

AD-A135 942

ISSUES IN SATELLITE PACKET VIDEO COMMUNICATION(U)

1/1

UNIVERSITY OF SOUTHERN CALIFORNIA MARINA DEL REY

INFORMATION SCIENCES INST S L CASNER ET AL JUL 83

UNCLASSIFIED

ISI/R5-83-5 NDA903-81-C-0335

F/G 17/2


NL

END

FILED

1-44

DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A135942

13

ISI Keprint Series
ISI/RS-83-5
July 1983



Stephen L. Casner
Danny Cohen
E. Randolph Cole

Issues in Satellite Packet Video Communication

DTIC FILE COPY

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DTIC
ELECTE
DEC 16 1983
B

INFORMATION SCIENCES INSTITUTE

4676 Admiralty Way/Marina del Rey/California 90267
(213) 833-7111

UNIVERSITY OF SOUTHERN CALIFORNIA



83 12 16 027

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

DD FORM 1473
1 JAN 73

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

20. ABSTRACT (continued)

Recent advances in packet communication over satellite links make it now possible, for the first time, to seriously consider the application of this technology to real-time video transmission.

Unfortunately, packet communication over satellite links poses several severe problems for real-time video applications: The data rate available on packet satellite links is usually below what is required for real-time video applications; this data rate is shared dynamically with many other users with unpredictable demands; the delay associated with any use of a geosynchronous satellite is very large; and the error characteristics vary as a result of many uncontrollable external factors.

This paper discusses the problems, suggests several techniques to cope with them, and describes the application of these techniques in a real-time packet video communication system being implemented at ISI.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

ISI Reprint Series
ISI/RS-83-5
July 1983



Stephen L. Casner
Danny Cohen
E. Randolph Cole

Issues in Satellite Packet Video Communication

UNIVERSITY OF SOUTHERN CALIFORNIA



INFORMATION SCIENCES INSTITUTE

4676 Admiralty Way/Marina del Rey/California 90291
(213) 822-1511

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 authors' 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 copies of articles 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 90291
USA

Contents

INTRODUCTION	1
THE PROBLEMS	1
APPROACH	2
Procedure	3
IMPLEMENTATION	4
Selected Algorithm	5
Simulation on General-Purpose Hardware	5
Real-Time Implementation	6
IKONAS graphics system	7
Bandwidth compression processor	7
Packetization functions	7
Wideband Network	7
Video System Optimization	8
CONCLUSIONS	9
REFERENCES	10

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



ISSUES IN SATELLITE PACKET VIDEO COMMUNICATION

Stephen L. Casner, Danny Cohen, and E. Randolph Cole

INTRODUCTION

Research in packet-switched communication has led to the development of extensive terrestrial networks with moderate data rates. Geographic coverage has been broadened through the use of satellites. Now, a quantum increase in data rate is provided by DARPA's Wideband Satellite Network (WBNET) which supports up to 3 Mb/s packet communication for data and real-time digital voice among several computers. This network provides adequate capacity to serve as a testbed for real-time packet video experiments, although that capacity must be shared with other users.

The basic strategy in packet satellite systems developed by DARPA has been to support rapid re-allocation 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 ground stations receive the broadcast signal and keep track of the demand for satellite capacity 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 digital voice traffic. The WBNET is an experiment intended to explore the effectiveness of using packet satellite techniques to support combined and dynamically varying voice and data communication requirements.

This paper explores a number of important issues which arise in connection with using packet satellite techniques to support real-time video communication in addition to voice and data services. The problems are discussed, and an approach for accommodating them is presented. This approach is being applied in a real-time packet video communication system we are implementing at ISI. An overview of the implementation is given and its interaction with the WBNET is described.

THE PROBLEMS

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

If the data rate available over the channel is not enough to meet the requirements of the video digitization scheme (typically 50 to 100 Mb/s), then 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 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 to be very bursty with rapidly varying communication demands, the data rate available may change suddenly without any prior warning.

