Proceedings of the 1985 Winter Simulation Conference
D. Gantz, G. Blais, S. Solomon (eds.)

"Multirate Voice Coding for Load Control on CSMA/CD Local Computer Networks"

Victor S. Frost, Edward M. Friedman, and Gary J. Minden

Telecommunications and Information Sciences Laboratory
2291 Irving Hill Drive
Lawrence, Kansas 66045

ABSTRACT

A study of mixed voice and data on a local area network under the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is presented. The study determined the feasibility of using multirate voice coding to control the traffic intensity on the network. By decreasing the voice coding rate for short periods of time (and thus the voice quality) when network traffic increases, a larger number of voice users can be realized. The premise is that short term voice quality can be traded for increased throughput on a CSMA/CD network. Collisions per millisecond is found to give a good indication of the traffic on the network. A feedback equation is used to set the voice coding rate, based on the measured collisions per millisecond. The coding rate determined from the feedback equation is rounded to one of four possible coding rates; other rates are clearly possible. The voice coding rates chosen were 48, 40, 32, and 24 kilo-bits-per-second. The simulation results show that approximately 10 additional voice users can be added to a 1 mega-bit-per-second network when the multirate algorithm is used.

Results are given relating the packet delay to the collisions per millisecond. A comparison between multirate and non-multirate is given. An indication of the voice quality for the multirate system is also discussed. The multirate load control algorithm provides a way to increase the amount of voice traffic on a CSMA/CD network without increasing the line capacity.

1. INTRODUCTION

There is considerable interest in integrated voice and data (IVD) communications network. One mechanism for providing IVD communications is to use PBX technology to provide both data and voice switching functions. Another possibility is to use local area network technology (LAN). In this paper we describe the performance of a LAN which is used for both voice and data transmission.

There have been previous studies of combined voice/data Local Area Networks [1,2,3]. The study done by DeTreville [1] was a comparison of the CSMA/CD type of network with a token bus. In that study the token bus was found to perform somewhat better than the CSMA/CD protocol. It was found in [1] that a 10 Mbps CSMA/CD LAN could support about 150 voice conversations with no data traffic present (for 100 ms of total delay through the system). A 5% data load reduced this to 125 conversations.

The Nutt and Bayer study [2] gives a comparison of various backoff algorithms appropriate for voice. In the study by Musser, et al., [3] the virtual token protocol (GBRAM) and the CSMA/CD protocol were compared. In addition, listener tests were performed which confirmed that two percent of the voice packets could be lost and the voice quality would still be acceptable to the user. It was found in [3] that a 1 Mbps CSMA/CD LAN could support 12 conversations or 94 conversations on a 10 Mbps network. The results presented in this study agree with the results given in [1,2,3].

This study was conducted to determine the feasibility in using multirate voice coding [4,5] to control the traffic intensity on the network. By decreasing the voice coding rate for short periods of time (and thus the voice quality) when the traffic on the network increased, a larger number of voice users could be realized. The premise is that short term voice quality can be traded for increased throughput on a CSMA/CD network. A model for this system is shown in Figure 1. It was clear that some convenient measure of the instantaneous traffic intensity was needed to drive the feedback loop. Collisions per millisecond was found to give a good indication of the traffic on the network. Furthermore, collisions per millisecond could be easily measured. A feedback equation was developed to set the voice coding rate, based on the measured collisions per millisecond. The coding rate was rounded to one of four possible coding rates (24,32,40,48 kbps); other rates are possible.

It was not the intent of this study to define the optimum voice coding rates for a multirate LAN; rather it will be shown that a significant improvement, i.e., increased number of users, can be obtained if multirate coding is used. Lower coding rates will clearly increase the network voice capacity. Computer simulation was used to test the multirate concepts. The simulation results show that approximately 10 additional voice users (an increase to 22 from 12) can be added to a 1 mega-bit-per-second network, if the multirate algorithm is used.

2. MULTIRATE VOICE CODING ON A LAN

Voice on a LAN is characterized by periodic arrivals of packets. Periodic voice packet arrivals models a voice source as being continuously in the talk-spurt mode, or a two-way conversation with silence detection. The voice packet is characterized by its generation period. The generation period is the packet length in bits divided by the voice coding rate in bits-per-second.

This work was supported by a National Science Foundation Young Investigation Award and AT&T Information Systems

$$G = P/R \quad (1)$$

where,

G = generation period in seconds,
P = packet length in bits
R = voice coding rate in bits-per-second

There is a continuous flow of voice packets from the source separated by the generation time. The time to transmit a voice packet over the network is much less than the generation period (the packet length in bits divided by the line capacity in bits-per-second). If a packet has not successfully transmitted over the network in the generation period, the packet will be discarded and is considered lost. It has been shown [3] that two percent of the voice packets can be lost and there will be no effect on the voice quality.

Voice is a compressible source, that is, voice can be coded at different rates and the voice quality decreases as the coding rate decreases [4]. A technique known as embedded coding [5] allows the same encoder/decoder structure to be used for a variety of coding rates. The idea is that voice generated at a high coding rate can have selected bits stripped away resulting in an effective lower rate. In this study we kept the voice packet size in bits constant (768 bits with no overhead bits included) and thus each packet represented a longer sequence of voice signal (i.e., the generation period increases) as more bits were stripped away and the coding rate was lowered.

