DRIFS-TDMA
A PROPOSAL FOR A SATELLITE ACCESS DISTRIBUTED-CONTROL ALGORITHM FOR MULTIMEDIA TRAFFIC IN A FADED ENVIRONMENT

N. Celandroni, Erina Ferro, Francesco Potortì

CNUCE, Institute of National Research Council
Via S. Maria 36 - 56126 Pisa - Italy
Phone: +39-50-593312 / 593203
Fax: +39-50-589354 - Telex: 500371
E-mail: {n.celandroni, e.ferro, f.potorti}@cnuce.cnr.it

SUMMARY

Most demand assignment TDMA(1) satellite access protocols use centralised-control access schemes rather than distributed ones because their simplicity and robustness usually compensate for the longer allocation delay. Starting from the FODA/IBEA(2) centralised-control protocol, we studied two distributed-control protocols, named DRIFS(3) and FEEDERS(4) respectively, for accessing a geostationary satellite channel. Multimedia traffic and faded environments were considered in the study of both the access schemes. This paper presents the DRIFS proposal, together with the recovery procedures from critical events, whose handling is central to the discussion of a distributed satellite access protocol. Probabilities of such events are also estimated. Reference [12] presents the FEEDER proposal, while in [13] the reader can find the results of a comparison between the two schemes obtained by means of simulation.

KEYWORDS Satellite access scheme Distributed control Capacity allocation on demand Fading Multimedia Traffic

1. Introduction

Allocation on demand protocols are based on the assumption that each station in the network receives a share of the global network capacity that depends on its needs. For TDMA satellite

---

(1) Time Division Multiple Access

(2) Fifo Ordered Demand Assignment/Information Bit Energy Adapter

(3) Distributed allocation with Request In Fixed Slots

(4) Faded Environments Effective Distributed Engineering Redundant Signalling

Preprint of an article published on the INTERNATIONAL JOURNAL OF SATELLITE COMMUNICATIONS, VOL. 15, No. 5

access, the share of each station corresponds to a time slice in the time frame. Some sort of control is needed to take into account the requests and to assign each station its transmission time. If a master station listens to all the requests, and then computes and broadcasts the resulting allocations, the system has *centralised control*. If all the stations independently listen to every other station request, and then compute the transmission time positions and lengths using a common algorithm, the system has *distributed control*.

Centralised control is more robust, because every station only needs to listen to the master in order to know when to transmit, while distributed control entails each station listening to all the requests. In a faded environment, where some stations may receive a low power level because of bad atmospheric conditions, a station may not be able to correctly receive the faded stations' requests. In this case, the burst time plan (BTP), i.e. the layout of the transmission allocations, cannot be computed, and the station cannot transmit without risking a collision, with a consequent loss of throughput. Some methods of recovery from this and other possible scenarios must be devised to make the system reliable. This explains why distributed control protocols are more complex than centralised ones, when a comparable robustness level has to be achieved. Even if robustness needs are satisfied, however, distributed control overheads are usually higher than those needed by centralised control. On the other hand, a distributed scheme is advantageous from a delay point of view, because each station can compute the BTP only one round trip time (RTT) after it has made a request for bandwidth, i.e. about a quarter of second for geostationary satellites.

In this paper we describe the DRIFS distributed scheduling algorithm together with the relevant recovery procedures. The access scheme presented allows the simultaneous transmission of both connection oriented real-time data (*stream data*) and connectionless data (*bursty data*), i.e. it supports multimedia traffic. Since different names are adopted to define concepts that are very similar in purpose, in the text we use the ATM Forum TM4.0 ("ATM Service Category") classification for the traffic categories, while in the footnotes the corresponding ITU-T I.371 ("ATM Transfer Capability") classification is given. The traffic deriving from constant bit rate (CBR)(5) and variable bit rate (VBR)(6) applications, such as videoconferencing, interactive audio, audio/video distribution, audio/video retrieval, native ATM voice, etc., belongs to the *stream* type of traffic, together with data generated by applications that have the ability to reduce or increase their information rate if the network requires them to do so (ABR -available bit rate- applications)(7).