The long inherent delay (more than 230 milliseconds), due to the high orbit of the geosynchronous satellite, reduces the system's capability to respond quickly to sudden changes. 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. The two major factors which increase the error rate are weather and electromagnetic emissions by other non-cooperative systems. While the weather changes relatively slowly, stray EM emissions may occur abruptly and intermittently. Any scheme which monitors the error rate should implement fast response to detected degradations, while possibly responding more slowly to detected improvements.

The channel errors fall into two categories: random bit errors and loss of entire packets. For a certain cost (in data rate), data can be protected to achieve a lower Bit Error Rate (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 to damaged addresses in particular.

APPROACH

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

- the observation that the required total data rate for video communication is the product of the spatial (xy) resolution, the shade/color (z) resolution, and the temporal (t) resolution;
- the assumption that the nature/dynamics of the scene will vary;
- the need to cope with sudden degradation of the channel performance (capacity, delays, BER, etc.): and
- the assumption that both the sender and the receiver have a significant amount of processing power and storage.

Our strategy is both dynamic and adaptive. It varies according to conditions observed in real time, such as the dynamic nature of the scene and 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. We assume 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 need not 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, delta modulation, or predictive coding) requires the transmission of control information to manage the connection, and of data to convey the signal. For packet video communication we propose schemes which, in addition to sending connection control and image data, also send "functions" telling the receiver what to do with the data. Therefore the entire encoding algorithm can be dynamically changed if the sender finds it desirable to do so. Or, the function may simply change parameters such as the granularity of coding tables or the block sizes.

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

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 a sudden decrease in channel performance may only cause the received image to suffer from quality degradation rather than total loss of parts of the image. This happens because of implementation strategies in the network packet switches which automatically drop packets of lower priority when the total demand exceeds the available capacity.

The successive approximations may obviously be along any dimension, or a combination of several dimensions. This includes dimensions from the image domain (xy , z , and t) or dimensions from any transform domain, such as a 1- or 2-dimensional Fourier or cosine transform.

Consider, for example, the communication of an image of size $xyz=512 \times 512 \times 8$ bits. One way to send the image is in the xyz order: first sending all the bits (z) for the first pixel in the first row, then stepping along the row (x) for all the pixels in the row, advancing through complete rows (y) until all the bits of the last pixel in the last row have been sent. This is probably the simplest way to communicate 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 rate of the channel increases, 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, the less significant packets may even be sacrificed entirely in order to afford the added data rate needed for the protection of the more significant packets.

If 512×512 -bit packets are still too big for reasonable communication, they can be divided into smaller packets, either by covering smaller areas, or by having each packet cover the entire 512×512 area but in a lower spatial resolution. The latter is similar to the way smaller memories are organized into bigger ones using various 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 may look like focusing a lens, where the entire image is transformed at once from a low quality image into a high quality image. The xyz -order scheme described earlier 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.

Procedure

The following procedures describe in general the operation of the transmitting and receiving stations implementing a successive approximation scheme.

For the transmitter:

1. Get a new image.
2. Decide which algorithm to use for that image.
3. Encode and send the algorithm code and its parameters with the encoded image
4. If there is no more time left, send a "Display Image" command and go to Step 1.
5. Else, subtract the transmitted image from the original image and go to Step 2.

For the receiver:

1. On receipt of a new function and data: use the specified function/parameters to combine this data with the image currently being constructed.
2. On receipt of "Display Image": perform LPF, if needed, on the image currently being constructed, display that image, and prepare to construct a new image.
3. On timeout: assume receipt of a "Display Image" command.
4. Occasionally send to the transmitter a performance evaluation of the channel, including parameters such as the actual throughput and observed error characteristics.

IMPLEMENTATION

A real-time packet video communication system is being implemented at ISI to test some of the techniques suggested in this paper. Our objective is to build a system capable of transmitting color video with moderate motion in real time at a data rate of 1.5 Mb/s. Monochromatic video will be used initially, with color to be added later.

