(12)

ISI Reprint Series
ISI/RS-84-136
June 1984

University of Southern California

Danny Cohen

Satellite Communication of Real-Time Packet Video Images

Preprinted from the Proceedings of the Seventh International Conference on Computer Communication, Sydney, Australia, 30 October–2 November 1984.

INFORMATION SCIENCES INSTITUTE ISI
213/822-1511
4676 Admiralty Way/Marina del Rey/California 90292-6695

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ISI/RS-84-136	2. GOVT ACCESSION NO. AD-A243 623	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Satellite Communication of Real-Time Packet Video Images		5. TYPE OF REPORT & PERIOD COVERED Research Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Danny Cohen		8. CONTRACT OR GRANT NUMBER(s) MDA903 81 C 0335
9. PERFORMING ORGANIZATION NAME AND ADDRESS USC/Information Sciences Institute 4676 Admiralty Way Marina del Rey, CA 90292-6695		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209		12. REPORT DATE June 1984
		13. NUMBER OF PAGES 14
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

This document is approved for public release; distribution is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

....

18. SUPPLEMENTARY NOTES

Preprinted from the Proceedings of the Seventh International Conference on Computer Communication, Sydney, Australia, 30 October-2 November 1984.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

communication protocols, packet communication, packet satellite communication, packet video, satellite communication, video communication, video compression, WBNET, wideband communication

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

(OVER)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. ABSTRACT (continued)

The paper describes major issues and problems in the use of packet satellite channels for real-time video images and suggests some solutions. The DARPA Wide Band Communication Network provides the testbed for this activity.

A major idea in this paper is the automatic tradeoff between the quality and the dynamics of the communicated images in order to optimally adapt to the changes of the channel performance and to the changes of the scene to be communicated.

Extensive work has been performed in the past both on packet satellite and on real-time digital video. The paper reports work which uses both of these technologies to achieve their marriage in the form of packet satellite real-time video communication. The paper suggests several techniques and describes some new approaches to encoding real-time video communication.

Section 2 describes issues which are in the application of packet satellite channels for real-time video communication. Section 3 shows a general approach which addresses these issues. Section 4 describes a general methodology for signal coding and compression, and shows how it may take advantage of time coherence by using time differences while reducing quantization errors at the same time. Section 5 shows a technique for encoding dynamic values, suited for the automatic tradeoff of motion (dynamics) and sharpness.

As a "public service" this paper is accompanied by list of references regarding packet satellite channels. This list was provided by Vint Cerf (then of DARPA).

ISI Reprint Series
ISI/RS-84-136
June 1984

University of Southern California

Danny Cohen

Satellite Communication of Real-Time Packet Video Images

Preprinted from the Proceedings of the Seventh International Conference on Computer Communication, Sydney, Australia, 30 October–2 November 1984.

INFORMATION SCIENCES INSTITUTE

213/822-1511
4676 Admiralty Way/Marina del Rey/California 90292-6695

This research is supported by the Defense Advanced Research Projects Agency under Contract No. MDA903 81 C 0335. Views and conclusions contained in this report are the author's and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them.

ISI Reprint Series

This report is one in a series of reprints of articles and papers written by ISI research staff and published in professional journals and conference proceedings. For a complete list of ISI reports, write to

Document Distribution
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
USA

SATELLITE COMMUNICATION OF REAL-TIME PACKET VIDEO IMAGES[1]

Danny Cohen
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
USA

Abstract

The paper describes major issues and problems in the use of packet satellite channels for real-time video images and suggests some solutions. The DARPA Wide Band Communication Network provides the testbed for this activity.

A major idea in this paper is the automatic tradeoff between the quality and the dynamics of the communicated images in order to optimally adapt to the changes of the channel performance and to the changes of the scene to be communicated.

Extensive work has been performed in the past both on packet-satellite and on real-time digital video. The paper reports work which uses both of these technologies to achieve their marriage in the form of packet satellite real-time video communication. The paper suggests several techniques and describes some new approaches to encoding real-time video communication.

