JOURNAL OF SPACECRAFT AND ROCKETS
Vol. 33, No. 1, January–February 1996

Systems Design of a Satellite Link Protocol

John Murphy,* Edward Chow,† and Richard Markley‡
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

We show how it might be possible to adapt asynchronous transfer mode (ATM) technology to satellite links. The ATM is a high-speed protocol designed with optical fiber as the intended transmission medium. Several problems arise when satellite channels are used. We propose to move the error recovery and detection from one layer of the ATM protocol to a higher layer. We base the retransmission strategy on the service carried, which utilizes the ability of the ATM protocol to differentiate services. Our simulation results show that we not only increase the raw data throughput for satellite channels to almost the theoretical limit, but we have improved the data transfer efficiency of the ATM protocol by 7.5%. Our results also show that it is possible to guarantee data services with no loss of data under certain conditions.

Introduction

THE asynchronous transfer mode (ATM), an emerging standard for broadband communications, involves the switching of small fixed-length packets.[1,2] It allows the sharing of resources between many different services by statistically multiplexing the sources. Although the ATM has been proposed for high-speed services, at rates of 155 and 622 Mb/s, there is increasing interest in its use at lower rates,[3] such as 1.544 Mb/s (T1). The interest is stimulated by two factors: 1) the lack of resources or of sufficient bandwidth for the higher speeds and 2) lack of services requiring them.

When dealing with mobile communication systems it is unlikely that the larger bandwidths will be available in the near future. However, there are a number of areas, civilian and military, where mobile or wireless operation is critical, and there is a desire to investigate the performance and advantages of the ATM in this area. There is also a current need to share the bandwidth resource in a more efficient manner. Multiplexing the services that are currently using the bandwidth is seen as an advantage of the ATM.

The Jet Propulsion Laboratory (JPL) operates the Deep Space Network (DSN) for NASA. The ground facilities of the DSN consists of three main antenna sites, around the world, and JPL, all connected by commercial satellite links. The bit rates being used are generally T1 or less. The services that use these links include command files for spacecraft and ground distribution of telemetry, as well as voice communications among operators and test data. There is also the future possibility of using these links for digital video. The accuracy of the data from and to spacecraft is critical, whereas other data and voice may not be as important. At present there is no automatic method of differentiating between the services that use the network.

The ATM has been designed with optical fiber as the expected transmission medium. This means that the expected errors will be produced by Gaussian noise and hence will introduce random geometrically distributed bit errors. With bit errors expected to be of the order of 10^{-10}, this means that at most one bit error will be expected in each 424-bit cell. Therefore error detection and one-bit error correction will almost fully protect the data. In satellite links, on the other hand, bursty errors are common. These bursts may be a result of Gaussian noise, but rather than producing randomly distributed bit errors, they produce a large group of bit errors, due possibly to the coding of the channel. It is also possible that degraded performance may occur on the link, causing increased bit errors, which may destroy the link for a fraction of a second or more. The overall performance of the satellite links is usually about 10^{-7}, and single-bit correction is not likely to protect the data. The loss for the data and the header of the ATM cell is shown in Fig. 1 vs the bit error rate. It is seen that for optical fiber there is about one cell lost every century for a 155-Mb/s link, whereas for satellite links the loss will be about one cell every minute. Therefore different error detection methods must be implemented when using the ATM over satellites. If the errors were likely to be all bursty, then the cell header could be used to discard incorrect cells. However, with both bursty and random error the information in the cells also needs to be protected from errors. There are many schemes proposed to adapt the ATM to non-optical-fiber environments[4]; however, these do not take into account the type of error found on the satellite links used by JPL.

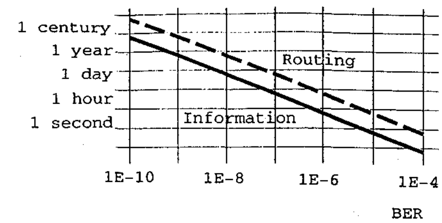

Fig. 1 Loss for 155-Mb/s link for various bit error rates.