To reduce the data rate of raw, digitized video to 1 or 2 Mb/s requires compression by a factor of at least 32. We can reasonably achieve a factor of 4 by reducing the spatial (xy) resolution to 240x256, which is approximately the practical resolution available from a typical home television receiver. Another factor of 2 can be squeezed from the temporal (t) resolution by reducing the frame rate from 30 to 15 frames/second and interpolating at the receiver for a 30 Hz update rate. However, further reduction by a factor of at least 4 is still required, calling for a bandwidth compression algorithm most likely based upon a domain transformation.

Many good techniques for bandwidth compression have been described in the literature, and some have been implemented in hardware [1]. Commercial bandwidth compression systems are available, but they cannot readily be adapted to take advantage of the nature of the packet satellite channel and they are very expensive (more than \$150,000 per installation). Therefore we intend to build a real-time compression system tailored for packet communication which utilizes the most appropriate of the available algorithms. We choose not to try to invent a new bandwidth compression technique to outperform those already developed. Instead we believe we can transmit better video in fewer bits by concentrating our efforts on coding and other aspects of the compression system, optimizing the packet satellite video system as a whole.

Our approach is to pick a compression technique based on:

- how well it works in the packet switching environment, allowing us to use the dynamic, adaptive approach described earlier; and
- how well it copes with the kind of errors the packet satellite channel is likely to exhibit.

Selected Algorithm

The bandwidth compression technique we have chosen is two-dimensional block transform coding using the Discrete Cosine Transform (DCT) [5]. For block transform coding, each incoming image frame is first subdivided into a matrix of blocks. The two-dimensional DCT of each block is computed and the resulting transform coefficients are coded in some number of bits according to their importance: coefficients corresponding to high frequencies may be discarded completely. The coded transforms are packetized, transmitted to the destination, decoded, and transformed back into subimage blocks.

Block transform coding is well suited to the types of adaptive techniques we have described. Only those blocks which have changed from one frame to the next need be transmitted. The blocks which have changed can be transmitted in several packets, with the most important coded transform coefficients transmitted in the highest priority packets, etc. Of course, errors or packet losses may result in old data being displayed for some time if updates to a block are infrequent, but this problem can be reduced by cycling through all the blocks and forcing a few to be updated in each frame. Packet communication also makes it easy for the receiver to return to the sender a notice about blocks which arrived with errors if the data is covered by checksums.

Larger block sizes result in higher image quality, but it is very difficult to build real hardware for block sizes larger than 16x16 pixels. Therefore our system, like almost all block transform coding schemes, will use blocks 8 or 16 pixels square or perhaps rectangles of 8x16 or 16x8 pixels. These block sizes also match well with the size of major features in typical images.

In addition to block transform coding using the DCT, we have considered other bandwidth compression algorithms, such as differential schemes including predictive coding and delta modulation, and hybrid methods which typically use a differential scheme in one dimension and transform coding in the other. Methods which do not require transform coding are generally much simpler to implement in hardware, but they cannot provide adequate bandwidth compression without severe loss of quality. Transforms other than the DCT, such as the Hadamard and Fourier transforms, can also be used for block transform coding, but the DCT gives provably optimum [6] compression at a slight increase in computational complexity. Images can typically be coded to one bit per pixel with little apparent loss of quality.

Block transform methods are also superior to differential and hybrid methods in their ability to suffer channel errors with a minimum of image quality degradation. Each error is confined to a single block; its effects, if visible at all, are distributed fairly uniformly over the block. In the case of differential coding methods, errors can cause visible effects throughout one or more scan lines of the image, causing very noticeable streaks.

Simulation on General-Purpose Hardware

To simulate the bandwidth compression techniques, we are using a general-purpose peripheral array processor, the Floating Point Systems AP-120B, which is interfaced to a Digital Equipment Corporation PDP 11/45 minicomputer host. The effects of varying xy and z can be simulated on the AP-120B, as well as the effects of channel errors. But the AP-120B is slower than real time by a factor of 10 or more, so processing of real-time sequences of images is limited to the small number of frames which can be captured in the frame memory of the display system. Significant testing of the dynamic techniques we have described will have to wait for real-time hardware to be built.