Section 2 describes issues which are in the application of packet satellite channels for real-time video communication. Section 3 shows a general approach which addresses these issues. Section 4 describes a general methodology for signal coding and compression, and shows how it may take advantage of time coherence by using time differences while reducing quantization errors at the same time. Section 5 shows a technique for encoding dynamic values, suited for the automatic tradeoff of motion (dynamics) and sharpness.

As a "public service" this paper is accompanied by list of references regarding packet satellite channels. This list was provided by Vint Cerf (then of ARPA).

1. Introduction

Recent advances in packet communication over satellite links have now made it possible, for the first time, to seriously consider the application of this technology to real-time video transmission. DARPA's Wide Band Network (WBNET) supports up to 3.0 Mb/s packet communication of data and real-time voice among many computers. This network provides a testbed for real-time packet video experiments.

The DARPA Wide Band Packet Satellite system is a USA-domestic satellite network which uses a 3.0 Mb/s channel of the Western Union Westar geosynchronous satellite. The network currently uses four ground stations to share the common wide band channel dynamically on a packet-by-packet basis. The network is based on principles developed and tested in an earlier 64 kb/s Atlantic Packet Satellite system (SATNET) which is still in operation serving DARPA's international experiments in command and control systems technology.

[1] This work was sponsored by the Defense Advanced Research Projects Agency (DARPA) under Contract No. MDA903-81-C-0335. Views and conclusions contained in this report are the author's and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them

The basic strategy in packet satellite systems developed by DARPA has been to support rapid reallocation of common (shared) satellite capacity among an arbitrarily large number of ground stations. To achieve this capability to adapt rapidly to varying demand, the system uses a contention, priority-oriented, demand allocation (CPODA) technique which allows each ground station to contend for access and use of the channel. All data entering the system is packetized by the computers which act as sources and sinks of traffic. Both "stream" and "datagram" services are supported. Datagram service accepts a single packet from the source computer and schedules it for transmission independently of all other traffic from the the source computer. Stream service permits a source host to request a certain fixed capacity for a given period of time. All packets presented to the ground station marked as associated with that stream are sent during the period reserved by all ground stations for the particular source ground station serving the source host.

The WBNET is a broadcast system; all transmissions are received by all ground stations, including the sending (barring local interference). As a consequence, it is possible for all ground stations to keep track of the demand for satellite capacity and to independently but cooperatively schedule access to the channel according to the requirements posted by each traffic source.

Designed initially to support large numbers of widely dispersed computer systems, the packet satellite technology has also been adapted to the support of point-to-point and conferenced voice traffic. The Wide Band Network is an experiment intended to explore the effectiveness of using packet satellite techniques to support combined and dynamically varying voice and data communication requirements.

The communication of real-time video over the WBNET deals with problems that are caused by several factors: the need to digitize the video signal, the further need to packetize it, compounded with the general characteristics of satellite channels and with the particularities of the WBNET.

Extensive work on digital video has been reported in the literature. This paper is not aimed at the issue of digitizing and compression. It uses known and proven video encoding techniques, especially the popular Discrete Fourier Transform.

Packet communication over satellite channels poses several problems for real-time video applications: The data-rate required for real-time video applications (above 55 Mb/s) is well above the data capacity on packet satellite channels available to us (0.75 to 3.0 Mb/s in our case). Furthermore, this limited capacity has to be shared dynamically with many other users with unpredictable demands. The packetization of the video bit-stream introduces several issues that are similar to those in packet speech. They deal with the variance in network performance as encountered by each packet, such as the variable-delay, ranging from an out-of-order arrival to a total loss of packets.

The delay associated with any use of a geosynchronous satellite is large relative to terrestrial delays (above 230 milliseconds); the error characteristics vary as a result of many uncontrollable external factors. This delay rules out retransmission of packets for real-time communication, either by timeout or by request (e.g., NACK).

The idiosyncracies of the WBNET, such as protocols and packet size limitations, add to the issues that have to be addressed.

One of the major design goals of our approach to packet image communication is to achieve a low data rate over the channel, while coping with its variable performance, as described in [1]. This approach is based on the communication of encoded transforms of the original images, and on performing the inverse operation at the receiver.