---

(5) Deterministic Bit Rate (DBR)

(6) Statistical Bit Rate (SBR)

(7) Same name in ITU-T

Applications for text/data/image transfer, messaging, distribution, retrieval and remote terminals typically generate traffic belonging to the unspecified bit rate (UBR)(8) service category, generally called *bursty data*.

Rain is a dominant factor in satellite communications at frequencies above 10 GHz. Transmission signal fading due to bad atmospheric conditions is countered in DRIFS by making the data redundant through the use of a variable coding and a variable symbol rate modem, as in the FODA/IBEA system [3, 4, 5]. The redundancy of the data a station transmits in its data bursts depends on the fade levels of both the sending station up-link and the receiving station down-link. The DRIFS system stability is studied under faded signal conditions, and the relevant recovery procedures are reported.

2. Some assumptions

Before describing the DRIFS access scheme, let us make some assumptions and give some definitions.

- The frame is the fundamental time unit used for the allocation scheme. Each station transmits at most one allocation request and one data burst per frame. All data transmitted (control or user data) are preceded by a unique word, specific for the type of data. The frame duration time is T_f.
- The allocation requests are broadcast to all the stations in the system and must be correctly received by all the stations to create a correct BTP.
- Each station repeats its stream and bursty requests at each occurrence of its control slot.

Even if the access scheme is theoretically independent of the TDMA controller hardware used, the possible availability of a preambleless modem would have a heavy impact on some design's choices. Below we refer to a modem whose acquisition preamble length is not negligible as far as the channel overhead is concerned. We have four TDMA controllers and modems available that can support different coding rates (1/2, 2/3, 4/5 and uncoded), symbol rates (512, 1024, 2048 and 4096 Kbaud using either BPSK or QPSK modulation schemes), and output power levels (-20÷0 dBm), all variables at the sub-burst level [1], [2]. This hardware was developed and used for implementing the FODA/IBEA centralised satellite access scheme. The protocol proposed in this paper can be used on the hardware presently available, but it would also benefit from the channel usage efficiency of a preambleless modem.

3. The DRIFS-TDMA access scheme

Each frame consists of a *control sub-frame* (CSF), used in F-TDMA(9) mode, followed by a *data sub-frame* (DSF), used in DA-TDMA(10) mode, and finally by a *first access slot* (FAS), used in contention mode.

---

(8) No equivalent in ITU-T

Each station has a fixed size *control slot* (CS) in the CSF, and it uses its own CS to send control information, while the application data are sent in the DSF. The size of the CSF is then proportional to the number of active stations. The position of a CS inside the CSF is maintained fixed, but it is shifted back when a preceding station leaves the system. Control data sent in a CS are preceded by a *control unique word* (CUW). The first CS in a CSF is preceded by a reference unique word (RUW) which allows the system's frame synchronization.

Each active station that has an allocation in the DSF sends application data in its data burst (DB). Each DB is identified by a *data unique word* (DUW) which precedes the user data. The DB can be divided into data sub-bursts (DSB), each addressed to a specific destination, with individual data bit and coding rates, according to the fade levels of both the sending and the destination stations.

The FAS is used by those stations willing to enter the system.