The basic idea of multirate voice coding for load control is that the voice quality can be traded for network load. When an increase in the traffic intensity on the network has been observed, the voice coding rate will be lowered. By decreasing the voice coding rate the generation period is increased causing fewer voice packets to access the network. The decrease in coding rate reduces the traffic on the network. Thus, a feedback system results. Practical embedded voice coders work at a finite set of rates. The rates used and their corresponding packet generation times are shown in Table 1.

Voice Coding Rate, kbps	Generation Period, ms
48.0	16.0
40.0	19.2
32.0	24.0
24.0	32.0

Table 1. The chosen voice coding rates and the corresponding generation periods, for a voice packet length of 768 bits.

3. THE LOAD CONTROL ALGORITHM

In the current section the feedback algorithm is described. A controlled system must measure a parameter indicative of the system's performance and then exert control over the system. In computer networks an indicator of network performance is the delay in getting a packet through. It was found that collisions per millisecond was a good indication of the packet delay, and therefore gave an indication of the amount of traffic on the network. Figure 2 shows that the collisions per millisecond increased at the same rate as the delay. Collisions per millisecond is easy to measure by counting the jam signals.

It was found that the rate can be changed using

$$\text{rate} = \text{ravg} + q\ (Q - \text{colpms}) \quad (2)$$

where,

rate = the new voice coding rates
ravg = the average voice coding rate = 33000 bps
q = multiplier = 13000 bps
Q = average rate of collisions per millisecond = 3.3
colpms = collisions per millisecond.

The values for the three parameters, ravg, q, and Q were chosen by observing the non-multirate network and by experience. The value for the multiplier, q, was chosen to insure that a small change in the (Q-colpms) term will cause a large change in the rate. In this study colpms was computed over 32 ms intervals.

The (Q-colpms) term of Equation 2 represents an error signal. When the colpms is greater than the average value the rate is driven down. When the error signal is positive the rate is driven higher. When the error signal is zero, colpms is equal to the average, the rate obtained is the average voice coding rate.

The voice coding rate obtained from Equation 2 must be quantized to one of the four possible values given in Table 1. The quantizing intervals for the coding rates are given in Table 2.

Rate Obtained From the Feedback Equation	Rate is Truncated to
44 < rate	48.0
36 < rate < 44	40.0
28 < rate < 36	32.0
rate < 28	24.0

Table 2. Truncation of the rate obtained from the feedback equation.

4. PROTOCOL MODIFICATIONS FOR MULTIRATE VOICE AND DATA

The standard CSMA/CD protocol must be modified slightly for voice packets and multirate techniques. See Figure 3. The voice packets can be discarded for two reasons (1) a voice packet experiences excess collisions within one generation period or (2) the packet lifetime (one generation time) is exceeded. The data packets are discarded only if too many collisions occur. If a voice or data packet experienced more than 16 collisions, the packet was discarded. In addition, if a voice packet was not transmitted in one generation period then the packet was discarded.

5. NETWORK PERFORMANCE

The performance of an IVD CSMA/CD network using multirate voice coding was measured using simulation techniques. A CSMA/CD simulation was constructed using SLAM [6] and validated [7]. Figure 4 compares our simulation prediction with those of Shoch [8] while Figure 5 compares Hughes and Li [9] results on packet delay to those predicted by the SLAM simulation. In both cases the SLAM simulation was found to provide a reasonable prediction of the network performance.

The network configuration used to obtain the simulation results is given below in Table 3.

line capacity:	1 mega-bit-per-second
bus length:	approximately 1 kilometer (4.5 microseconds)
slot time:	round-trip bus delay = 9.0 microseconds
backoff algorithm for data:	binary exponential backoff truncated at 2^{10}
backoff algorithm for voice:	binary exponential backoff truncated at 2^{9}
jam time:	4.8 microseconds
interframe spacing:	9.6 microseconds
data packet length:	4096 bits
voice packet length:	768 bits

five data modes generating Poisson traffic

Table 3. Network Configuration

The multirate voice coding system was extensively tested using the simulation model. Results are presented from three cases: 1) constant data traffic load and coding rate, 2) constant data traffic load and multirate voice coding and 3) bursty (random) data traffic load and multirate voice.

The percentage of lost voice packets is an important performance indicator for integrated voice/data network. It gives an indication of the feasibility in using packet voice on a particular network. It has been shown that the voice quality degrades rapidly when the percentage of lost voice packets exceeds the 2% level [10]. In Figure 6 the percentage of lost voice packets versus the number of simulation conversations is shown for the multirate coding case and the non-multirate coding case. When the traffic on the network increases, the voice coding rate will be decreased. A decrease in the voice coding will cause less traffic on the network, and therefore reduces the percentage of lost voice packets at high loads.