Real-Time Implementation

We are currently designing hardware for a real-time implementation of the video bandwidth compression. The complete packet video system includes the following components:

- A commercial video digitizer, frame buffer and display.
- Special-purpose bandwidth compression hardware.
- A high-bandwidth multi-processor to packetize the data.
- The DARPA Wideband Satellite Network consisting of packet switches, error-correcting modems, and ground stations.

The following paragraphs describe these components which are pictured in figure 1.




Figure 1: Block Diagram of Real-Time Implementation

IMPLEMENTATION

IKONAS graphics system

The commercial video system, an IKONAS RDS-3000, provides video input and output functions and a frame buffer which can operate at full video frame rates. The IKONAS system was chosen because it has a bus structure which allows us to add compression and coding hardware in a modular fashion. The system can be configured with one or two bit-slice processors; these will be used for control and encoding/decoding functions.

Bandwidth compression processor

There are several ways to achieve the processing capacity required for real-time video. The two main ways we considered were:

1. one or two very fast, powerful, specially-designed processors; or
2. several small, general-purpose signal processors, executing the same program at the same time on different data streams.

We have chosen the second approach, because it appears to be simpler to implement, less expensive, and more flexible than the first approach. We are designing a board to plug into the bus of the IKONAS system. The processors to be used are Texas Instruments TMS320 signal-processing-oriented microprocessors, each capable of 5 million operations per second. An 8-processor system can do either the DCT or its inverse in real time on a 240x256 image segmented into blocks up to 16x16 pixels in size.

Packetization functions

As raw video data is transformed and encoded, data segments will be produced for transmission across the channel. The segments will contain minimal header (control) information since the processing available in the IKONAS bit-slice processor will be limited. The packetization function will be completed by sending the segments over a 2 Mb/s serial link to a multi-processor packet switch called a Voice Funnel, where complete header information for the Packet Video Protocol (PVP) [2] and Stream (ST) [4] protocols will be added.

The Voice Funnel packet switch [7] is built by Bolt Beranek and Newman, Inc. (BBN). Its hardware consists of multiple MC68000 microprocessors which intercommunicate via a high-bandwidth butterfly switch. The standard Voice Funnel software functions as a gateway between multiple hosts or local networks and the Wideband Satellite Network. It supplies the necessary protocol for communication with the network node (PSAT) and obtains resource reservations as needed.

We plan to implement additional software for the Voice Funnel to "packetize" the video data being transmitted and "depacketize" the data being received. The receiver's task is more complicated than the transmitter's: buffering is required to smooth out variance in network transmission delays, and it may be necessary to sort out-of-order packets or remove duplicates caused by retransmission. The receiver will also streamline the header information before transmitting the video data over the serial link to the IKONAS system for processing.

Wideband Network

The Wideband Satellite Network consists of many subsystems itself. Included are the packet-switch node called a Pluribus Satellite Interface Message Processor (PSAT), also built by BBN [3]; the Earth Station Interface (ESI), which is a high-speed burst modem plus an encoder/decoder for forward error correction, supplied by Linkabit Corporation; the earth station transmitter, receiver, and antenna supplied by Western Union; and, of course, a channel on the Westar III satellite.

The PSATs work together to partition time on the shared satellite channel. Time is allocated for two classes of

traffic: "streams" can be reserved for data with a slowly-varying rate, and "datagram" service is provided on demand for bursty traffic. Before a packet can be transmitted as a datagram, a reservation must first be sent, resulting in an overall delay of two satellite round trip times. To avoid this delay for real-time data such as voice or video, a stream can be preallocated to match the periodic transmission rate of the data packets.