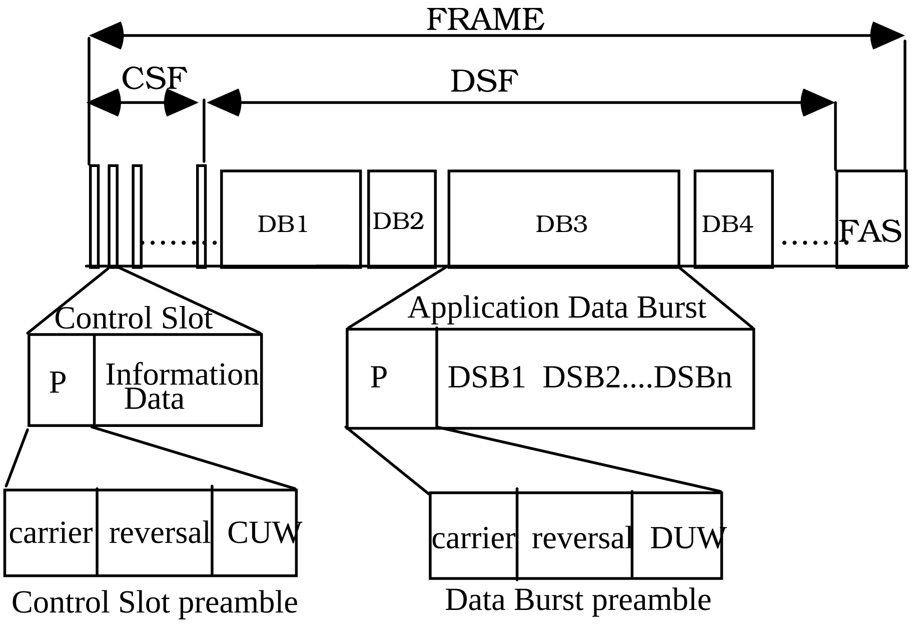

Fig. 1. The DRIFS frame structure

3.1 Station states

A station can be in one of the following states:

• *Switched-off.*

• *Listening.* After a station is switched on it begins listening to acquire the CUWs, to read the allocation requests contained in each control slot, and to compute the BTP. No control slot is

---

(9) Fixed-TDMA

(10) Demand Assignment-TDMA

assigned to the listening station, so it cannot transmit. The station is only collecting information on the other stations and making dummy channel schedulings to be verified one round trip later. After a listening station has computed a certain number of correct BTPs, it is eligible to become *active* and it switches into synchronised state.

- *Synchronised*. To become active a synchronised station must declare its presence by sending a burst in the FAS, located at the end of each frame. Data transmitted in the FAS contain the first stream and bursty requests of the station; the format of the FAS data is therefore exactly equal to the CS data. The FAS is accessed in contention with all the synchronised stations that want to become active. When a synchronised station has successfully used the FAS to enter the network it becomes active and all the active stations update the channel scheduling in order to include one more CS (in the last position of the CSF) for the new active station.
- *Active*. An active station computes the BTP at each frame, it uses its CS to send stream and bursty allocation requests, and it has bandwidth allocated for its data transmission in the frame following the reception of its allocation request.
- *Waiting-Active*. In this state an active station has no bandwidth allocated for data transmissions but the CS must still be used to send its requests in order not to be declared dead.
- *Going Down.* A station which explicitly declares that it will stop any activity is considered *going down.* This does not affect the other active stations. After the completion of the going down procedure, the control slot associated with that station is deallocated. The control slots of all the following stations are then moved back in order to reduce in size the CSF, thus enlarging the DSF.
- *Disappeared.* A station which does not transmit in its CS for a certain number M of frames is considered to have *disappeared*. This causes some problems for the other stations which start the ending procedure (see below). After the completion of this procedure, the CS associated with the disappeared station is deallocated, and the CSs of the following stations are moved back.

Appendix A contains the diagram of the state transitions of a station.

3.2 Use of the control slot

Each control slot has a fixed position in the CSF, as long as no station leaves the system. Every station knows which station is using which control slot, how many active stations there are and the control slot position of each one. When a station stops, all the stations must recompact the CSF, i.e. all the control slots following the stopped station must go back one position. An entering station is always assigned a control slot at the end of the CSF. In the example of Figure 2, N is the number of active stations, whose addresses are $i, j, .. k$, respectively; the INF_i field contains both the stream and the bursty requests of station i , its fade level, and a space for one *new born* or one *dead* station (one at a time) that station *i* has detected.

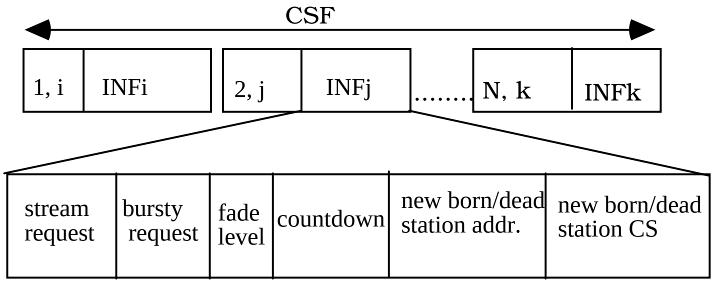

Fig. 2. The CSF format

Since the CS length is not negligible with respect to the frame length (at least using traditional modems that need a preamble), no more than *U* control slots are accommodated in a frame. If more than *U* active stations are present in the system, then the control slots are spread in $C_c =]N/U[$ frames, where *N* is the number of stations and $]x[$ is the smallest integer not smaller than *x*. If *U* is equal to 8, up to 256 stations can be accommodated in 32 frames. C_c is the *slot assignment cycle*, and $F_c = 1/C_c$ is the *assignment frequency*.