The rest of this article is organized as follows: the model for the satellite channel is described; modifications to the ATM protocol are investigated; a model for the system and the simulation are shown, and the results are presented.

Channel Model

The satellite link is characterized by the delay across the link and the bit rate of the link. In this article the delay is taken to be 270 ms, and the bit rate is T1, or 1.544 Mb/s. There is also the error due to incorrect decisions on the bits at the receiver. There are many models proposed to model this bit error, and most are based on reasoning why the errors occur. Rather than take one of the predefined models or be involved in the actual coding and transmission schemes being used on the link, we use an empirical model based on actual test data taken from the DSN, and we also assume Gaussian noise. Therefore we end up with a combination of burst errors and single-bit errors due to Gaussian noise.

The experimental data used to model the burst-error distribution are from fractional T1 links, which vary in bit rate from 56 to 224 kb/s, and the tests were taken over a period on the order of two days. A total of 10^{11} bits were used from the test. On investigation of these results we first made a distinction between nonburst events and burst events. During the nonburst events only single-bit errors occur in a random manner, and the satellite link is in this mode of operation about 99.75% of the time. Secondly, while in a burst event, groups

Received Feb. 14, 1994; revision received May 17, 1994; accepted for publication May 20, 1995.

*Academic Part Time, Advanced Information Systems Section; currently Lecturer, Department of Electronic Engineering, Dublin City University, Dublin, Ireland. Email: murphy@systems.caltech.edu.

†Member of Technical Staff, Advanced Information Systems Section. Email: edward-chow@isd.jpl.nasa.gov.

‡Technical Manager, Advanced Information Systems Section. Email: richard.w.markley@jpl.nasa.gov.

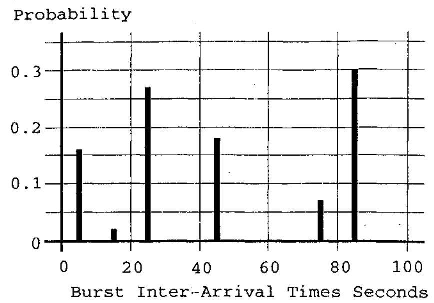

Fig. 2 Probability distribution of burst interarrival times (IAT).

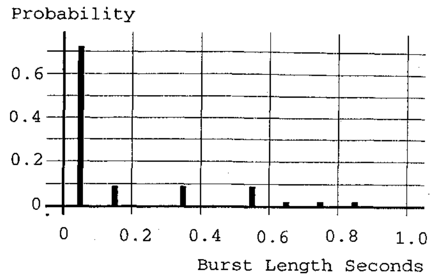

Fig. 3 Probability distribution of burst length.

of continuous bit errors occur. These burst errors do not occur in a random fashion, and so we model the interarrival time of the burst errors by another empirical distribution. Thus the resulting model is thus made up of random single-bit events and burst events. In the burst event the probability distribution for the interarrival time of burst errors is shown in Fig 2. The mean time between bursts is approximately 36 s, and the variance is 24 s. The duration of the burst itself is given by the burst length, and the probability distribution for this is shown in Fig 3. The mean burst duration is 0.15 s with a variance of 0.26 s. Thus a burst may have tens of thousands of bits in error continuously. We model both the burst length and the burst interarrival time in terms of time rather than in terms of bits, so that this can be applied to links of different bit rate.

Proposed ATM Adaptation Layer

The ATM protocol has three layers, corresponding approximately to the bottom two layers of the OSI model.[5] The bottom layer, which is called the physical layer, is concerned with the physical transmission of the bits. The second layer, the ATM layer, takes a payload of 48-byte cells and puts a 5-byte header on to it to form a 53-byte cell. This header has the routing and addressing information in it, as well as an 8-bit cyclic redundancy code (CRC) that detects header errors[2] and corrects single-bit errors. The 48-byte payload comes from the third layer, the ATM adaptation layer (AAL), which adapts the cells to the different services. The lower part of the AAL, called the segmentation-and-reassembling (SAR) sublayer, breaks a message up into cells. There are a number of proposed AALs specified for different applications.[1,2] The proposed AAL for data transfer is AAL 3/4, which uses a 44-byte information load and a 2-byte header and trailer. In the trailer of AAL 3/4 is a 10-bit CRC that protects the information load from errors. However, on account of the small CRC, the potential undetected error rate is high (10^{-3}), and there is no correction for double-bit errors. Furthermore, the four bits allowed for a sequence number will limit the throughput on the long-delay satellite links, e.g., to 10.5 kb/s when the delay is 270 ms.