The disadvantage of a stream is that it takes longer to set up or modify the allocation. If the required data rate drops, part of the allocation will be wasted. If the data rate increases above the allocation, the PSAT may be forced to discard some of the data packets. Fortunately, changes in the video data rate occur relatively infrequently, allowing the allocation to be reduced when appropriate for increased efficiency. If the video data rate increases suddenly some data may be lost before the allocation can be increased. But, by using the priority structure provided by the PSAT, it is possible to ensure that the most significant image information gets through so that the quality degrades gracefully.

In addition to packet loss, the video data may be damaged by bit errors. The bit error rate of the wideband satellite channel may be as high as 5 errors in 1000 bits (5×10^{-3} BER) due to noise on the signal. To compensate, the ESI includes a convolutional encoder and a sequential decoder for forward error correction at the four coding rates 1, 7/8, 3/4, and 1/2. These rates represent the ratio between the number of bits in the raw information and in the error-correcting-coded signal which is transmitted over the channel. For example, rate 1 coding corresponds to no correction at all, while rate 1/2 coding transmits twice as many bits to achieve a substantial reduction in the error rate.

Video System Optimization

By treating all the components of the real-time packet video system together, the overall performance can be optimized by adjusting the parameters of each unit to accommodate the changing data-rate requirements of the scene and the available channel bandwidth and error characteristics.

In addition to tuning of the bandwidth compression algorithm itself, there are factors whose adjustment affects several components in the system. A good example is the choice of packet size. To reduce the load on the packet switches and for optimum utilization of the channel, packets should be made as large as possible to keep down the rate at which they are transmitted. But, the loss of a large packet is more damaging than the loss of a small one. Therefore a balance must be struck and perhaps changed as the situation changes.

Selection of the ESI coding rate is another important part of the optimization. For a fixed channel capacity, we must determine which of the following strategies improves overall quality:

1. increasing the ESI coding (e.g. to rate 1/2) to reduce the bit error rate, which requires more video bandwidth compression to fit the reduced data rate available; or
2. decreasing the ESI coding (e.g. to rate 1) to allow a higher data rate and less stringent video bandwidth compression, while suffering a corresponding increase in bit errors.

One might expect the first choice to be the better one since it employs more "intelligence". On the other hand, if the compressed data can be encoded so that it is relatively insensitive to bit errors, the second choice might give better overall quality. The optimum solution might be a combination of increased coding for the most significant data and decreased coding for the less significant data.

CONCLUSIONS

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 important to dynamically monitor the channel characteristics (available data rate, total demand, BER, etc.) in order to optimize channel use by performing the proper tradeoffs.

To verify these assertions, a real-time packet video system is being developed at ISI which will demonstrate packet video communication across the DARPA Wideband Satellite Network.

REFERENCES

1. Chen, Wen-Hsiung, "Scene Adaptive Coder," in *IEEE 1981 International Conference on Communications*, pp. 22.5.1-6, IEEE, June 1981.
2. Cole, E. R., *PIP - A Packet Video Protocol*, USC/Information Sciences Institute, W-Note 28, August 1981.
3. Falk, G., S. Groff, R. Koolish, and W. Milliken, *PSAT Technical Report*, Bolt Beranek and Newman Inc., Report No. 4469, May 1981.
4. Forgie, James W., *ST - A Proposed Internet Stream Protocol*, M.I.T. Lincoln Laboratory, IEN 119, September 1979.
5. Jain, A. K., "Image Data Compression: A Review," *Proceedings of the IEEE* 69, (3), March 1981, 349-389.
6. Pratt, W. K., *Digital Image Processing*, Wiley-Interscience, New York, 1978. [See especially chapter 23.]
7. Rettberg, R., C. Wyman, D. Hunt, M. Hoffman, P. Carvey, B. Hyde, W. Clark, and M. Kralej, *Development of a Voice Funnel System: Design Report*, Bolt Beranek and Newman Inc., Report No. 4098, August 1979.

END

FILMED

1-84

DTIC