

International Journal of

Volume 15 Number 4

July–August 1997

Satellite Communications

Chief Editor: **Barry G. Evans**

Editor, North America: **Peter Nuspl**

ISSN 0737-2884

IJSCEF 15(4) 141–195 (1997)



FEEDERS-TDMA: A DISTRIBUTED-CONTROL ALGORITHM FOR SATELLITE CHANNEL CAPACITY ASSIGNMENT IN A MIXED TRAFFIC AND FADED ENVIRONMENT

NEDO CELANDRONI*, ERINA FERRO AND FRANCESCO POTORTÌ
 CNUCE/CNR Institute, Via S. Maria, 36 56126 Pisa, Italy

SUMMARY

This paper presents Faded Environments Effective Distributed Engineering Redundant Signalling (FEEDERS), an access scheme for sharing, in time division multiple access (TDMA) mode, the capacity of a satellite channel among a number of stations, on the basis of user demand. This scheme and its companion Distributed Allocation with Request In Fixed Slots (DRIFS), result from a study carried out by the authors on distributed-control protocols for geostationary satellite access. Both protocols derive from the Fifo Ordered Demand Assignment/Information Bit Energy Adapter (FODA/IBEA) centralized-control system and have the same features. The distributed technique to compute the capacity allocation adopted by FEEDERS improves some performance in FODA/IBEA, but raises a problem about system stability. Techniques to solve this problem are presented, together with system performance. A comparison is also made with the FODA/IBEA system. © 1997 by John Wiley & Sons, Ltd.

Int. J. Satell. Commun., 15, 185-195 (1997)

No. of Figures: 7 No. of Tables: 2 No. of References: 16

KEY WORDS: satellite access scheme; distributed control; capacity allocation on demand; fading; isochronous traffic; anisochronous traffic

1. INTRODUCTION

The drawback of a system which performs the capacity assignment on demand of a satellite channel in centralized mode is the rather high request-allocation delay. In fact, in a centralized control system, a master station is charged with computing the bandwidth allocation layout, which is broadcast to the traffic stations after their requests have been received. This causes a request-allocation delay of at least two round trips.

In a system designed for multimedia traffic, such as fifo ordered demand assignment/information bit energy adapter (FODA/IBEA),¹⁻³ this has an impact on both the channel set-up and release times of the real-time (stream) traffic, as well as on the end-to-end delay experienced by the packets during the transient, due to a step of the non-real time (bursty or datagram) traffic.⁴ Furthermore, since in the FODA/IBEA system each station adapts the coding and the bit rates of the transmitted data to the link quality, the attenuation level of each link must be estimated and predicted for a time at least equal to the request-allocation delay. The accuracy of the estimation very much depends on this delay, so the

higher this quantity, the higher the margin on the link budget to be adopted, in order to prevent an under-estimation of the link attenuation.

Faded environments effective distributed engineering redundant signalling (FEEDERS) results from a study carried out to reduce the request-allocation delay. To achieve acceptable system stability, the control information must be more reliable than in centralized control. For this reason the control information is made redundant.

As in FODA/IBEA, the fade countermeasure system adopted by FEEDERS is based on the adaptation of the energy per information bit to each individual link status, which depends on atmospheric conditions. The total attenuation of each link (up-plus down-link) is compensated for by varying the transmission power, coding and bit rates. Assuming a multi-channel time division multiple access (TDMA) access to the satellite, the transmission power variation must ensure a suitable constant back-off at the satellite transponder input, to avoid excessive intermodulation noise. The power control can thus be used to compensate for all or part of the up-link attenuation, while the total compensation is completed by varying the coding and the bit rates as well.

In Section 2 the proposed allocation scheme is described, and in Section 3 techniques to recover contingencies due to control information misses are studied. Criteria to size the most important parameters, together with some performance evaluations of the system, implemented on the hardware avail-

*Correspondence to: Nedo Celandroni, CNUCE/CNR Institute, Via S. Maria, 36 56126 Pisa, Italy
 E-mail: N.Celandroni@cnuce.cnr.it

¹In Reference 4, the ATM Forum TM4.0 ('ATM Service Category') classification for the stream and bursty traffic categories, and the corresponding ITU-T I.371 ('ATM Transfer Capability') classification is given.