We suggest that it is more efficient to consider moving the CRC and the sequence numbers to the higher convergence sublayer[6] (CS). By moving the sequence number to this layer we gain larger blocks to put error detection and correction data on. For the AAL 3/4 the overhead is about 8% but to modify that for a T1 satellite link we would need at least an 11-bit sequence number. Also, by moving the CRC to the CS sublayer the undetected error is reduced to less than 10^{-9}. The ATM standards specify many AALs that can be used, but it is permissible to specify a different AAL for a particular application. We suggest that high-speed, long-delay satellite links need a unique AAL. This can be achieved by inserting the sequence number and the CRC at the CS sublayer, requiring a type of framing. The AAL that we propose is capable of supporting 155-Mb/s throughput with a 270-ms delay on a single channel. The proposed framing structure is shown in Fig. 4.

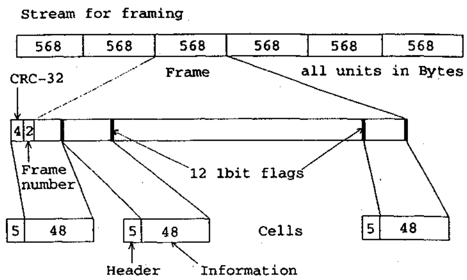

Fig. 4 Proposed convergence sublayer AAL.

In this article it is also suggested that there be service selection to improve efficiency for reliable delivery. This is needed because real-time data and voice usually cannot tolerate delays of one and a half times the round-trip propagation delay. For example, voice will tolerate the 270-ms propagation delay, but if it is in error, it will take another 540 ms for retransmission, even assuming no congestion delay. The total delay is 810 ms, which is unacceptable. Therefore there is no use in trying to detect errors or retransmit voice cells. In fact, any service that is delay-sensitive in the sense of requiring delays to be less than a few seconds is considered here to be real-time and so is not framed for error detection and retransmission. For real-time services no framing is required. Even if the link is fully occupied, we propose to make room for the retransmission of reliable data by discarding some of the real-time cells, e.g., the voice cells.

It can be seen that the frame is at the CS sublayer it takes 568 bytes of user data and adds 8 bytes of overhead to form a 12 by 48 byte frame to be split up by the SAR into 12 cells. The 8 bytes of overhead is made up of a 32-bit CRC, which could be the standard IEEE 32-CRC.[6] There is a sequence number of 16 bits, which allows a full 155 Mb/s of data on a link with the delay as given above. The sequence number needed for 155 Mb/s depends on the size of the frame, and the error characteristics need to be known for an optimally sized frame. It is shown in the Results section that a frame size of about 12 cells is a fair estimate for the conditions of our channels.

At the receiver there is a problem when a cell is lost in that it is not known whether it is a voice cell or is one of the framed cells. To overcome this and to maintain the framing synchronization, a one-bit flag is proposed to show the start of the frame. This flag is the last bit of the cell and tells whether the cell is the first in the frame or not.

The efficiency that we gain is twofold: first, we have only 8 bytes of overhead on a frame of 568 user information bytes, and secondly, we have the sequence numbers to transmit over long delay links. The gain in efficiency compared to AAL 3/4 is that instead of carrying 528 bytes of user information bytes in 12 cells, we now carry 568 bytes, which is a 7.5% increase in efficiency.

Simulation Details

The AAL model and the selective retransmission model were simulated using SES/*workbench*, a discrete-event package that allows hardware and software simulation. The model of the retransmission scheme and the error probabilities in this paper are created by use of the graphical interface. The SES/*workbench* compiles the code to C and runs on a four-processor SparcServer 629. Each simulated data point is the equivalent of 17 days on a T1 satellite link, but rather than modeling the 2×10^{12} bits, only the burst event is simulated. This reduces the simulation to 5×10^9, bits and further optimization of the model compresses this to 5×10^7 events. In real time this can take between 3 and 36 h of run time. In total 1300 h of CPU time was expended in the simulation tests.

