratio for data traffic of 10, we would need 3.13 MB/sec to handle the 3.3 KB/sec average data traffic in our example. Thus, we would require a link at 4.73 MB/sec which would be utilized at 29.5% to support the load. Delay for both voice and data traffic would be low, as well as the amount of buffering required.

Since no pure packet switched networks have approached the load discussed here, comparison is more difficult. Choosing the ARPANET as an example, we can scale up loading to carry our assumed load. Using data from Kleinrock [1] which predicts ARPANET performance on a fully loaded 50KB/sec link will support about 11.7KB/sec of user data traffic and 6.3 KB/sec of host-host protocol, or a total of 18 KB/sec of host-host traffic, we can extrapolate the link capacity needed to support our load by a linear scaling. This leads to an estimate of a link capacity of 869 KB/sec to support 313KB/sec of data traffic, and 3011 KB/sec to support 100 16 KB CVSD speakers, or a total link capacity of 3880 KB/sec, which would be utilized at 36.0%. Previous measurements of ARPANET delay [3] under typical current load conditions find a median per node delay of about 28 millisec/packet. This could be expected to increase with significant increases in line utilization.

V. CRITIQUE

As we have proceeded through the analysis, we have made a number of assumptions which are properly questioned:

- (1) Voice traffic has absolute priority of service on the channel.
- (2) Data traffic is a Poisson stream.
- (3) Successive service times on the data channel are independent.

In this section, I would like to explore the meaning of these restrictions, and investigate means of relaxing them.

- (1) Priority of Voice Traffic

This assumption allowed us to divide the problem into two separate pieces representing the voice and data traffic. If, indeed, we are operating at a probability of overflow for data traffic of 10^{-4} , we can, without significantly impacting the analysis, decide that instead of losing data packets, we will discard voice packets. Because of the low probability of occurrence, this technique should not significantly affect voice channel performance while noticeably aiding data performance.

Perhaps more interestingly, the data waiting time and buffering space both are directly dependent on the variance of the service time. If we control the voice channel to reduce the maximum gap between data packets, then the variance of the service time can be noticeably reduced, affecting both performance measures. Depending upon the voice load, this will cause either a slight increase in delays for the voice traffic, or may push the voice channel

into a region where loss of voice packets occurs more frequently.

(2) Poisson Data Traffic Assumption

We have assumed that data traffic is generated according to a Poisson rule. Since it is well known that data traffic is bursty in nature, the assumption requires further discussion.

One possibility is to view the data traffic as made up of many separate users. For example, if the 313 KB/sec in the example is divided among about 63 unsynchronized users each having a 10 KB/sec link to the network which is used 50% of the time, then the burst nature of an individual data user becomes insignificant. The total collection of users now provides an approximately Poisson population for analysis.

In conjunction with this, we postulate a rate control on each data user to control his peak data rate, and monitor his average data rate. Network routing and thus flow control would be performed on the basis of average data rate for each user, while each user can instantaneously use the peak rate. So long as the number of data users is large enough to provide some advantage from the law of large numbers, this approach will work. The peak rate for any individual user must be well below the rate of the data channel to avoid the possibility that a burst from a small number of users could swamp the channel.

If it is desired to provide peak rate service approaching the rate of the data channel, then some control can be imposed on the duration of peak rate bursts to prevent channel saturation. For example, if a user is capable at peak rate of generating 1000 packets/second, and his average rate is 100

packets/second, then the limiter might restrict his usage so that he may send a maximum of 50 packets in 50 milliseconds, a maximum of 100 packets in 250 milliseconds, and a maximum of 150 packets in one second. Such a scheme might be implemented by keeping counts which decay exponentially with different time constants. If any count is instantaneously above its limit, then no further input may be accepted until the count has decayed below the limit.

(3) Independent Service Times

The analysis assumes that service times for successive packets on the data channel are independent. Since the number of active speakers in the voice channel is a slowly varying function of time, this assumption may not be valid. There is enough data available to test this question in the simulator, and these tests should be performed in the future. Stated formally, the question is whether the sequence of empty slots in the voice channel forms a renewal process. Following Cox and Lewis [4], there are two major approaches to testing the hypothesis that this data was generated by a renewal process. The simple test involves calculation of the serial correlation coefficient for successive intergap times in the channel. More extensive tests include autocorrelation of the series of intergap times for many different lags by FFT techniques, spectral analysis of this series, and tests for independent in Markov chains. Cox and Lewis provides the necessary details but the tests have not been carried out.

The effects of the failure of the independence assumption would be to change the amount of buffering required for the data channel. Although this change may be in either direction, the direction of concern is obviously an increase in the buffering requirements. Although some deviation from independence may be expected, this deviation appears to be small and should not significantly change the results represented here. Since the amounts of data buffering explored in this memo were fairly small (~3000 16-bit words), a significant change in this number, even as much as a factor of 5, would not change the overall conclusions.

ACKNOWLEDGEMENTS

The original concept for the PVC network comes from Jim Forgie. Without that original concept, this work could not have been performed. Ronald Sonderegger and Gino Coviello of the Defense Communications Engineering Center (DCEC) have kindly supported this work and offered a number of helpful suggestions. Thanks to Dr. Ben Gold and Dr. Ted Bially for picking an error out of an early draft and for comments on presentation. Finally, thanks to Chris Youell for keeping up with the continuous changes in the manuscript.

REFERENCES

- [1] L. Kleinrock, Queueing Systems, Volume 2: Computer Applications (Wiley, New York, 1976).
- [2] P.T. Brady, "A Technique for Investigating On-Off Patterns of Speech," Bell Syst. Tech. J. XLIV (1965).
- [3] J.W. Forgie and C.K. McElwain, "ARPANET Delay Measurements," NSC Note #70 (July 1975).
- [4] D.R. Cox and P.A.W. Lewis, The Statistical Analysis of Series of Events, (Methuen, London, 1966).