The control slots must somehow be protected against transmission errors. Assuming a modem whose maximum bit rate is 8 Mbit/s, we devise a bit rate of 2 Mbit/s and a coding rate of 4/5 for the control slots.

3.3 Leaving and entering the satellite network

3.3.1 Leaving the network

A station may stop its transmissions either spontaneously or unwillingly, due to an abrupt crash or a sudden deep fade level which have occurred in the station. A station is usually able to leave the system in an orderly fashion when the fade becomes greater than the *outage level* O_l that is, the fade level above which transmission and reception become unreliable. In case of crash, the station disappears without sending any end-message, and this causes a certain number of CUW miss events in all the other active stations. On detecting the first CUW miss event, an active station cannot compute the BTP and therefore switches to the *waiting active* state. After *M* consecutive CUW misses, each active station declares the disappearing station dead, by broadcasting in its control slot a countdown from *F* to 1 (*fault countdown*). The number *M* is chosen so that the probability of missing the CUWs of another live station for *M* consecutive times is below a given threshold P_M; the countdown length *F* is such that the probability that another station does not understand any of the *F* countdown CSs falls below another given threshold P_F .

Both thresholds can easily be forced to be small at will (say 10^{-8}), thus making most inconsistency problems negligible. Since both the numbers *M* and *F* depend on the signal quality of the station itself, different fault countdowns may typically be going on at the same time. For example, one station may countdown starting at 3, another one at 4 or 5, and not all of them starting in the same frame. When more than one fault countdown is going on, the lowest is the reference one. The time needed by the system to remove the CS of the disappeared station is $RTT + (M + F)T_f$, where M and F are those of the station with the lowest fade. After that time the CS of the dead station is deallocated, the CSF recompacted, and each station can begin transmitting data again.

When a station spontaneously goes down, it sends an F countdown of end-messages in its control slot (*leaving countdown*). When the countdown reaches zero, the station stops all transmissions. Its control slot is then deallocated and the CSF recompacted. While the normal closing procedure does not affect the remaining active stations, the disappearance of a station causes a temporary interruption in the data transmissions of all the stations.

3.3.2 Entering the network

It is assumed that an entering station A is not in fade conditions. Station A is in synchronised state, and it uses the FAS to broadcast its address, the number of the CS which it will use, plus its stream and bursty requests. Data in the FAS are transmitted with the same protection as the CS data. The entry operation must be repeated for a number E of consecutive frames, each time indicating in the countdown of the INF field how many more times the message will be repeated (*entering countdown*). Since the FAS is accessed in contention mode, let us consider the case that station A collides with another station B (we neglect multiple collisions), and compute such a probability. Let us assume that there is a large number of stations whose requests are generated by a Poisson process with an average equal to Λ per time unit, T_f is the FAS repetition period, and E is the request repetition number. The probability of zero collisions is $e^{-\Lambda(2E-1)T_f}$. Let us call p_{nu} the probability of not understanding an FAS because of a non-acquisition of the carrier or symbol clock, or unique word (UW) mismatch, or CRC error. Assuming $p_{nu} << 1$, and $(2E-1)\Lambda T_f << 1$, the total probability of a station not understanding the entering countdown of another station (p_{ne}) can be approximated by

$$p_{ne} = p_{nu}\, e \;\text{[illegible]}\; T_f\, e \;\text{[illegible]}\; \sum_{i=1} p_{nu} \;\text{[illegible]} = \text{[illegible]}\, T_f\, (1+2p_{nu}) \qquad (1)$$

Collision does not necessarily mean failure, because not all the E FAS transmissions need to be received by the active stations. In fact, one good reception by all the stations is enough. However, the higher the number of colliding transmissions, the higher the probability that some stations does not understand the request, leading to a misunderstanding on the number of active stations and, consequently, to the start of a recovery procedure.