The term "dynamics" throughout the paper refers only to any changes with respect to time and not to the dynamic range of the images.

2. Issues in Packet Satellite Communication of Video Images

Packet satellite channels pose some challenges for real-time packet video communication. First and foremost is performance, as measured by data-rate (or capacity or throughput), delay, error characteristics, and the great variation of these features.

Since the data-rate available over the channel is not enough to meet the requirements of the video encoding scheme, special measures have to be taken to compress the data and to optimize the utility of the available capacity. These measures imply some compromises. In our case, the total capacity of the WBNET packet satellite channel, which has to be shared among all of its users, is only 0.75 to 3.0 Mb/s. Therefore, we cannot use any of the standard digital video communication techniques without using special measures and compromises.

The available data-rate depends on the total utilization of the channel which is a function of the total demand of all the simultaneous users. Since computer applications tend be very bursty and to vary rapidly their communication demands, the amount of available data-rate may vary suddenly, without any prior warning.

The long inherent delay (above 230 milliseconds) due to the high orbit of the geosynchronous satellite reduces the system capability to respond quickly to sudden changes, which are relatively frequent in computer communication. It is more important to respond properly and quickly to a reduction in data-rate than to an increase.

The error characteristics of the channel depend on many environmental factors and also may change rapidly. The channel errors fall into two categories, Bit Error Rate (BER) and loss of entire packets. For a certain cost (in data-rate) data can be protected to achieve a lower BER. It is assumed that the system builder already has protected the headers of packets to minimize packet loss due to damaged headers in general and damaged addresses in particular.

The two major factors which increase the error rate are weather and electromagnetic emissions by other non-cooperative systems, like radars and neighboring microwave systems. Experience shows that these factors do not change very rapidly. Any scheme which monitors the error rate should implement fast response to detected degradations, while possibly responding more slowly to detected improvements.

3. Approach to Encoding Real-Time Video Images

Our approach to encoding real-time video information is based on:

1. The observation that the required total data rate for this communication is the product of the spatial (xy) resolution, the shade/color (z) resolution, and the temporal (t) resolution,

2. The assumption that the nature and the dynamics of the scene vary,

3. The need to cope with sudden changes of channel performance (capacity, delays, BER, etc.), and

4. The assumption that both the sender and the receiver have a certain non-trivial amount of processing power and storage.

Our strategy is both dynamic and adaptive, in the sense that it varies according to conditions observed in real-time, such as the dynamic nature of the scene and of the current load and performance of the channel.

If the available capacity is less than what is needed for communicating the full xy, z, and t resolutions, the system must (automatically or manually) decide what to sacrifice. It is well known that for fast motion the spatial and the shade/color resolutions are secondary to the temporal resolution, while the inverse is true for static images. However, any sacrifice must not necessarily be along the xy, z and t boundaries but may take other dimensions such as loss of sharpness.

Conventional encoded communication, performed according to a fixed algorithm (such as PCM, CVSD, and LPC) requires the transmission of control information to manage the connection, and of data to convey the signal.

Packet video communication may use schemes which, in addition to sending connection control and image data, send also "functions" telling the receiver what to do with the data. Therefore, the entire encoding algorithm can be dynamically changed if the sender finds it optimal to do so, typically as a result of scene type change. This includes changes of either entire algorithms or just of parameters like the granularity of coding tables or the block sizes.

It is possible to send successive approximations (or iterations) which lead the receiver to the desired image. If these successive approximations are marked with decreasing priority, then in case of a sudden decrease in channel performance the received images are more likely to suffer from quality degradation only - rather than from total loss of image parts. This may happen because of common implementation strategies which automatically drop packets of lower priority when the total demand exceeds the available capacity.

When such successive approximations are used, it is imperative to ensure that high priority packets are successfully delivered. When such a packet is lost or severely damaged, the entire image may have to be discarded.

When fast motion is detected, the data may be sent using a lower spatial and/or shade resolution but a higher temporal resolution than usual, at the same time instructing the receiver to perform a spatial low-pass filter operation to "smear" the image in the way that human vision and TV cameras do.