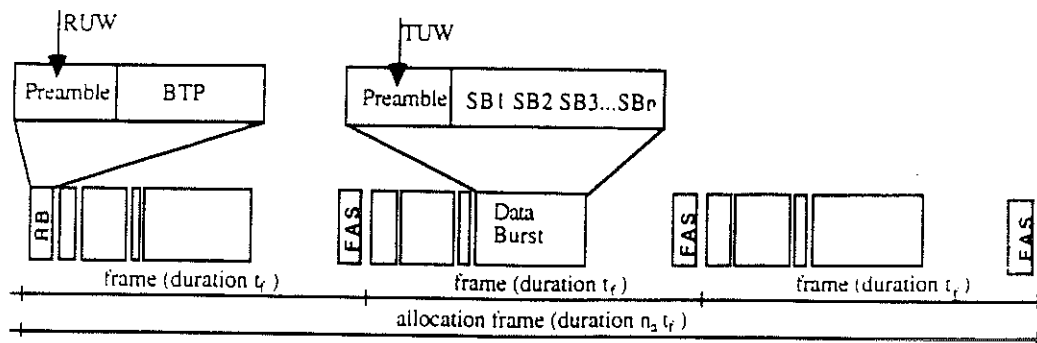


Figure 1. FEEDERS frame layout. There are four active stations and four data bursts are present in each frame. SB denotes a sub-burst and TUW is the traffic unique word which identifies a data burst

able at present, are given in Section 4 and conclusions are drawn in Section 5. Some other access schemes are referenced for comparison in References 5–12.

2. THE ALLOCATION SCHEME

Some definitions are given below together with the access scheme description.

—The frame is the interval of time in which every station transmits a burst. The station data burst may contain both stream and bursty packets. The frame duration time is denoted by t_f .

—The burst time plan (BTP) is the allocation layout where time is allocated for transmitting both bursty and stream data. The BTP must be known by each station, so as to transmit, as well as receive all incoming bursts. This constraint is imposed by the hardware available at present, in which the modem must be quenched immediately before the next burst arrives, in order to avoid false acquisitions of the carrier frequency. The incoming burst bit rate must also be known in advance by the modem. The access scheme is designed considering these limits. Hardware without such limits would simplify some procedures.

—The control message contains the allocation requests, plus some extra information, such as the up- and down-link attenuations. The control message is broadcast by each station in all frames inside the transmission burst header. It is piggy-backed with data, if any.

—The allocation frame (AF) is the time between two consecutive BTP applications. The AF coincides with the time between two consecutive reference bursts (RB) sent by the reference station. The reference burst is sent using a reference unique word (RUW), allowing all traffic stations to synchronize with the network. Inside the reference burst a reference BTP is also sent to allow new stations, which want to enter the system, to set up the BTP for receiving. The reference BTP is also needed in the algorithm to recover missed allocation requests. The AF duration time is denoted by $t_a = n_a t_f$, where n_a is the number of frames contained in an AF (see Figure 1). In each AF, the stations have n_a chances

to transmit, so each individual control message (and consequently each allocation request) is repeated n_a times. The BTP is applied to all the frames in an AF.

The allocation request for stream capacity is computed as in FODA/IBEA and in distributed allocation with request in fixed slots (DRIFS)¹⁴, i.e. it is equal to the bandwidth needed (the sum of all the application needs). The allocation request for bursty traffic is computed by each station as:

$$r = \min[f(q + H i), r_{\max}] \quad (1)$$

where q is the volume of bursty waiting to be sent, i is the bursty traffic coming into the system and f (less than one) and H are the coefficients to be chosen. The request does not have a lower limit, unlike in FODA/IBEA, because there is already one burst per station, so there is no need to allocate a minimum amount for efficiency reasons.

Both types of allocation are computed in a distributed way. They can vary on an allocation frame basis. Each station must be able to compute the BTP for the next AF at the end of the current one. Usually, the BTP is computable after the first frame in each AF. However, if bursts are missed, data needs to be collected in the subsequent frames, in order to build a complete request plan with the allocation requests from all stations.

The stream allocation algorithm is equal to the one adopted in FODA/IBEA. The bandwidth is allocated equal to the request and maintained until it is left.

The bursty allocation is computed as follows. For each station an allocation equal to the request, increased by the preamble overhead,* is prelimin-

*As also stated in Reference 4, even if both distributed-control access schemes studied by us are theoretically independent of the TDMA controller hardware used, we refer to a modem with an acquisition preamble length not negligible as far as the channel overhead is concerned. This is due to the fact that we have the availability of four TDMA controllers and modems that can support different coding rates (1/2, 2/3, 4/5 and uncoded), symbol rates (from 1 to 8 Mbit/s) and output power levels (-20 ± 0 dBm), all variables at the sub-burst level. This hardware has been developed and used for the implementation of the FODA/IBEA centralized satellite access scheme.

arily allocated in each frame. If a residual capacity C_r still remains after a complete allocation cycle, it is shared evenly among all the active stations, even those which sent a null request. If C_r is negative, all requests are reduced by a suitable factor, to make the allocations occupy exactly the capacity C_d reserved for bursty data. In this case, the stations which sent a null request receive an allocation which is only enough to send one control message per frame.