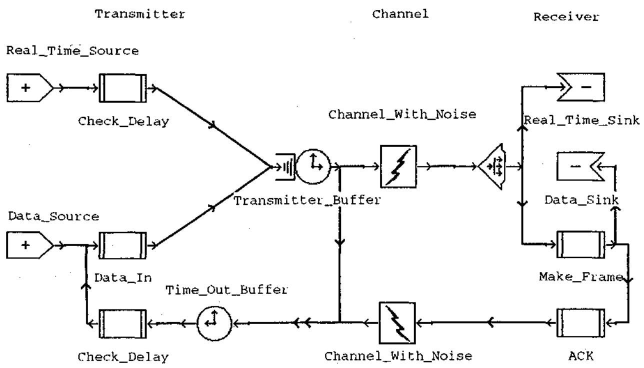

Fig. 5 Simulation model.

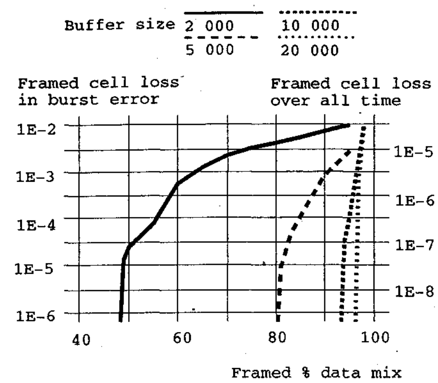

Fig. 6 Framed-data loss for burst error, varying buffer size.

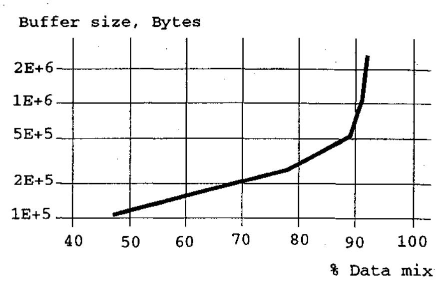

Fig. 7 Transmitter buffer required for no framed-cell loss.

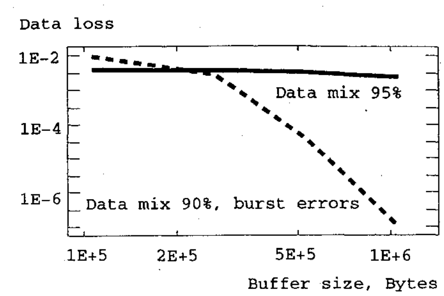

Fig. 8 Effect of varying the buffer size.

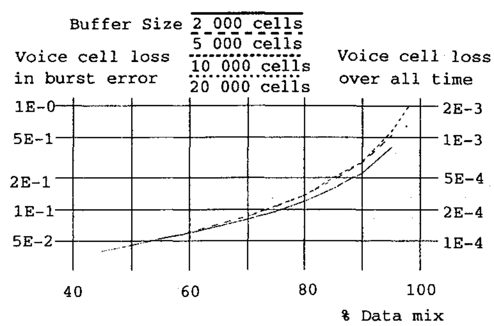

Fig. 9 Real-time cell loss, varying buffer size.

The model consists of the generation of real-time and framed cells as shown in Fig. 5. There is a buffer at the transmitter, which varies in size, that can be used to influence the throughput. The frames are numbered just as they leave the transmitter buffer, modulo 65,563. The channel has two sources of noise: burst errors and Gaussian noise. At the receiver the real-time cells and all the errored cells are released, and the others are framed. If any bits in the frame are in error, the whole frame is discarded. If the frame CRC is correct, then an acknowledgment is sent back to the transmitter, through a channel in which it also may encounter errors. The transmitter, after transmitting the cells onto the channel, keeps the framed cells in a time-out buffer. This buffer is 111.3 Kbyte, which accommodates the two-way propagation delay plus a small processing delay. When an acknowledgment arrives, the frame is discarded from the time-out buffer. Otherwise, after the time-out delay, the frame skips to the head of the transmitter buffer, after checking that there is sufficient room in the buffer. If there is not, then real-time cells are discarded from the buffer, and if this does not release enough space, then the frame is discarded. The real-time source cells check the delay in the transmit buffer before joining the queue. If the delay is more than a specified value, in this case 50 ms, then the channel is assumed to be in burst error and the cells are discarded.