The result shows a substantial increase in the number of simulated coversations over the case where no multirate voice coding was used. The 2% level of lost voice packets when no multirate coding was used is at approximately 12 conversations, and that with multirate coding was approximately 22 conversations. A comparison of the constant coding rate case to the results given in [1,2,3] shows good agreement. The curve of the lost packet percentage for the multirate case is not smooth due to the coding rate changes. When the number of simulated conversations increases, the percentage of lost voice packets is expected to increase. However, the percentage of lost packets drops from 1.9 when 15 conversations are simulated to 1.3 when 20 conversations are simulated. This occurs because the coding rate decreases when the number of conversations increases, which causes fewer voice packets to access the network. Thereby, decreasing the percentage of lost packets, allowing a greater number of conversations to take place.

The voice delay and the collisions per millisecond are shown in Figure 7. Here the collisions per millisecond increased with the voice packet delay, which verified that the collisions per millisecond was giving a good indication of the network traffic. The voice delay decreased from 9.5 milliseconds when 25 conversations were simulated, to 8.5 milliseconds when 30 conversations were simulated. The delay should increase as the number of conversations increases, however, this figure does not show the number of lost packets. When a packet was lost the delay statistics were not collected for that particular packet. So, at the higher loads more packets were lost and the packets that were discarded were not included in the delay statistics.

Throughput is another important network performance indicator. Throughput is an indication of network utilization. Throughput was calculated as the time the network was successfully carrying packets divided by the total time that was simulated. A comparison of throughput for the non-multirate case and the multirate case is shown in Figure 8. When the load was high, the multirate case gives a lower total throughput than the non-multirate case. This occurred because fewer voice packets attempted transmission at high traffic loads.

The segmental signal-to-noise ratio gives an indication of the voice quality [11]. The segmental signal-to-noise for the non-multirate case using a 48 kbps rate is 30 dB. The voice quality, and therefore the segmental signal-to-noise should drop as the voice coding rate decreases. The voice coding rate should drop down to the lowest rate as the number of conversations increases. The voice coding rate and segmental signal-to-noise are shown in Figure 9, where the expected result was obtained.

Data packet delay is another performance indicator. Data delay should remain as low as possible from the point of view of a network user [12]. A comparison of data packet delay for the constant coding rate case and the multirate coding case is shown in Figure 10. As seen in Figure 10, the data packet delay is decreased when the multirate voice coding techniques are employed.

Network performance in the presence of bursty traffic was also investigated. To model this we forced the data load to vary between 5% and 25% in 5% increments, with an average of 15%. The data load was changed every five seconds. It was found [7] that there was little difference in performance between constant and bursty traffic in our model. Figure 11 shows the segmental signal to noise ratio versus the number of conversations. There was not a significant difference between the constant and random data load cases. The time response of the system was useful for verifying that the feedback algorithm was operating properly. Figure 12 a and b illustrate the time response of the system. Figure 12a shows how the coding rate changes versus a function of time while Figure 12b shows the corresponding level of collisions per millisecond. In this simulation experiment the data load changed from 25% to 5% at 5 seconds. It is clear that the feedback system responded quickly to the change in data load. Note that at the high data load the voice coding rate changed rapidly. It is expected that the embedded voice coder performance will not be degraded due to these variations; however, this has yet to be verified.

6. CONCLUSION

A LAN configuration that has a combined loading of voice and data has been evaluated. The objective of the study was to determine the feasibility in using a multirate voice coding scheme to control the network load.

The multirate voice coding system performed well under varying load conditions. A comparison was made between the multirate voice coding system and a voice/data network that used a constant voice coding rate. The results of that comparison show that the multirate system gives an improved performance. Specifically, an increase of ten conversations was achieved when the variable rate coding scheme was applied to a system that consisted of a 1 Mbps line and a constant data load of 15%.

In addition, the constant data load and random data load cases were compared. This comparison showed that the constant load case closely approximated the random load case. Of more significance, the random load case simulated the bursty traffic associated with data. Therefore, the multirate coding techniques performed well under bursty traffic. It was demonstrated from a network perspective, e.g., lost packets, delay, etc., that multirate voice coding improved the system performance. However, further study is needed to assess the voice quality of a CSMA/CD network which used multirate voice coding.

7. REFERENCES

[1] J. D. DeTreville, "A Simulation-Based Comparison of Voice Transmission on CSMA/CD Networks and on Token Buses," AT&T Bell Laboratories Technical Journal, Vol. 63, No. 1, January 1984, pp. 33-55.

[2] G. J. Nutt and D. L. Bayer, "Performance of CSMA/CD Networks Under Combined Voice and Data Loads," IEEE Trans. on Communications, Vol. COM-30, No. 1, January 1982, pp. 6-11.

[3] J. M. Musser, T. T. Liu, L. Li, and G. J. Boggs, "A Local Area Network as a Telephone Subscriber Loop," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 6, December 1983, pp. 1046-1054.

[4] H. M. Hessestad, R. J. McAulay, and J. Tiernay, "Practical Considerations for Speech Digitizing Systems at Rates from 64.0 to 0.6 kbps," IEEE Global Telecommunications Conf., Miami, Florida, November 1982, pp. 349-356.

[5] D. J. Goodman, "Embedded DPCM for Variable Bit Rate Transmission," IEEE Trans. on Commun., Vol. COM-28, No. 7, July 1980, pp. 1040-1046.