To enter the system, a special first access slot (FAS) is present in the last position of each frame. It is accessed in contention mode. If there might be a collision with another station entering the system, the colliding stations wait for a random number of AFs before repeating the operation. When the maximum number of active stations is reached, the FAS space is temporarily deallocated. In fact, during this period, no more stations are allowed to enter the system. The FAS is allocated again when at least one more entry is possible in the system.

The possible states in which a station can be are:

—Switched off.

—Synchronizing: after a station is switched on, it begins listening in order to acquire the RUW and to read the reference BTP, which is used to receive all the other stations, starting from the second frame in each AF.

—Synchronized: the station correctly receives the allocation requests from all the active stations, so that it can compute the BTP.

—Active: after being synchronized, at the first need for transmission (stream or bursty traffic), the station sends a control message using the FAS in the next AF. This control message is thus repeated n_a times. The first control message may already contain stream and bursty requests. On receiving its own control message in the FAS, the station computes the BTP, considering its own request as well. The station is then able to transmit in the next AF. After receiving a request in the FAS, all stations consider that request as if it were repeated in all the following AFs, until the second request has been received from the neo-active station. This allows the continuity of allocations after the initial request has been received. In fact, the second request is received one round trip later. After a certain time has elapsed without traffic to transmit, an active station quits the system by sending a specific control message. The other stations consider the quitting message and deallocate the space reserved for the quitting station, starting from the next AF. A station must quit the system when it is in outage, i.e. when the measure of its E_b/N_0 ratio at the minimum bit rate, on receiving its own data, is below the acceptable threshold.

—Going-down: when a station misses the allocation requests of at least one station, for an entire AF, it is named missing station. The missed station, i.e. the station that is not heard, is then flagged by the missing station as a going-down station. The missing

station sets the going-down station back to the active state when the next message from the going-down station itself is received.

—Down: a going-down station is declared down when set as going-down in the received reference BTP. The allocation of the down station is definitively removed. A station that is no longer present in the reference BTP must quit the system and re-enter, if necessary, via the FAS.

3. RECOVERY ALGORITHMS FOR MISSED ALLOCATION REQUESTS

The probability of a station going down without warning the system is extremely low. Nevertheless, we must consider that such an event could happen and thus make it recoverable. Missing the allocation requests of one station for a whole AF is another very rare event. In such a case, the missing station can assume either that the missed station died, or that it was not able to receive the missed station. Let us now analyze these two hypotheses and the relevant recovery algorithms. Some items valid for both hypotheses are given below. A distinction between the hypotheses is then made, with the relevant analysis.

A station which receives good allocation requests from all the active stations in an AF is able to compute a *valid* BTP for the next AF. Otherwise, the BTP is only *presumed* and on receiving the relevant AF, an attempt is made to receive the first frame of that AF with the presumed BTP. If the presumed BTP is wrong (this is proved by the presence of missed bursts), the reference BTP just received, is used to receive the following frames of that AF.

Independently of its application for transmitting in the next AF, each station must always compute a BTP (valid or presumed) for the reception of that AF.

When a station is declared down, its allocation is removed when the next valid BTP is computed, i.e. the BTP is computed with the correct allocation requests received from all other stations. That time is considered the end of the recovery algorithm for a down station.

A collision event is detected in an AF when the station down-link is good, at least one station has been missed, and the number of corrupted packets (CRC incorrect) exceeds a certain threshold. Let us define as previous BTP (PBTP) the reference BTP received in the current AF. When a station detects a collision in an AF, it uses the PBTP to transmit in the next AF. If a valid BTP is different from the reference BTP, this means the reference station used a BTP which was presumed, and its choice was wrong. The reference BTP is ignored for the reception of that AF, but, even though it was wrong, it still has to be used in case the contingency recovery algorithm imposes the assumption of the PBTP.

H1 hypothesis: a missing station assumes that the

missed station has gone down. The missing station prepares the BTP for the next AF as if the missed station had sent a null request. This is necessary in order to give the missed station further chances of transmitting the control message. The missed station is flagged as going-down by the missing station.