The two terms of (1) represent the probability of a total collision (collisions on E frames) and the probability of collision in all the E frames but one. In case of total collision, all the already active stations do not understand the entering countdown of both stations A and B, so, after a random number of frames, evenly distributed in a range from 1 to *R*, both stations A and B will try again the entering procedure.

The second case is more dangerous, because some active stations may understand the entering countdown of A (which started first than B) and some other may not. This event depends on p_{nu}. This situation may cause misunderstandings on the number of active stations with consequent collisions in the CSs. In order to avoid this, a robust procedure must be adopted to be sure that all the active stations in the network receive the information that station A is entering the network. The procedure adopted is similar to that one used when a station leaves the network (see 3.2.1). As soon as an active station receives a good FAS from A, it starts a *confirmation countdown* in its control slot, repeating for *Q* times (from Q to 1) the address of the station A and the control slot position it will occupy. The value of *Q* depends on the individual signal quality of each station. This procedure must be executed by all the active stations, while station A is still in *synchronised* state, listening for the confirmation countdowns of the other stations. When more than one confirmation countdown is going on, the lowest is the reference one. In absence of collisions, after a time equal to $RTT + T_f(E + Q + C_c)$ the new control slot is added and station A switches in *active* state.

At this point, if at least one FAS transmitted by station B has been correctly received, the active stations start the confirmation countdown for station B, assigning to it the control slot next to station A.

3.4 The allocation algorithm for stream data

The algorithm keeps an account of the stream requests generated by CBR, ABR, and VBR applications. The requests consist of two values (X,Y), whose meaning is as follows:

CBR request:	$X = Y$. Here X and Y represents the requested bandwidth.
ABR request:	$X > Y$. Here X represents the maximum requested bandwidth and Y is the minimum usable bandwidth.
VBR request:	$X < Y$. X represents the bandwidth allocated to the station in each frame. Upon station request, the system allocates to the station the additional Y-X bandwidth previously booked. When the station does not request it, the system can use the extra Y-X bandwidth for the bursty applications of all the stations. The station can use its stream unused bandwidth for the local bursty applications.

When an application makes a request to the station it is connected to, the station makes the same request to the system. Once a station request has been accepted by the system, the stream bandwidth is guaranteed and a time slot is assigned in each frame. If not accepted, the request is simply ignored by the system. On the other hand, as a station is assumed to know the channel occupancy before broadcasting its request, it is reasonable to think that unacceptable requests are not issued at all. The only case when the system actually refuses a request is when several requests are simultaneously issued by different stations, whose sum is greater than the space available. Any modification to a stream request is flagged as such by the requesting station.

3.5 The allocation algorithm for bursty data

The allocation request for the bursty capacity is computed by each station exactly as in FODA/IBEA, i.e. proportional to the traffic coming into the station (i) plus the backlog (q), i.e. the volume of data already waiting for transmission to the satellite:

$$r = q + H\, i \tag{2}$$

where H is a temporal constant. Simulations for the centralised FODA/IBEA algorithm indicated the best value of H to be 0.4, but in the presented access scheme this value may need to be retuned. See [4] and [5] for details on the bursty allocation algorithm. The allocation algorithm requires the assignment cycle, i.e. the time between two consecutive allocations to the same station, to be equal to C_c frames.

4. Bad control slot reception

A station must be able to distinguish a CS reception error caused by a noisy channel from one caused by an empty channel in order to make the right scheduling adjustments. In the first case, a *CRC error* is detected when the data contained in a CS is received; in the second case a *CUW miss* is signalled. In both cases, the station is not able to compute the BTP and, in order to avoid possible collisions, the station does not transmit data in the DSF for the next assignment cycle, but it continues to maintain its control slot, in order to avoid causing synchronisation problems for the other stations.

A sequence of M consecutive CUW misses from station B authorises a station A to begin the fault countdown that declares the station B dead (as described in 3.2.1). The miss probability is given by

$$p_{miss} = p_{na} + ((1 - p_{na})\, p_{wUW}) \tag{3}$$

where p_{na} is the modem non-acquisition probability, and p_{wUW} is the wrong UW probability due to a number of errors bigger than the acquisition threshold T, whose expression is

$$p_{wUW} = \sum_{i=T+1}^{L} \binom{L}{i} BER^i (1 - BER)^{L-i} \tag{4}$$

where L is the UW bit length. Table 1 reports the UW p_{miss} probability obtained from sperimental data by using our available hardware [1, 2].