[6] A. Alan and B. Pritsker, "Introduction to Simulation and SLAM II," Wiley, 1984.

[7] V. S. Frost and E. M. Friedman, "Computer Simulation Model for a CSMA/CD Local Area Network Utilizing a Multirate Voice Coding Technique for Load Control," Telecommunications and Information Sciences Laboratory Technical Report TR-TISL 6731-2, University of Kansas, May 1985.

[8] J. F. Shoch and J. A. Hupp, "Measured Performance of an Ethernet Local Network," Comm of the ACM, Vol. 23, No. 12, Dec. 1980, pp. 711-721.

[9] Herman D. Hughes and Liang Li, "A Simulation Model of the Ethernet," Dept. of Computer Science, Michigan State University East Lansing, Technical Report TR 82-008, 1982.

[10] N.S. Jayant and S. W. Christensen, "Effects of Packet Losses in Waveform Coded Speech and Improvements Due to an Odd Even Sample-Interpolation Procedure," IEEE Trans. on Communications, Vol. COM-29, No. 2, February 1981, pp. 101-109.

[11] N. S. Jayant and P. Noll, "Digital Coding of Waveforms," Prentice-Hall, 1984.

[12] M. Ilyas and H. T. Mouftah, "Performance Evaluation of Computer Communications Networks," IEEE Communications Magazine, Vol. 23, No. 4, April 1985, pp. 18-29.

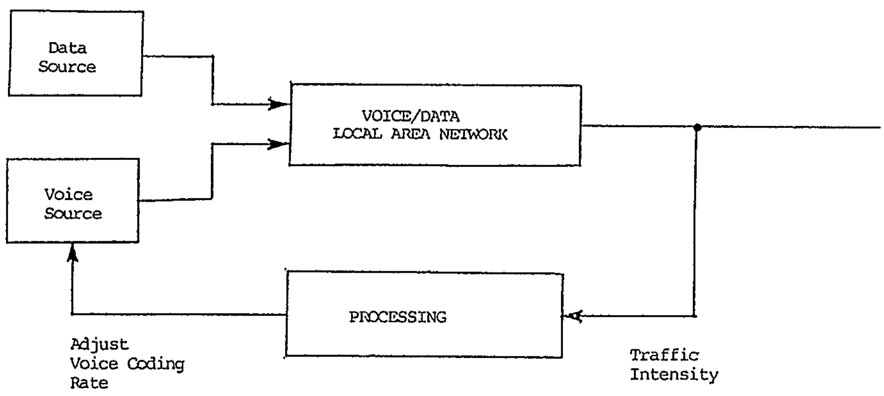

Figure 1. Block diagram of a voice/data Local Area Network where the voice coding rate is adjusted according to the traffic intensity on the network.

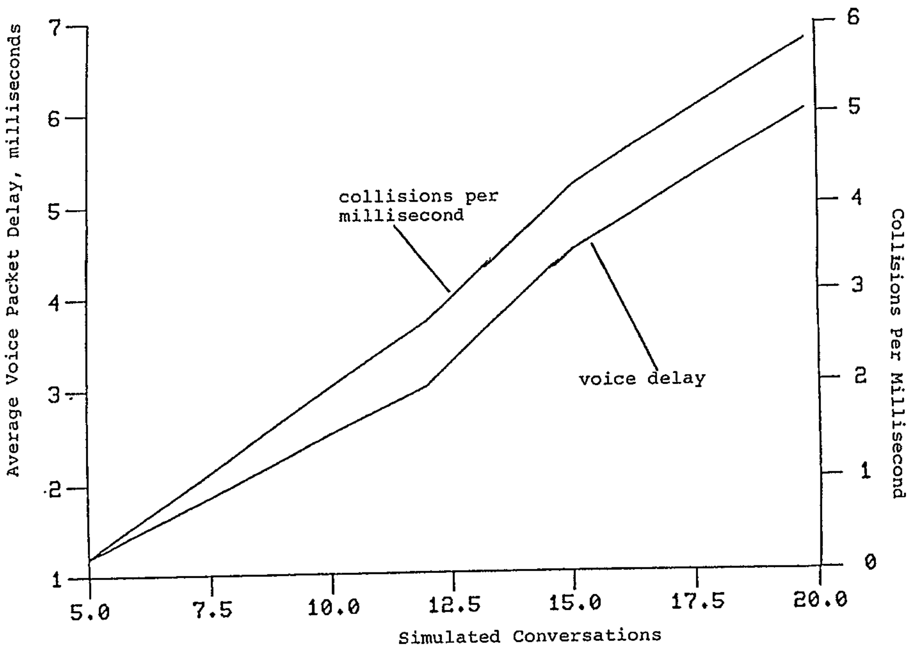

Figure 2. Comparison of the voice packet delay in milliseconds to the average rate of collisions per millisecond, with a constant voice coding rate of 48 Kbps, and a constant data load of 15%.