H2 hypothesis: a missing station assumes that it was not able to receive the missed station correctly; thus, it is not allowed to transmit in the next AF. The missing station sets the missed station as going-down. The presumed BTP, which will be used at the reception time, is computed as if a null request from the missed station had been received. After n_d consecutive AFs in which a station is missed, the missing station is allowed to transmit in the next AF, assuming the PBTP. If all stations are missed in an AF, but the reference burst is correctly received, the missing station is allowed to transmit in the next AF, assuming the PBTP.

Figures 2 and 3 show three different contingencies which are recovered by applying the algorithms relevant to hypotheses H1 and H2, respectively. Contingency (a) represents a station that actually goes down, (b) represents a station which misses the control information of another station for one AF, and (c) shows a station which is missed by all stations, including itself, for a number of AFs insufficient to be declared down. Although we do not give an exhaustive scenario of all possible contingencies, many other situations which need to be recovered can be explained by combining a number of these cases. The downward arrows in these figures indicate that the BTP computed in a received AF is applied for transmitting in the next AF. The upward arrows indicate when a transmitted AF is received. In all the three cases, the first miss is detected in the AF numbered $n - 6$. Let us now look at the different behaviours of the two algorithms for each individual case.

(a) The missed station is declared as going-down until confirmation is received in the reference burst; that station is then considered down. In H1 no transmission is missed by any station. The only anomaly is that the minimum allocation is given to the down station until the AF $n + 5$. In H2 the missing station causes a missing transmission by all stations for two AFs. After the third AF has been received without the missed station, the transmissions are resumed assuming the PBTP for six consecutive AFs. The first PBTP is the one computed by the reference station for transmitting in the AF $n - 4$. The PBTPs assumed for the AFs $n + 6$ and $n + 7$ contain the minimum allocation for the down station, because they are computed after the first miss. The AF $n + 2$ is considered complete, because the missing station is already declared down, so a valid BTP can be computed from it, and the AF $n + 8$ is completely normal.

(b) The missed station is flagged as going-down by the missing station, but this condition is reset in the next AF. In H1, the missing station generally causes a collision on transmitting in the AF n , due to the computing of the BTP which gives the minimum allocation to the missed station. The collision is detected by all stations, who also assume the PBTP for the transmission in the AF $n + 6$. In H2 a missed transmission is caused by the missing station in the AF n . When the AF n is received, no station can compute the BTP, so the AF $n + 6$ is empty. The reception of this empty AF causes all stations to assume the PBTP for transmitting in the AF $n + 12$. This PBTP contains the minimum allocation for the missing station, because of the hole in the AF n .

(c) For five consecutive AFs a station is missed and declared as going-down by all the other stations. This condition is then reset because the missed station reappears before the reception of the reference burst where it is flagged as going-down. In H1, this event causes only the minimum allocation to the missed station, for the AFs in which the station is missed. Thus, after the AF $n + 4$, the situation is normalized. Assuming the H2 hypothesis, in the AFs n and $n + 1$, no station is allowed to transmit, and in the AF $n + 2$, the PBTP is assumed by all the stations as in case (a) above. On transmission of the AF $n + 5$ everything is normal for the complete reception of the AF $n - 1$, but the empty AFs n and $n + 1$ mean that all stations assume the PBTP to transmit in the AFs $n + 6$ and $n + 7$. These two PBTPs contain the minimum allocation for the missed station, because they are computed by the reference station on reception of the AFs $n - 6$ and $n - 5$, in which the station is missed. The reception of the AF $n + 2$ is complete and the situation is normalized after the transmission of the AF $n + 7$.

The H1 algorithm generally produces shorter recovery periods and fewer problems in terms of transmission inhibitions than the H2 algorithm. In particular, contingency (a) does not cause any trouble at all in H1. The drawback of the H1 algorithm affects contingency (b). In fact, H1 here involves a collision (with consequent corruption of data) and a collision detection, which is not a secure event like the missed reception due to missing transmission as considered by H2. Indeed, the collision may not be detected when the power level of the colliding station, at the transponder input, is many dBs lower than that of other stations. This may happen when there is severe up-link fade. However, if the collision is not detected in the first collided AF, a new collision is caused in the next AF, with a greater chance of being detected by all stations, because it

PBTP = previous burst time plan CMA = control message allocation (the min. allocation is given for a CM)
 H = hole (missed transmission) C = collision (on transmission) CD = collision detected (on reception)
 MA = missed by all MDS = missed station GDS = going down station DS = down station