Results

It is possible to guarantee delivery of the framed cells, with the given error probabilities, by means of retransmission. However, if the ratio of framed cells to real-time cells is high, then some framed data will be lost to burst errors. Tests were carried out for three type of errors: first only bursty errors, then only Gaussian errors, and finally both Gaussian and bursty errors. For each of these types of errors the effect of the buffer size is investigated. Also, the effect of changing the proportion of framed-data on the cell loss is investigated.

For burst errors, the cell loss for framed data depends on the proportion of framed data and to a lesser extent on the transmitter buffer, as is shown in Fig. 6. The cell loss rate for framed data can be decreased to whatever value is required by either backing off the proportion of framed data or increasing the transmitter buffer size. However, as the percentage of framed data approaches 100, the effect of the buffer becomes small and the graph becomes almost vertical, meaning that either all the data get through or all get lost. It becomes important to find the value of the proportion of framed data that allows all the framed data to get through without error. For both bursty error and Gaussian noise, the buffer required for zero framed-data loss is shown in Fig. 7. As expected, as the percentage of framed data increases towards 100, the buffer size required increases faster than exponentially. It is therefore apparent that the best scheme is to have a reasonable mixture of framed data and real-time traffic.

Gaussian noise introduces only single-frame errors for bit error rates lower than 10^{-5}. Hence the Gaussian error model on its own will not be sensitive to transmitter buffer size. This can be seen by comparing how the bursty error and the Gaussian error differ in the dependence of framed-data loss on transmitter buffer size, as shown in Fig. 8. The variation in loss on varying the transmitter buffer size is almost negligible in the Gaussian-error case, but is large in the burst-error case. To combat Gaussian errors a buffer the same size as the time-out buffer will suffice.

The gain of getting no loss for the framed data cells is paid for by having to retransmit the cells in error. If there is no excess capacity, as is assumed in this article, then there is going to be loss in the real-time cells. To see the effect of this loss, a plot of real-time cells, or voice cells, lost due to burst events for various transmitter buffer sizes and against the proportion of framed-data is shown in Fig. 9. In the burst event it is possible to lose all the voice cells if the proportion of framed data is high enough. This means that we have run out of bandwidth to retransmit. However, this loss when averaged out over the day is of the order of 10^{-3}. Also, the loss is concentrated into intervals of seconds. There may be schemes to lessen the effect of this loss of voice and spread it out more evenly. However, if the critical objective is not to lose framed data, that is the price that may have to be paid. Also, the size of the transmitter buffer has almost no effect. This is because when the delay expected in the transmitter buffer is more than 50 ms, the real-time cells are discarded. The loss

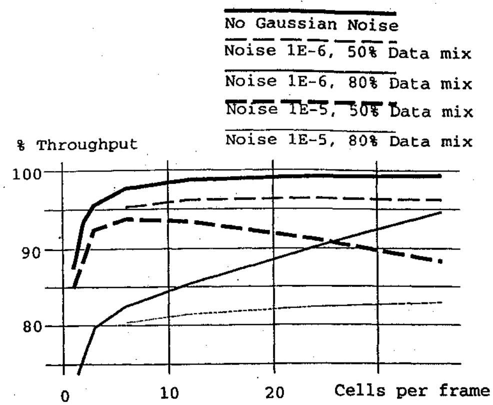

Fig. 10 Variation of the number of cells in a frame.