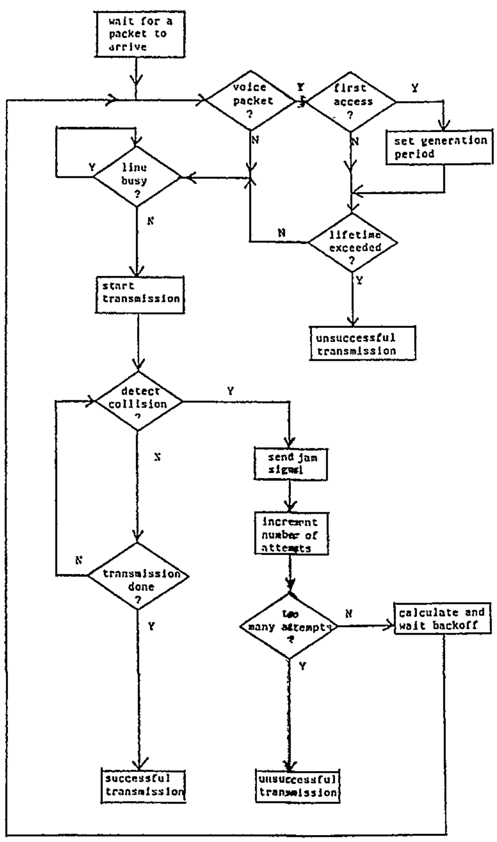

Figure 3. The decision logic for the CMSA/CD protocol that has been modified for multirate voice.

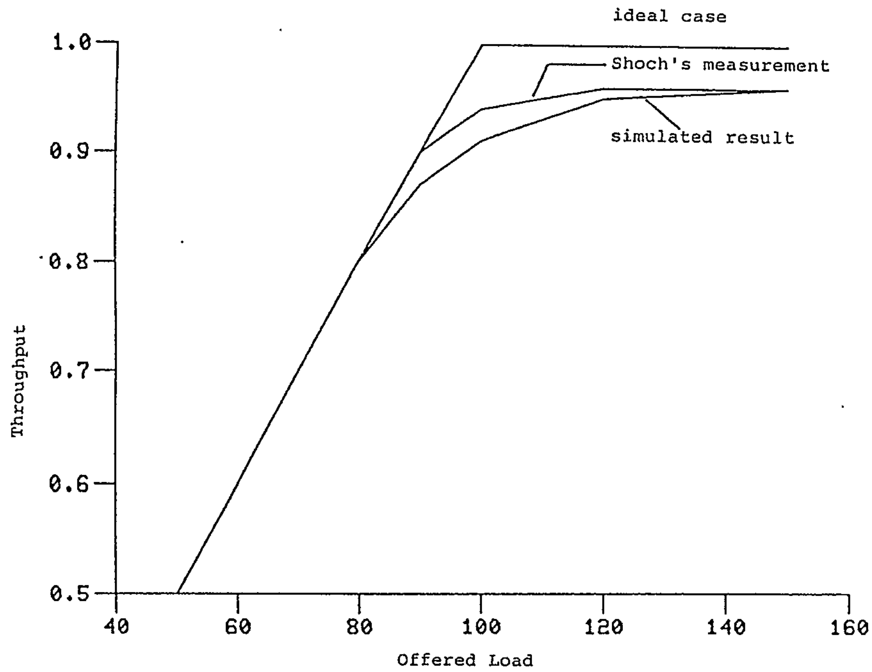

Figure 4. Comparison of the simulated throughput to Shoch's measurements, under various traffic loads and packet length of 4096 bits.

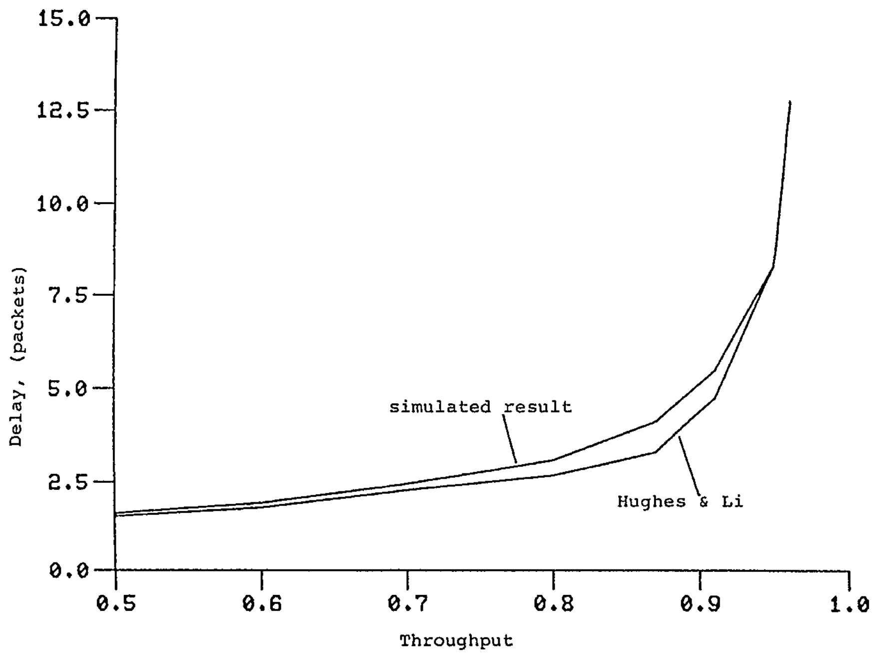

Figure 5. Comparison of the simulated average system delay to the delay reported by Hughes and Li, under various traffic loads and a packet length of 4096 bits.