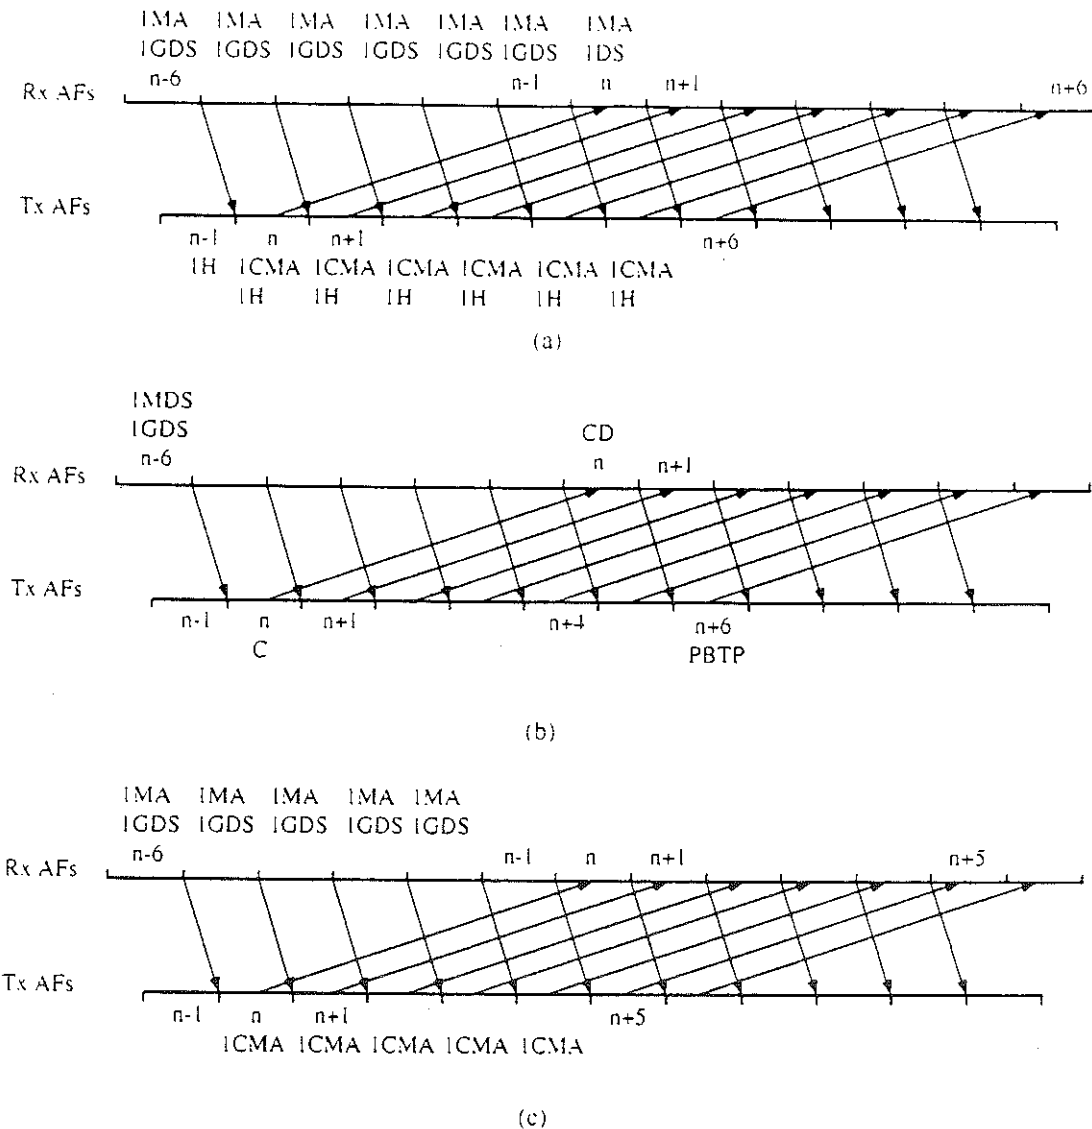


Figure 2. Contingencies recovered using the HI algorithm. (a) a station goes down; (b) a station misses the control message of another station for one AF; (c) a station is not heard by any of the others for a number of AFs. Parameters used: $t_f = 20$ ms, $n_u = 3$, round trip delay = 270 ms

is very likely that this time more stations (which detected the first collision) will collide.

4. SYSTEM SIZING AND PERFORMANCE

The probability that events (a) and (c) occur, depends on station reliability and efficiency of the fade detection and prediction systems. The evaluation of such probabilities is beyond the scope of this paper. Event (b), on the other hand, depends on certain system parameters and the criteria for choosing them are given here. The exact calculation of the probability of event (b) is too complex, so only an increase in this quantity is evaluated.

Hereafter, variables written with capital letters

denote the expressions in decibels of the same variables written in small letters, i.e. $X = 10 \text{Log}_{10} x$.