E_b / N_0 [dB]	p_{miss}
7	10^{-2}
9	$3*10^{-4}$
11	$5*10^{-6}$
12	$5*10^{-7}$

Table 1. Unique word miss probability, measured with 4 kHz frequency offset.

In Figure 3 the typical value of the BER [3] is reported as a function of the channel E_b / N_0, and for different coding rates. The probability of a CRC error on receiving the data in a control slot is given by

$$p_{crc} = 1-(1-BER)^{L_c} \qquad (5)$$

where L_c is the number of bits in a control slot. This probability is much lower than p_{miss} in the presented case, therefore the event of a sequence of consecutive CRC errors is extremely unlikely. If the case of a mixed sequence of CRC errors and UW misses, the same procedure as for the UW miss sequence is undertaken.

On the other hand, if station A has bad CS receptions from all the other active stations (therefore even by itself) it must leave the network, according to the network leaving procedure.

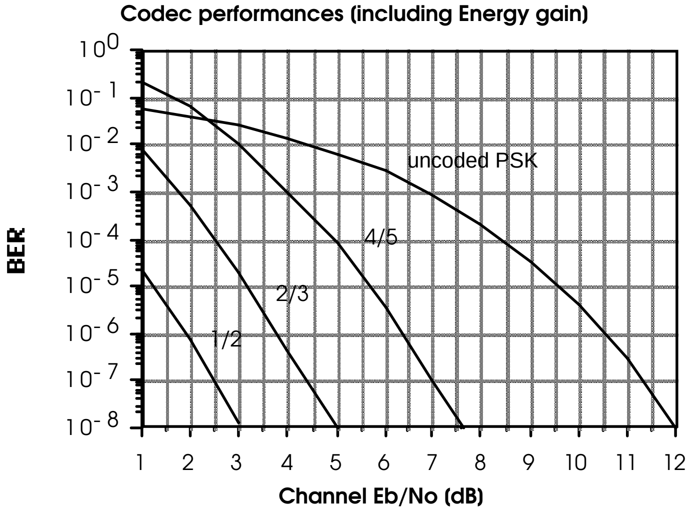

Fig. 3. BER versus E_b / N_0 for different coding rates

4.1 The non-transmission probability

The probability of a CUW miss or a control slot CRC error is approximately given by:

$$P_{nu} \cong p_{crc}(1-p_{miss})+p_{miss} \qquad (6)$$

The non-transmission probability due to a CUW miss or to a control slot CRC error is given by

$$p_{nt} = 1 - (1 - p_{nu})^N , \qquad (7)$$

where N is the number of active stations.

Figure 4 shows the non-transmission probability for several numbers of active stations using our available hardware. A preamble length equal to 216 total symbols (144 carrier + 40 reversal + 32 UW) with the CS information data 192 bits long and 4/5 coded has been assumed. The fade levels refer to the control slot E_b / N_0 reference value of 18 dB at 2 Mbit/s. All the stations are considered at the same fade level. Two couples of 16 bit UWs (one for each in-phase and quadrature channels of QPSK modulation) with a 4 bit threshold are considered.

Figure 5 shows the probability that a station is at a fade level higher than the abscissa value. The attenuation probabilities were obtained using the CCIR interpolation formula [14], giving the attenuation exceeded for a percentage of an average year

$$A(p_{\%}) = A_{001} \; 0.12 \; p_{\%}{}^{-(0.546+0.043 \; Log_{10} \; p_{\%})}$$

where $p_{\%}$ is the probability as a percentage, and A_{001} is the attenuation exceeded for 0.01 % of the time. A_{001} was taken equal to 22.5 dB for the up-link and 12 dB for the down-link, respectively, according to data measured at the Italian station in Fucino [7, 8]. The fade level at the receiver input is computed by combining the up and down atmospheric attenuations, but assuming that each station is capable of countering up to 10 dB of up-link fade attenuation by means of up-power control. If we cross the data in Fig. 5 with the data in Fig. 4 we see, for example, that the probability that a station has a fade deeper than 7 dB during a year's period is less than 10^{-3}, and that the probability of non transmitting is less than $3 \cdot 10^{-5}$ if 8 stations at the same fade level are present in the system. In fact, unless the stations are geographically close together, the non transmitting probability is smaller, because each station experiences its own fade level, while the plot in Fig. 4 assumes that all the stations have the same fade level.