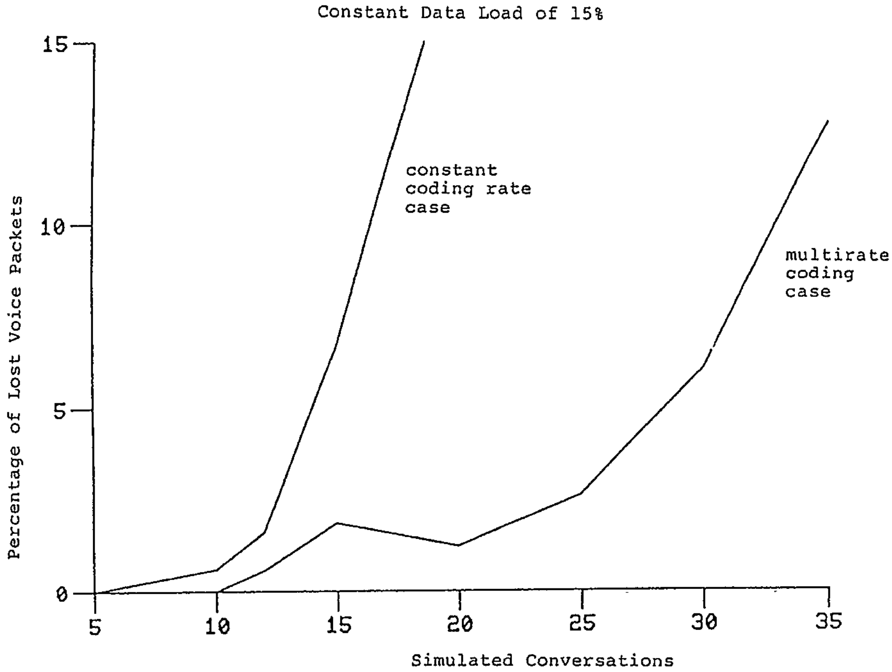

Figure 6. Comparison of the percentage of lost voice packets for the multirate and the non-multirate cases.

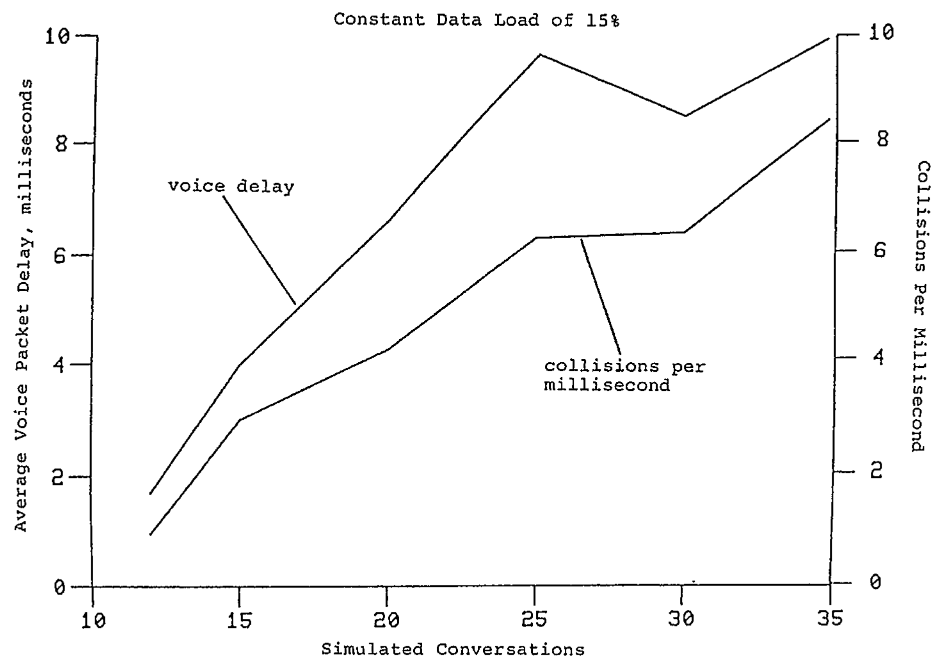

Figure 7. Rate of collisions per millisecond and the average voice packet system delay versus the number of simulated conversations.