Let us denote by e the E_b/N_0 ratio available at the receiver of each station when the maximum transmission rate is used. We have

$$e = \frac{e_u e_d}{e_u + e_d + 1} \tag{2}$$

where e_u and e_d are the up- and down-link E_b/N_0 ratios at the maximum bit rate, respectively. Let us denote by r_i the transmission rate with $i = 1, \dots, J$, where r_1 is the maximum and r_J the minimum bit rate. The E_b/N_0 ratio available at the station input

Denoting by b_{ovh} the bit overhead of each burst, we have:

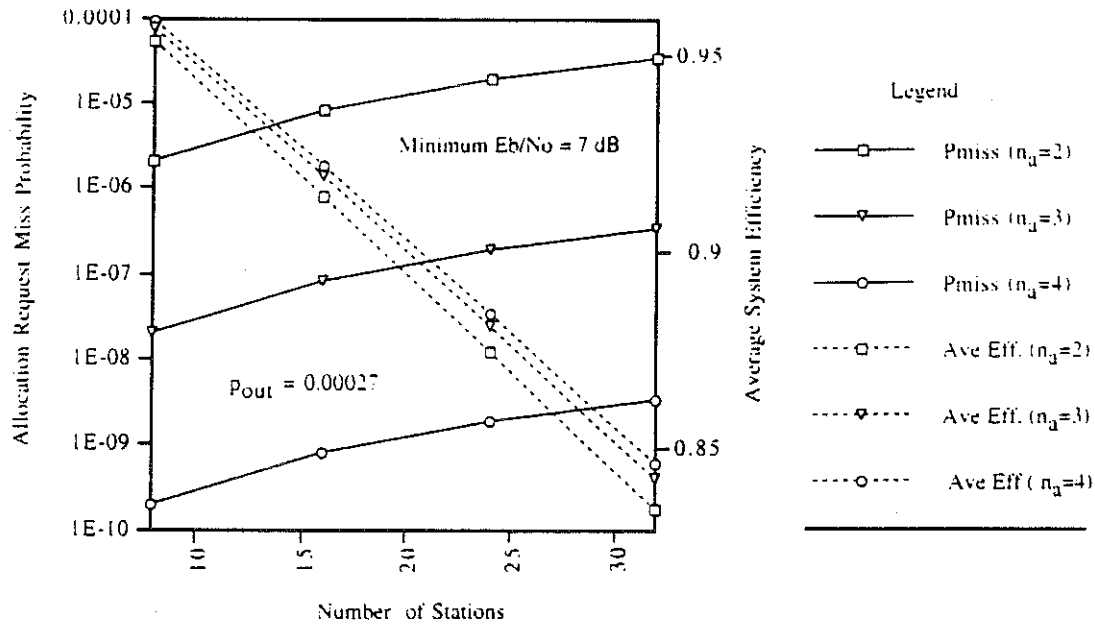
$$b_{ovh} = \text{guard_gap} + \text{preamble} + (\text{burst_header} + \text{crc})/\text{coding_rate}. \quad (10)$$

The burst header contains the control message. In the present hardware the burst header (called control sub-burst) also contains the bit and coding rates of all the sub-bursts contained in the burst. These two transmission parameters are selected according to

the destination station and the class of service (the required BER) of each individual sub-burst. In the following calculations, the size of the burst header is set to the minimum size imposed by the hardware itself.

With the position $p_0 = 1$, the average burst overhead \bar{b}_{ovh} , when none of the stations is in outage, is given by

$$\bar{b}_{ovh} = \frac{b_{ovh}}{1 - p_j} \sum_{i=1}^j g_i(p_{i-1} - p_i). \quad (11)$$



(a)




Figure 5. Allocation request miss probability and average system efficiency as functions of the number of active stations and the number of frames in an allocation frame (n_a). Two different values of the minimum E_b/N_0 are represented: (a) 7 dB; (b) 8 dB. The outage probability (p_{out}) is also indicated

Table I. Parameters used for the numerical example

Reference burst bit rate	2 Mb/s
Reference burst coding rate	4/5
Traffic bursts bit rate	8, 4, 2 Mb/s ($J=3$)
Burst headers cod. rate	4/5
Reference burst length (l_{RB})	$2300 + 160n_s$, max bit rate bits
Burst overhead (b_{ovh})	700 max bit rate bits
Frame length (T_f)	20 ms, 160000 max bit rate bits
Round trip time	270 ms
E_{uc}	15.1 dB
E_{dc}	15.1 dB
E_c	12 dB
A_{001up}	22.5 dB (Ka band)
$A_{001down}$	12 dB (Ka band)
U_f	10 dB

The average system efficiency $\bar{\eta}$ is then expressed by

$$\bar{\eta} = 1 - [l_{RB}(n_s) + \bar{b}_{ovh}n_s n_a] \frac{1}{n_a t_{fr1}} \quad (12)$$

where $l_{RB}(n_s)$ is the reference burst bit length, which depends on the number of stations.