The successive approximations may be along any dimension, or along any combination of several dimensions. This includes dimensions from the image domain (xy, z, and t) or dimensions from any transformation domain such as the one- and two-dimensional Discrete Cosine Transform (DCT).

Consider, for example, the communication of an image of size xyz = 512x512x8 bits. One way to send the image is in the zxy-order: sending first all the bits (z) for each pixel from the first pixel in the first row, through the second pixel in the first row (x), through all the bits of the last pixel in the last row (y). This is probably the simplest way of communicating an image.

Another possibility is to use xyz-order, where the most significant bit of every pixel is sent first, then the second most significant bit of every pixel, and so on to the least significant bit of every pixel.

The transmission of all these bits requires the use of many packets. In the latter scheme the priority of the packets should decrease with the significance of the corresponding bits, such that if the channel performance degrades suddenly with no advance warning, some of the least significant packets could be discarded while all the more significant ones would arrive safely.

Similarly, if the error characteristics of the channel are degraded, then more significant packets can be marked for better (hence more data intensive) error protection/correction coding while the less significant ones are left with less protection (or none at all), or may even be sacrificed entirely in order to afford the added data rate needed for the protection of the more significant packets.

If 512x512 bit packets are still too big for reasonable communications they can be divided into smaller packets, either by covering smaller areas, or by having each packet cover the entire 512x512 area but in a lower spatial resolution, similar to memory interleaving schemes.

It is best to have the successive approximations converge to the target image not at a uniform rate, but at a rate which starts high and decreases later, such that the first approximations convey "most" of the information and the later ones serve to enhance it. This iterative process should look like focusing a lens, where the entire image is transformed from a low quality image into a high quality image. The above xyz-order scheme has this property.

The purpose of using such schemes is to be able to react dynamically to performance changes, both in the available data-rate and in the error characteristics.

4. Signal Encoding and Compression

Let $C = S(X)$ be the encoding of the signal X, and let $Y = S'(C) = S'S(X) = H(X)$ be the signal reconstructed from the code C. The distortion caused by the coding is $X-H(X) = [I-H]X$, which is not as important as the difference between the perception of the original signal and the perception of the reconstructed signal. If the distortion, $[I-H]$, is in the "null-space" of the perception then the coding causes no perceived distortion.

The degree of freedom exploited by the compression is the deliberate use of the $[I-H]$ distortion if its effect on the perception is small, e.g., loss of low energy high frequency signal components.

An efficient coding approach is to transform the signal into another space (or domain) whose components have known perceptual significance, and to allocate code bits to each of these dimensions according to its importance to the perception process, and to the signal statistics (a la Huffman coding).

In many cases, the frequency analysis (temporal or spatial) of the signals suggests the use of block coding. In the case of speech these blocks are usually intervals of about 20 milliseconds, whereas for video images the most interesting blocks are of sizes 8x8, 8x16 and 16x16 pixels.

From a pure information-theory point of view, the space of the coded representation cannot be of a lesser dimensionality than the original time domain representation (i.e., the block size) of a general signal without loss of information.

However, two factors make signal compression possible. First, the fact that a certain family of signals (rather than the most general signals) is being coded, and second, the high level of redundancy in the signal (or the forgiveness of the human perception).

If the signals are restricted to a certain family which can be modeled as the output of some process (or the range of some operator) then the dimensionality of the time domain representation is not as high as the block size but only as the rank of that operator. The inversion of this operator yields parameters which, when given to the process, reproduce the original signal.

In image communication, with a fixed block size, increasing the x resolution of an image usually yields lower spatial frequencies and higher pixel-to-pixel correlation. Similarly higher y resolution yields higher line-to-line correlation, and higher t resolution yields higher frame-to-frame correlation. The higher these correlations are the better the compression which may be devised. For example, when facsimile machines doubled their resolution from 192 ppi to 384 ppi no increase in the number of bits in the coded representation was required.

High correlations may be exploited either by simple minded differential schemes along the appropriate axis (e.g., run length) or by more complicated transforms into a frequency domain where high spatial correlation concentrates energy in the low spatial frequencies so that very few bits are needed for coding the higher frequencies.