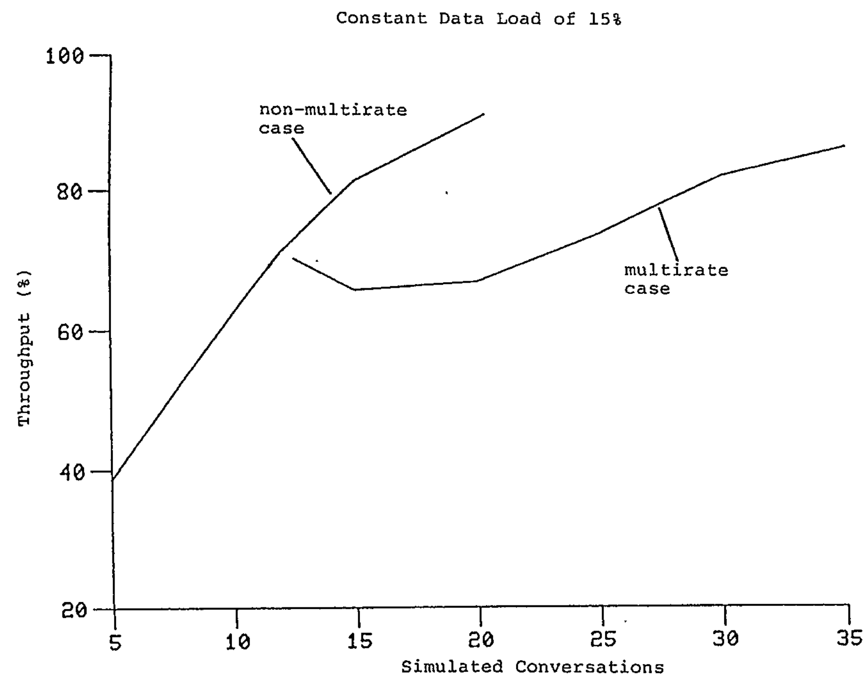

Figure 8. Comparison of throughput for the multirate and nonmultirate cases.

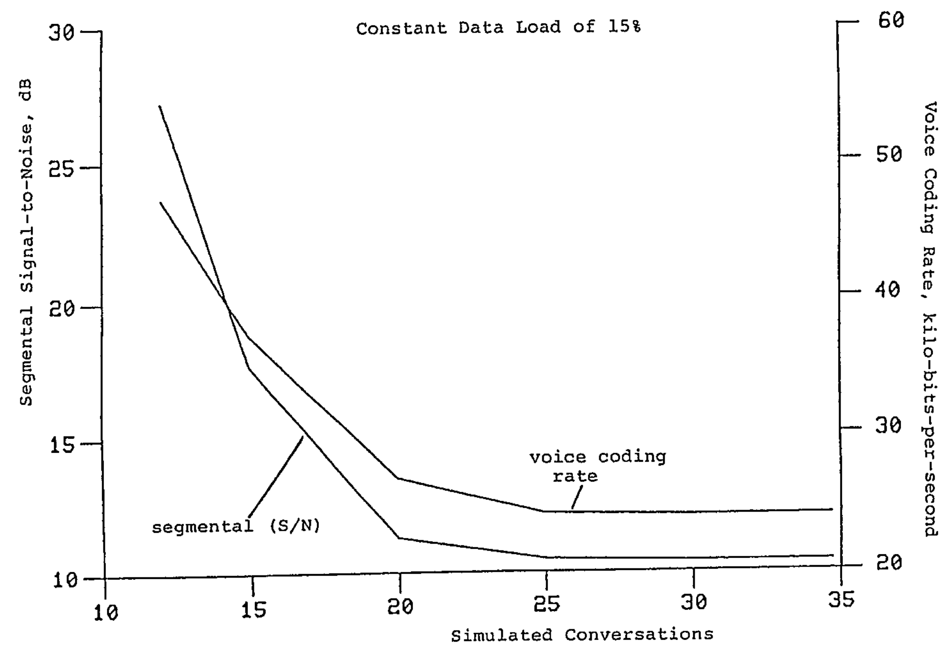

Figure 9. Comparison of voice coding rate in kilo-bits-per-second and the segmental signal-to-noise in decibels.

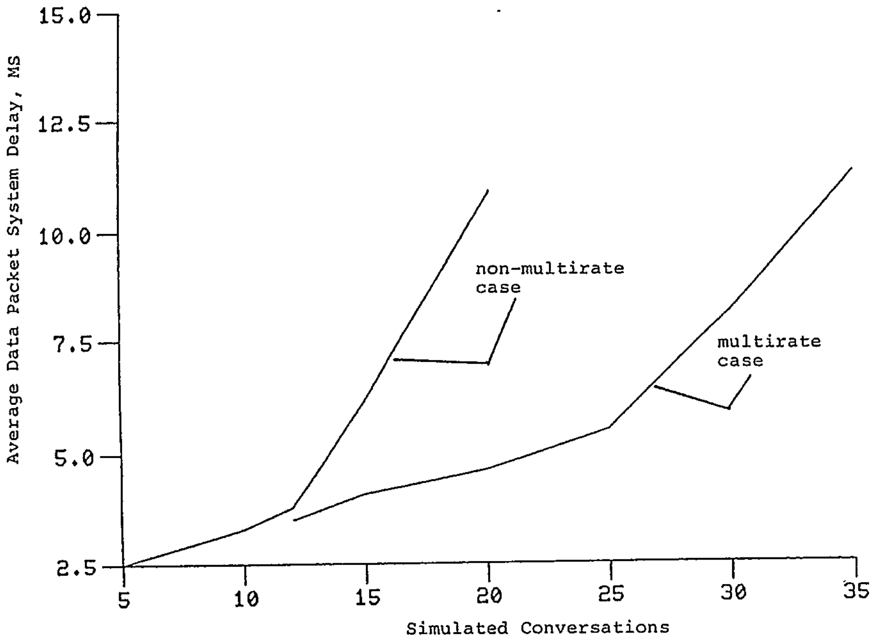

Figure 10. Comparison of the total system delay of data packets for the multirate and nonmultirate voice coding cases.