Figure 5 shows a numerical example, with p_{miss} and $\bar{\eta}$ as functions of the number of active stations. Sections (a) and (b) of this figure relate to the choices of 7 and 8 dB for parameter E_{min} , respectively. In the former, a reasonable compromise between the delay gain and the control message miss probability is reached with $n_a = 3$, while $n_a = 2$ gives acceptable results only if E_{min} is limited to 8 dB. This example is tailored for the current hardware operating in Ka band. The parameters used are reported in Table I, while Table II shows the measured burst miss probabilities, with a burst-to-burst frequency offset of 4 kHz. The function used for $p_m(E)$, obtained by fitting the data of Table II with a second order polynomial, is

$$p_m(E) = 10^{1.22 - 0.227E - 0.0333E^2} \quad (13)$$

There are four bit rates possible with the hardware considered. Indeed, the 1 Mb/s rate is used when the link fade is so severe that the 2 Mb/s rate would not be able to guarantee the packet loss rate needed

Table II. Burst miss probability, measured with 4 kHz frequency offset

E_s/N_0 [dB]	Log p_m
7	-2.0
9	-3.5
11	-5.3
12	-6.3

by a high class of service data. Figure 6 shows the probability distribution of system efficiency, in relation to common preamble rates of 8, 4 and 2 Mb/s, when 32 active stations are considered, and $n_a = 3$. We can see that for this number of stations, system efficiency may be rather low, but these events occur for a very small portion of the time.

Both stream and bursty applications benefit from the reduced request-allocation delay of the FEEDERS system in comparison with a centralized one like FODA/IBEA. However, the most evident gain is in the reduction of the delay of bursty traffic during a transient step of traffic coming into a station. A comparison between the two systems is made by using the analytical results from Reference 4, where the system operates in linear conditions. A step of traffic of a sufficiently small size to guarantee linearity is applied to one station alone, while all other stations work in stationary conditions. In order to apply the expressions derived




Figure 6. System efficiency probability distribution for $n_s = 32$ and $n_a = 3$

in Reference 4, it is necessary to determine the value of the request-allocation delay d_{RA} , which, for the distributed case, is given by

$$d_{RA} = \frac{t_a}{2} + t_a + k t_a \quad (14)$$

where the first term is the average time that data must wait for before being considered in the next request, the second term is the duration of the AF, and the third is the propagation delay of the request, where k is the minimum integer, such that $k t_a \geq$ round trip delay. On the other hand, for the centralized case we have

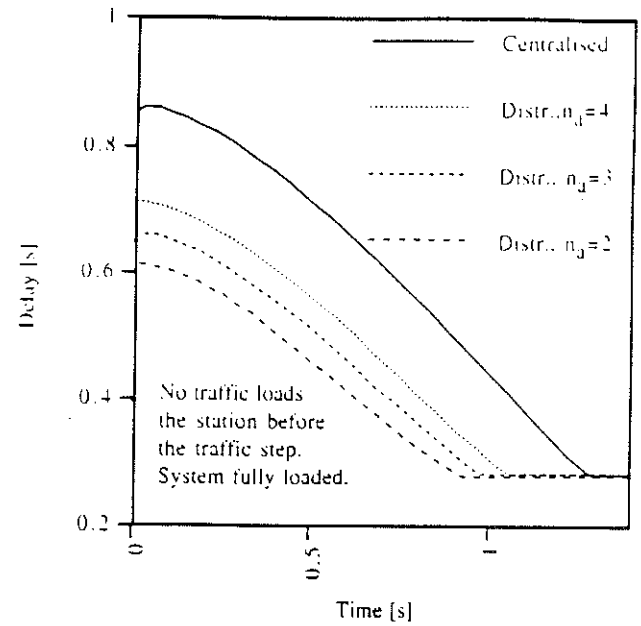
$$d_{RA} = \frac{t_f}{2} + t_f + 2 k t_f \quad (15)$$

where k is the minimum integer, such that $k t_f \geq$ round trip delay.