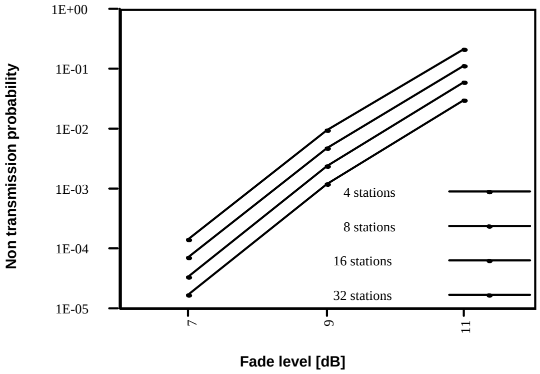

Fig. 4. Non-transmission probability for stations at the same fade level

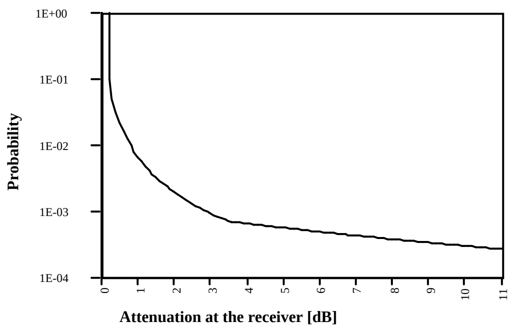

Fig. 5. Probability of a station with a fade deeper than a given value

While a comprehensive evaluation of the non transmission probability of a station for a period of a year would require a Monte Carlo integration—which will be part of future work—a reasonable approximation can be found with the same criteria used in [12], where the FEEDERS/TDMA scheme is presented. Since we have no experimental data for p_{miss} in the 12÷18 E_b / N_0 range, because the CUW miss events are too rare to be measured, we assume that p_{miss} in the 12÷18 dB range is the same as p_{miss} at 12 dB. We also assume that all the other parameters are equal to those used in the FEEDERS paper [12], with a minimum E_b / N_0 level of 7 dB at the receiver input. Such hypotheses lead us to a conservative estimate of the average p_{miss} of $5 \cdot 10^{-6}$ per station in a system with 32 active stations.

4.2 Critical events

We will now analyse the cases when a station corrupts the data transmitted by another station (collision) because of a miscomputing of the BTP. Such events are very unlikely because a station always transmits its CS in a fixed place in the frame, and it refrains from transmitting data if it is not able to build the BTP (non transmission event). Consequently, a collision can only happen when a station’s CS position changes; specifically, when some other station leaves or enters the system. Let us examine the possible cases:

- Station A unwillingly stops its transmissions. Station B is in fade, so either its fault countdown is not the first to start (because its M value is high), or it is not the shortest one, or both. Station B does not understand the fault countdown of the other stations, so it does not move its CS to the new location at the right time. This means that the probability of this event happening is always less than P_F which is typically less than 10^{-8} per dying station.
- Station A is erroneously declared dead by station B, and station A does not understand station B' s fault countdown. The probability of this event happening is always less than $P_M \cdot P_F$ which is typically less than 10^{-16} per dying station.
- Station A willingly leaves the system, while station B does not understand its countdown. The probability of this event happening is always less than P_F which is typically less than 10^{-8} per each station leaving the system.

- Station A enters the system, and station B does not understand neither station A nor all the other stations' confirmation countdown. The probability of such an event per each entering station is less than $2\Lambda T_f p_{nu} P_Q$ where P_Q is the probability of not understanding the confirmation countdowns, which can be reduced at will.

5. The distributed control overhead

The channel overhead due to the distributed synchronisation is equal to the CSF bit length. If U is the number of stations accommodated in each frame, we have

$$\text{CSF length} = U(L_c + T_{ovh}) \qquad (8)$$

where L_c is the control slot bit length and T_{ovh} is the control slot transmission overhead due to the UW, the preamble and the guard time between two control slots.