Exploiting high frame-to-frame correlation by block transforms cannot be done without a great amount of storage (to hold several frames) and without additional delay (to wait for more frames) and processing. Therefore the simple minded differential one-step schemes which require the storage of only one frame appear to be a reasonable compromise between compression, delay, processing, and memory.

In images with static parts the compression may be significantly improved by using special block-oriented codes, like a single code to indicate an all-zero block.

Most of the images of interest may be "compressed" by using some spatial spectral tools (e.g., DFT and DCT) which are linear.

In many applications of interest there is high time-coherence which may be exploited by the use of incremental techniques which communicate the differences between successive images.

A little known interesting mathematical property is that the linear spatial transformation and the time differentiation operators commute.

Therefore, it is possible to devise several communication schemes based on communicating values of the signal, of the differential signal, of the transformed signal, of the transformed differential signal, etc.

All of these techniques suffer from the quantization error caused by the need to limit the communication bandwidth. We alleviate this problem by redefining the time differences. Instead of computing time differences as the difference between the current signal and its previous value, as seen by the sender, the differences are computed between the most current signal as available at the sender and the previous signal value as received by the receiver.

The digital transmission with its error control mechanisms eliminates the traditional channel noise, but introduces quantization errors, which may be easily computed by the sender and be corrected for.

Figure 1 shows a scheme where the sender estimates the receiver's quantization errors and corrects for them.

The coupling of this system with coding techniques such as the one described in the next section allows the system to converge fast and coarsely to dynamic images and in fine resolution to slow images.

Hence, static images are always communicated with the maximal resolution available to the system whereas moving objects are communicated in more coarse resolution.

Note how well this approach fits the visual perception process since the eyes (and TV cameras) provide blurry images of moving objects and sharp images of stationary objects only.

This method, like all other incremental methods, is very sensitive to loss of packets, and must take steps to recover from such losses, e.g., by a periodic communication of absolute, rather than incremental values.

5. Coding Dynamic Quantities

This section discusses a compression and coding scheme for general (random) data. When it is used in a differential way it always converges perfectly to static values at the expense of coarser approximations to rapidly changing quantities.

Let X be a single quantity which has to be communicated. The standard coding technique is for the sender to encode X, i.e., map X into a code $C = S(X)$, to send the resulting code, C, and for the receiver to decode C to get $Y = S'(C)$.

If the data rate available over the channel is sufficient, then the internal representation of X may also be used as its code, $S = I$, yielding an encoding which loses no information. If the receiver shares the same presentation then $S' = I$ and the communication process introduces no distortion.

If the code, C, has to have less bits than the original signal, X, then the encoding must lose some information which cannot be recovered by the decoder.

The S operator is a many-to-one mapping, such that entire intervals are assigned to the same code. The inverse operator, S', is a one-to-one mapping which assigns a unique "representative" value to the entire interval, represented by the code. Hence, the entire coding procedure, $H = [S'S]$, is a step function which assigns single values to intervals.

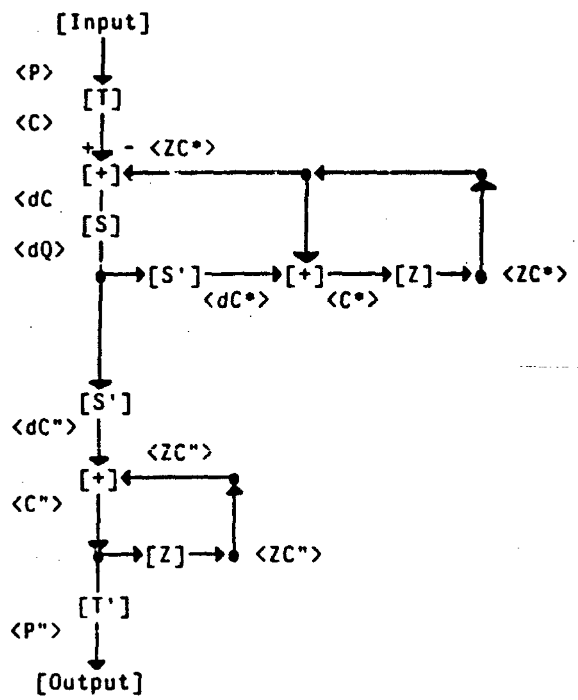

Figure 1: A Scheme Based on Adaptive-Differential Transformation

In cases where dynamic values have to be communicated the coding of the differences, instead of the raw data, should be considered.