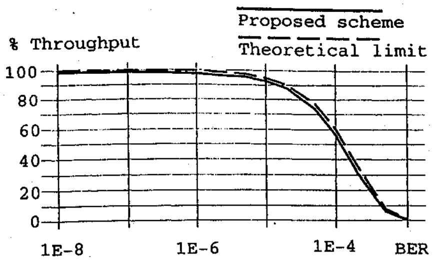

Fig. 11 Comparison of proposed scheme with theoretical limit.

of the voice cells is dependent almost entirely on the error statistics of the model and not on the retransmission model.

The frame size is an important decision, and the optimum value depends on the error rate of the link as well as the distribution of errors. Small frames have high protocol overhead, but have the advantage that when an error occurs only a small number of bits are lost. Large frames have lower protocol overhead, but when one is lost, a large number of bits are lost. It is really the Gaussian error that determines the size of the frame that should be used. The burst errors affect the large and small frames almost the same, because a number of continuous frames will be in error. The throughput when the Gaussian error rate is varied between 10^{-5} and 10^{-6} and the percentage of framed data varied between 50 and 80 is shown in Fig. 10. What is noticed is that for a low proportion of framed cells (for example, 50%) a lower frame size is preferred, but for a higher proportion of (for example, 80%) larger frames would achieve higher throughput. There is a need to compromise the throughput by picking a fixed frame size. If there is only one link, then there could be a optimum frame size, or even an adaptive frame size such as is already implemented in other applications.

The theoretical throughput limit of any retransmission scheme plotted against the proposed scheme is presented in Fig. 11. What is noticed is that we approach the limit almost over the whole range of bit error rates. Therefore we conclude that for these satellite links there is no need to use more complicated retransmission schemes,[7] as there is only a small amount of efficiency left to be captured with an increasing cost of transmitters and receiver protocols.

Conclusions

In this paper we have shown that it is possible to use the ATM over satellite links. We have proposed to place the error control in the convergence sublayer of the ATM rather than the segmentation and reassembly sublayer. This improves both the efficiency of the protocol and the efficiency of the error detection. The relocation of the error control functions is compliant with the ATM standards. We also propose to differentiate the recovery mechanism, or retransmissions, according to the service. This allows a guarantee of no loss, to be given to data services, while not affecting the real-time services. Although the simulation was for a T1 link, this approach can be incorporated into a data stream that is part of a larger link, such as a 45-Mb/s link.

Acknowledgments

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. The authors wish to express their appreciation to Michael Chelian, JPL, who investigated the error distributions on the DSN. They also express sincere thanks to Julia George, JPL, and Ed Upchurch, JPL, who helped with many useful and helpful comments on the SES/*workbench*. The first author would like to thank Peter Colaluca, of SES/*Europe*, for support of the SES/*workbench*, and Charles McCorkell, Dublin City University, for his continuing support and encouragement.

References

[1]de Prycker, M., *Asynchronous Transfer Mode: Solution for Broadband ISDN*, 2nd ed., Ellis Horwood, London, 1993, pp. 102–145.

[2]Lane, J., *Asynchronous Transfer Mode: Bandwidth for the Future*, 1st ed., Telco, 1992.

[3]Chow, E. T., and Markley, R. W., "Asynchronous Transfer Mode Link Performance over Ground Networks," Jet Propulsion Lab., TDA Progress Rept. 42-114, Pasadena, CA, Aug. 1993, pp. 185–189.

[4]Biersack, E. W., "Performance Evaluation of Forward Error Correction in an ATM Environment," *IEEE Journal of Selected Areas in Communication*, Vol. 11, No. 4, 1993, pp. 631–640.

[5]de Prycker, M., Peschi, R., and van Landegem, T., "B-ISDN and the OSI Protocol Reference Model," *IEEE Network*, Vol. 7, No. 2, 1993, pp. 10–18.

[6]Bertsekas, D., and Gallagher, R., *Data Networks*, 2nd ed., Prentice–Hall, Englewood Cliffs, NJ, 1992.

[7]Schwartz, M., *Telecommunication Networks: Protocols, Modelling and Analysis*, 1st ed., Addison–Wesley, Reading, MA, 1987, pp. 119–160.

A. L. Vampola
Associate Editor