With the hardware currently at our disposal one control slot occupancy is equal to the 2.5% of the frame, considering:

- 4*16 bit UW,
- 144 carrier symbols
- 40 reversal symbols
- 192 bit control slot data (including CRC) ,
- 384 bit guard time,

all of which are transmitted at 2 Mbit/s, which is four times slower than the maximum bit rate of 8 Mbit/s allowed by the hardware considered. If we assume that no more than eight CSs are set up in each frame, then the CSF produces an overhead of 20% of the available bandwidth.

6. Conclusions

A distributed-control TDMA satellite access scheme has been presented; the possible system inconsistency problems and collision events have been analysed.

The system has the potential advantage—over the centralised FODA/IBEA system from which it derives—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 K$_a$ 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.

The simulation results reported in [13] highlight that with any of the types of traffic used (Poisson, Two-states Markov-modulated Poisson, and Fractal Gaussian Noise), DRIFS performs better than FEEDERS (the other distributed-control access scheme that we studied) from 32 stations on. This is because DRIFS can have an allocation cycle bigger than the frame size, so the overhead is lower than the FEEDERS’s one for a large number of stations. The simulation results presented in [13] also give suggestions for an improved distributed algorithm, which is a compromise between FEEDERS and DRIFS. This scheme should work like FEEDERS when the number of stations is less than 32, then if the number of stations is increased, it should enlarge the allocation frame to C_c frames, as designed for DRIFS.

References

[1] GEC-MARCONI Research Centre
"Olympus TDMA equipment. Transmit and receive controllers. Hardware manual. Issue 2.", MTR 91/60A, Y/212/10442, September 1991.

[2] GEC-MARCONI Research Centre
"Olympus TDMA equipment. Burst mode modem equipment manual", MTR 91/60A, Y/212/10443, September 1991.

[3] Celandroni N., Ferro E., James N., Potorti' F.
"FODA/IBEA: a flexible fade countermeasure system in user oriented networks", International Journal of Satellite Communications, Vol. 10, No. 6, pp. 309-323, November-December 1992.

[4] Celandroni N., Ferro E., Potorti' F.
"FODA/IBEA System description. Final Report.", CNUCE Report C94-18, September 1994.

[5] Celandroni N., Ferro E., Potorti' F.
"Experimental results of a demand-assignment thin route TDMA system", International Journal on Satellite Communications, Vol. 14, N0. 2, pp. 113-126, March-April 1996.

[6] Nai-Ting Hsu, Lin-Nan Lee
"Channel scheduling synchronisation for the PODA protocol", ICC' 78, pp. 42.3.1-42.3.5.

[7] Capsoni C.,Matricciani E., Mauri M., Paraboni A.
Il "Frequency Scaling" - Programma Sirio - I principali risultati dell' esperimento di propagazione di onde elettromagnetiche a frequenza superiore ai 10 GHz. Vol II pp. 117-156. CNR-CSTS - Roma, 1983.

[8] Carassa F., Matricciani E., Tartara G.
"Frequency diversity and its applications", International Journal of Satellite Communications, Vol. 6 No. 3, pp. 313-322, July-September 1988.

[9] Carassa F.
"Technical aspects in the future development of satellite communication systems with particular reference to the use of frequencies above 10 GHz', proceedings of the 19th Int. Conf. 'Scientifico sullo spazio', Rome, March 1979, pp. 22-41.

[10] Lam S.S.
"Satellite packet communication. Multiple access protocol and performance", IEEE Transaction on Communications, Vol. COM-28, No. 4, April 1980, pp. 468-488.

[11] Papantoni-Kazakos P., Hall Thornton
"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] Celandroni N., Ferro E., Potorti' F.
“FEEDERS-TDMA. A distributed-control algorithm for satellite channel capacity assignment in a mixed traffic and faded environment”, submitted to the International Journal on Satellite Communications.

[13] Celandroni N., Ferro E., Potorti' F.
"Comparison between distributed and centralised demand assignment TDMA satellite access schemes", International Journal on Satellite Communications, Vol. 14, N0. 2, pp. 95-112, March-April 1996.

[14] CCIR Report 564-4, "Propagation data and prediction methods required for earth-space telecommunication systems", 1990.

APPENDIX A. STATE TRANSITIONS OF A STATION