If X keeps changing then the received value may be different, but when it slows down and stabilizes the received value catches up with it. Hence, blurred images of dynamic objects, and sharp images of static objects, which is a desired tradeoff.

In order to guarantee the convergence of the coding scheme for static values, it is necessary to have $0 < S'S(X)/X < 2$ for all values of X. This condition guarantees that when X stabilizes each successive iteration is closer to X than the previous one, and Y converges to X.

Proof: Define $H = S'S$ and $k(X) = H(X)/X$. The above condition is $0 < k < 2$.

Let Yn be the nth iteration for the static value X:

$$Y_n = Y_{n-1} + H(X - Y_{n-1})$$

and similarly:

$$Y_{n+1} = Y_n + H(X - Y_n) = Y_n + (X - Y_n)\, k(X - Y_n)$$

The new error is:

$$E_{n+1} = X - Y_{n+1} = X - Y_n - H(X - Y_n) =$$

$$= (1-k)\,(X - Y_n) = (1-k)\,E_n$$

$|1-k| < 1$ because $0 < k < 2$.

Hence $|E_{n+1}| < |E_n|$. [QED]

The closer k(X) is to +1 the faster the iterations converge. A good coding scheme must assure that the value of k is close to +1 for all X.

The minimum code requirement to assure convergence is a 3-valued code:

$$H(dX > 0) = +1$$
$$H(dX = 0) = 0$$
$$H(dX < 0) = -1$$

This code always converges, but the larger the value of X the slower the convergence due to the low value of $k = H(X)/X = 1/|X|$.

6. Conclusion

Conventional video compression techniques have to be augmented to cope with the special problems inherent to communication over packet satellite channels. However, with proper attention to the nature of these problems it is possible to support real-time video image communication by adapting dynamically to the variances in the channel performance, such that the system gracefully changes its behavior according to the available resources and to the dynamics of the scene.

It is most important to monitor dynamically the channel (available data-rate, total demand, BER, etc.) in order to optimize its use by performing the proper tradeoffs which depend also on the varying nature of the scene.

7. Acknowledgments

Randy Cole of ISI and Vint Cerf of MCI provided great help in clarifying these issues through many discussions. Steve Casner and Jon Postel, both of ISI, joined them in helping in the preparation of this paper by providing much needed critique.

Tom Strini and Sheila Coyazo helped to convert this paper into readable English.

I wish to thank all of them.

8. References

1. S.L. Casner, D. Cohen and E.R. Cole, "Issues in Packet Satellite Video Communication", ICC-83 proceeding.

2. L.G. Roberts and B.D. Wessler, "Computer Network Development to Achieve Resource Sharing", Spring Joint Computer Conference, AFIPS Proceedings, Vol. 36, 1970, pp. 543-549.

3. R.E. Kahn, "Resource Sharing Computer Networks", Proceedings of the IEEE, Vol. 60, No. 11, November 1972, pp. 1397-1407.

4. N. Abramson, "The Aloha System—Another Alternative for Computer Communications", AFIPS Conference Proceedings, Vol. 36, 1970, pp. 295-298.

5. F.E. Heart, et al., "The Interface Message Processor of the ARPA Computer Network, Spring Joint Computer Conference, AFIPS Proceedings, 1970, pp. 551-567.

6. R. Binder, et al., "Aloha Packet Broadcasting—a retrospect", AFIPS Conference Proceedings, National Computer Conference, 1975, pp. 203-215.

7. N. Abramson and F. Kuo, Editors, Computer Communication Networks, Prentice Hall, Englewood Cliffs, N.J., 1973.

8. L. Kleinrock and S.S. Lam, "Packet Switching in a Slotted Satellite Channel", AFIPS Conference Proceedings, NCC, 1973, pp. 703-710.