Figure 7 reports the delay versus time curves for both the centralized and distributed cases, with different values of the parameter n_a . Figure 7(a) represents a case in which the system is fully loaded and the station has no traffic at all before the traffic step (assumed to occur at the time $t = 0$). Figure 7(b) represents a case in which the system is loaded at 83 per cent of its capacity and the station considered is already loaded with a constant traffic rate of a size equal to the traffic step, before the traffic step itself begins. The delay gain of the distributed case, compared to the centralized case, is evident. This result is confirmed by comparison of the performance of the two systems made with simulation in Reference 15.

5. CONCLUSIONS

The FEEDERS system has the potential advantage over the centralized FODA/IBEA system, of improving the system response to traffic transients and of reducing the connection set-up time. Moreover, the dynamic response accuracy to channel fading may be significantly improved. The only drawback is system efficiency, for a large number of active stations and when the Ka band is used. This is principally due to the long preamble that a traditional burst modem needs, in order to keep the burst miss probability below an acceptable threshold, at low E_b/N_0 values. The use of a preamble-less modem, which performs better from this point of view, would remove this inconvenience, thus allowing a higher number of active stations in the system. In addition, if the exact knowledge of the BTP to receive all the bursts were not required, the procedures to recover the contingencies discussed above would be simplified and system robustness increased. These features would further improve the competitiveness of FEEDERS, which already seems promising even when current hardware is used.



(a)



(b)

Figure 7. Delay versus time curves during a traffic step, for both centralized and distributed systems, with different values of n_a . (a) the system is fully loaded and the station is not preloaded; (b) the system is loaded at 83 per cent and the station is preloaded with a constant rate traffic

REFERENCES

1. N. Celandroni and E. Ferro, 'The FODA-TDMA satellite access scheme: presentation, study of the system and results', *IEEE Trans. Commun.*, **39** (12), 1823-1831, (1991).
2. N. Celandroni, E. Ferro, N. James and F. Potortì, 'FODA/IBEA: a flexible fade countermeasure system in user oriented networks', *Int. J. Sat. Commun.*, **10** (6), 309-323, (1992).
3. N. Celandroni, E. Ferro and F. Potortì, 'Experimental results of a demand-assignment thin route TDMA system', *Int. J. Sat. Commun.*, **14** (2), 113-126, (1996).
4. N. Celandroni, 'Analysis of transient of the datagram traffic for a demand assignment TDMA satellite access scheme', to appear in *Performance Evaluation*.

5. D. W. Petr, K. M. S. Murthy, V. S. Frost and L. A. Neir, 'Modelling and simulation of the resource allocation process in a bandwidth-on-demand satellite communication network', *IEEE J. Select Areas Commun.*, **10** (2), 465-477, (1992).
6. J.-F. Chang and S.-H. Lin, 'Delay performance of VSAT-based satellite wide area networks', *Int. J. Sat. Commun.*, **11** (1), 1-12, (1993).
7. S. Tasaka and Y. Ishibashi, 'A reservation protocol for satellite packet communication.—A performance analysis and stability considerations', *IEEE Trans. Commun.*, **Com-32** (8), 920-927, (1984).
8. F. Delli Priscoli, M. Listanti, A. Roveri and A. Vernucci, 'A distributed access protocol for an ATM user oriented satellite system', ICC '89, Boston (USA), June 11-14, Vol. 2, pp. 690-694.
9. A. Baiocchi, M. Carosi, M. Listanti and A. Roveri, 'Modelling of a distributed access protocol for an ATM satellite system: An algorithmic approach', *IEEE J. Select. Areas Commun.*, **9** (1), 65-75, (1991).
10. S. S. Lam, 'Satellite packet communication. Multiple access protocol and performance', *IEEE Trans. Commun.*, **COM-28** (4), 468-488, (1980).
11. P. Papantoni-Kazakos and T. Hall, 'Multiple access algorithms for a system with mixed traffic: high and low priority', ICC '89, Boston (USA), June 11-14, Vol. 2, pp. 665-671.
12. N.-T. Hsu and L.-N. Lee, 'Channel scheduling synchronization for the PODA protocol', ICC '78, pp. 42.3.1-42.3.5.
13. CCIR Report 564-4, *Propagation data and prediction methods required for earth-space telecommunication systems*, 1990.
14. N. Celandroni, E. Ferro and F. Potorti, 'DRIFS-TDMA. A proposal for a satellite access distributed-control algorithm for multimedia traffic in a faded environment', to appear in *Int. J. Sat. Commun.*
15. N. Celandroni, E. Ferro and F. Potorti, 'Comparison between distributed and centralised demand assignment TDMA satellite access schemes', *Int. J. Sat. Commun.*, **14** (2), 95-112, (1996).