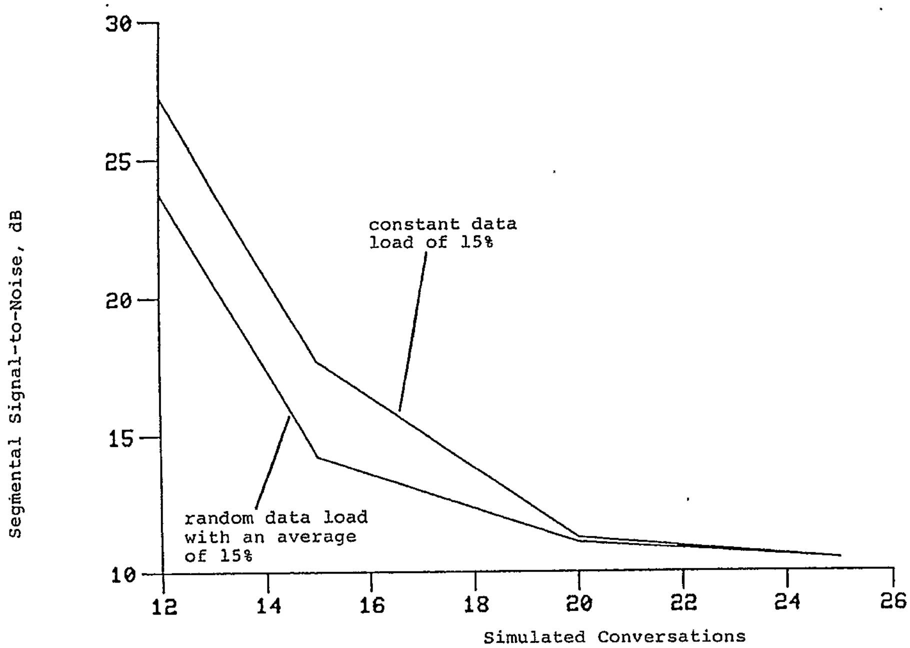

Figure 11. Comparison of the segmental signal-to-noise for a constant and random data load.

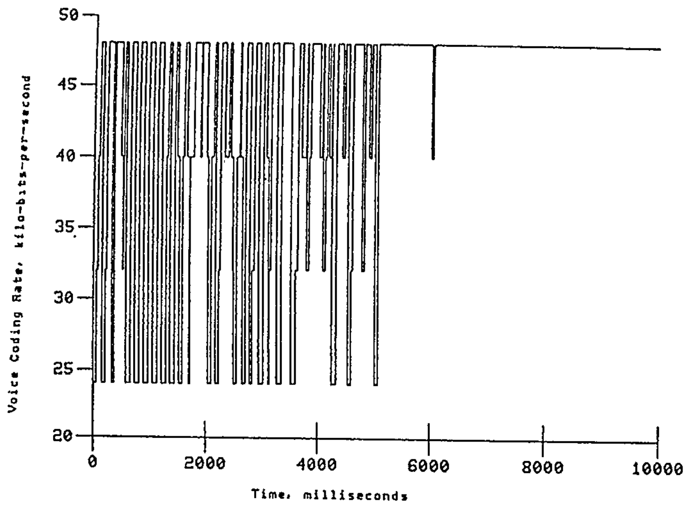

Figure 12a. Voice coding rate versus time, for 12 simulated conversations and a data load that changes from 25% to 5% at 5000 milliseconds.

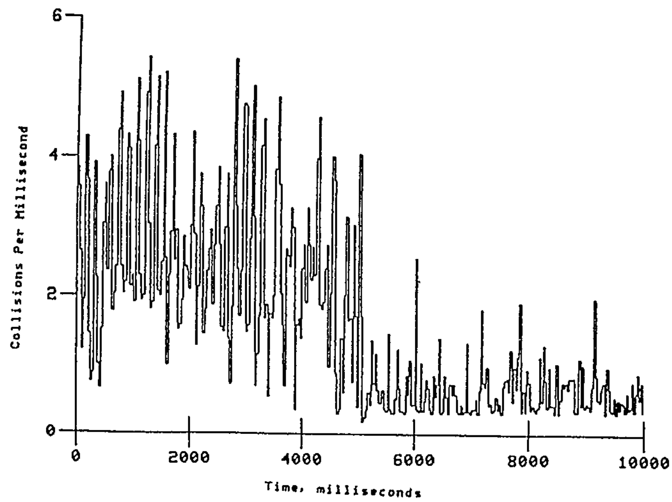

Figure 12b. Collisions per millisecond versus time, for 12 simulated conversations and a data load that changes from 25% to 5% at 5000 milliseconds.