9. L. Kleinrock and F. Tobagi, "Random Access Techniques for Data Transmission over Packet Switched Radio Channels", AFIPS Conference Proceedings, NCC, 1975, pp. 187-201.

10. L. Kleinrock and S.S. Lam, "Packet Switching in a Multiaccess Broadcast Channel: Performance Evaluation", IEEE Transactions on Communication, Vol. COM-23, 1975, pp. 410-423.

11. L.G. Roberts, "Aloha Packet System with and without Slots and Capture", ACM SIGCOMM, Computer Communication Review, Vol 5, No. 2, April 1975.

12. S.S. Lam and L. Kleinrock, "Packet Switching in a Multi-access Broadcast Channel: Dynamic Control Procedures", IEEE Transactions on Communication, Vol Com-23, September, 1975.

13. L.G. Roberts, "Dynamic Allocation of Satellite Capacity through Packet Reservation", AFIPS Conference Proceedings, NCC, 1973, pp. 711-716.

14. N. Abramson, "Packet Switching with Satellites", AFIPS Conference Proceedings, NCC, 1973, pp. 695-702.

15. R.E. Kahn, "The Introduction of Packet Satellite Communications", National Telecommunications Conference, Nov. 1979, pp. 45.1.1-8.

16. I.M. Jacobs, et al, "CPODA—A Demand Assignment Protocol for SATNET", Fifth Data Communications Symposium, Snowbird, Utah, 1977.

17. I.M. Jacobs, et.al, "General Purpose Satellite Networks", Proceedings IEEE, Vol 66, No. 11, November 1978, pp. 1448–1467.

18. I.M. Jacobs, et al, "Packet Satellite Network Design Issues", Proceedings, NTC, November 1979.

19. L. Palmer, J. Kaiser, S. Rothschild and D. Mills, "SATNET Packet Data Transmission", COMSAT Technical Review, Volume 12, No. 1, Spring 1982, pp. 181–212.

20. E. Killian and R. Binder, "Control Issues in a PODA Voice/Data Satellite Network", ICC, June 1980.

21. C. Heegaard, J. Heller and A. Viterbi, "A Microprocessor-based PSK Modem for Packet Transmission over Satellite Channels", IEEE Transactions on Communications, COM-26, No. 5, May 1978, pp. 552–564.

22. P. Cudhea, D. McNeill, D. Mills, "SATNET Operations", AIAA 9th Communications Satellite Systems Conference, Collection of Technical Papers, 1982, pp. 100–105.

23. D.A. McNeill, et al, "SATNET Monitoring and Control", Proceedings of the NTC, November 1979.

24. P.T. Kirstein, et al, "SATNET Applications Activities", Proceedings of the NTC, November 1979.

25. W.W. Chu, et al, "Experimental Results on the Packet Satellite Network", Proceedings of the NTC, November 1979.

26. E.V. Hoversten and H.L. Van Trees, "International Broadcast Packet Satellite Services", ICCC Conference Proceedings, Kyoto, Japan, September 1978.

27. V.G. Cerf and R.E. Kahn, "A Protocol for Packet Network Intercommunication", IEEE Trans. on Comm., Vol. COM-23, May 1974, pp. 637–648.

28. V.G. Cerf and P.T. Kirstein, "Issues in Packet Network Interconnection", IEEE Proceedings, Vol. 66, No. 11, Nov 1978, pp. 1386–1408.

29. L. Evenchik, D. McNeill, R.P. Rice, F. Deckelman, et al, "MATNET, an Experimental Navy Shipboard Satellite Communications Network, "IEEE INFOCOM 82 Proceedings, March, 1982.

30. M.L. Molle and L. Kleinrock, "Analysis of Concentrated ALOHA Satellite Links", Sixth Data Communications Symposium, Nov 27–29, 1979, pp. 87–95.

31. Weissler, et al, "Synchronization and Multiple Access Protocol in the Initial Satellite IMP", COMPCON, September 1978.

32. Hsu and Lee, 'Channel Scheduling Snychronization for the PODA Protocol", ICC, June 1978.