DELAY ANALYSIS FOR INTERLAN TRAFFIC USING TWO SUITABLE TDMA SATELLITE ACCESS SCHEMES

Nedo Celandroni*, Erina Ferro*, Francesco Potortì*
and
Gerard Maral#

*CNUCE, Institute of National Research Council
Via S. Maria 36 - 56126 Pisa - Italy
Phone: +39-50-593207/593312/593203
Fax: +39-50-589354 - Telex: 500371
N.Celandroni|E.Ferro|F.Potorti@cnuce.cnr.it

#Ecole Nationale Superieure des Telecommunications (TELECOM Paris)
Site de Toulouse, BP 4004, 31028 Toulouse Cedex, France
Maral@tlse.enst.fr

SUMMARY

Two demand assignment TDMA satellite access schemes that are suitable for providing flexible interconnections of local area networks are compared by means of simulation: FODA/IBEA and CFRA. A unique simulation tool was used, a meshed network of VSATs under the control of a master station was considered and similar networking conditions were applied to both access schemes. The generated traffic had two components having different level of burstiness: bulk traffic and interactive traffic. The performance evaluated is the delay for each traffic component. We show that CFRA is best when connecting clusters of only a few stations and even individual ones exchanging light traffic. FODA/IBEA, on the other hand, is better suited to interconnect networks, or within networks with many hosts, where heavy traffic is more likely.

Keywords: satellite, TDMA access schemes, comparison, bursty traffic, demand assignment, simulation.

1. Introduction

Satellite communications have evolved with respect to the use of satellite capacity. Initially, satellites were used to provide bandwidth to support television services and trunking telephony. The most popular access technique was Frequency Division Multiple Access (FDMA), whereby carriers shared fixedly the bandwidth of satellite transponders. Then demand assignment (DA) was introduced in order to use more efficiently the satellite capacity in the case of many telephone users with low demand, by varying the bandwidth

Preprint of an article published on the INTERNATIONAL JOURNAL OF SATELLITE COMMUNICATIONS, VOL. 15, No. 4
Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

allocated to earth stations with respect to variations in traffic demand. Demand assignment became even more popular with the advent of Time Division Multiple Access (TDMA) since it is easy to implement within TDMA satellite systems. With TDMA variable capacity can be allocated simply to earth stations either by varying the duration of the burst transmitted by each station, or by allocating a variable number of constant duration bursts per frame to each station, depending on the traffic demand from each station at the considered time.

Today, wideband local communication networks are being installed all over the world, and in many cases these Local Area Networks (LANs) need to be interconnected by satellite links. LAN traffic can be considered as an *aggregate traffic*, with two main components: *real time* traffic (telephony, video), and *non real time* traffic (computer data exchange). Real time traffic is generated at a nearly constant bit rate, and is therefore often referred to as *stream* traffic, while non real time traffic consists typically of data bursts, with silence in between, and therefore is often called '*bursty*' traffic.

The five applications responsible for bursty traffic leaving LANs are FTP (file transfer), NNTP (news distribution), SNTP (e-mail), TELNET (on-line conversations), and RLOGIN (remote login). These bursty applications can be classified into two categories:

- Bulk traffic (FTP, NNTP, SNTP)
- Interactive traffic (TELNET, RLOGIN)

Bulk traffic is characterised by a large quantity of bytes conveyed within a single burst, and a high value of interarrival time between bursts. On the other hand, *interactive* traffic has a small amount of bytes per burst and a low value of inter arrival time between bursts. Table 1 illustrates the typical characteristics of LAN traffic, where *burstiness* is defined as the peak-to-mean bit rate ratio.

Traffic Category		Characteristics	burstiness (on a frame time)
Stream		quasi-constant bit rate	1- 5
Bursty	bulk	many bytes per burst long inter arrival time	20 -900
	interactive	few bytes per burst short inter arrival time	4000 - 5000

Table 1. LAN traffic characteristics.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

Each traffic component also has different requirements: for instance, stream traffic should be routed from source users to destination users with a nearly constant delay, while delay is less critical for bursty traffic. The optimised use of the satellite capacity entails flexible demand assignment schemes that are efficient with respect to the various traffic components of the aggregated traffic, yet simple to implement.

Several TDMA schemes have been proposed for providing flexible interconnection of local area networks. These are:

- the Advanced Business Communication System (ABCS) [17],
- the Combined Free/Demand Assignment Multiple Access (CFDAMA) [19] [21] [22] [23],
- the FIFO Ordered Demand Assignment/Information Bit Energy Adaptive (FODA/IBEA) [14] [15] [16] [20] [18],
- the Combined/Fixed Reservation Assignment (CFRA) protocol [5].

ABCS is a protocol allowing access to the satellite by a network comprising a few tens of stations (typically 30) on a time share basis. Every station can reserve capacity according to its traffic load, and the fraction of the overall available capacity is allocated proportionally to the demand. Therefore, the distribution of capacity among network stations is fair. However, the protocol makes no distinction between traffic types, and therefore stream traffic suffers from similar delay and jitter as the bursty traffic, although it is more sensitive to such impairments.

CFDAMA combines free assignment with demand assignment techniques. At low and medium traffic loads it provides an average transmission delay shorter than that of the demand assignment schemes, while maintaining the high channel utilization.

This paper focusses on the two latter protocols and compares their performance under similar conditions of traffic and for a given satellite network. These protocols were elaborated by the two research teams that collaborated in the comparison study reported in this paper, and as such the detailed implementations of these protocols are better known to the authors than the other ones cited. Both FODA/IBEA and CFRA were simulated by using a simulation tool named *FRACAS*[1], and the results are compared.

---

[1] FRAmed Channel Access Simulator

The paper is organised as follows. Section 2 introduces the network organisation and the two access schemes to be compared. Section 3 presents the problems encountered in setting up a common testbed. Section 4 outlines the simulator, while Section 5 describes the network models. Section 6 shows the simulation results, and Section 7 contains the conclusions.

2. Network organisation and protocol description

Figure 1 illustrates the network organisation considered in this work. It consists of 16 traffic earth stations (VSATs: Very Small Aperture Terminals) accessing a portion of the transponder bandwidth in time division multiple access (TDMA) mode, and a central network management station (or master) allocating the capacity to the traffic earth stations according to their requests. Signalling messages generated by the VSATs can be conveyed to the master either on an out-of-band signalling channel, or piggybacked on the traffic. Allocation messages generated by the master station are transmitted to the VSATs on a dedicated broadcast signalling channel. The VSAT network is a meshed network, i.e. connections between VSATs are direct and established on demand. The protocols considered in this paper to access the transponder partition are FODA/IBEA and CFRA. Each protocol has its own TDMA frame structure, and a different burst structure. Figures 2 and 3 show the TDMA frame structure for FODA/IBEA and CFRA, respectively.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

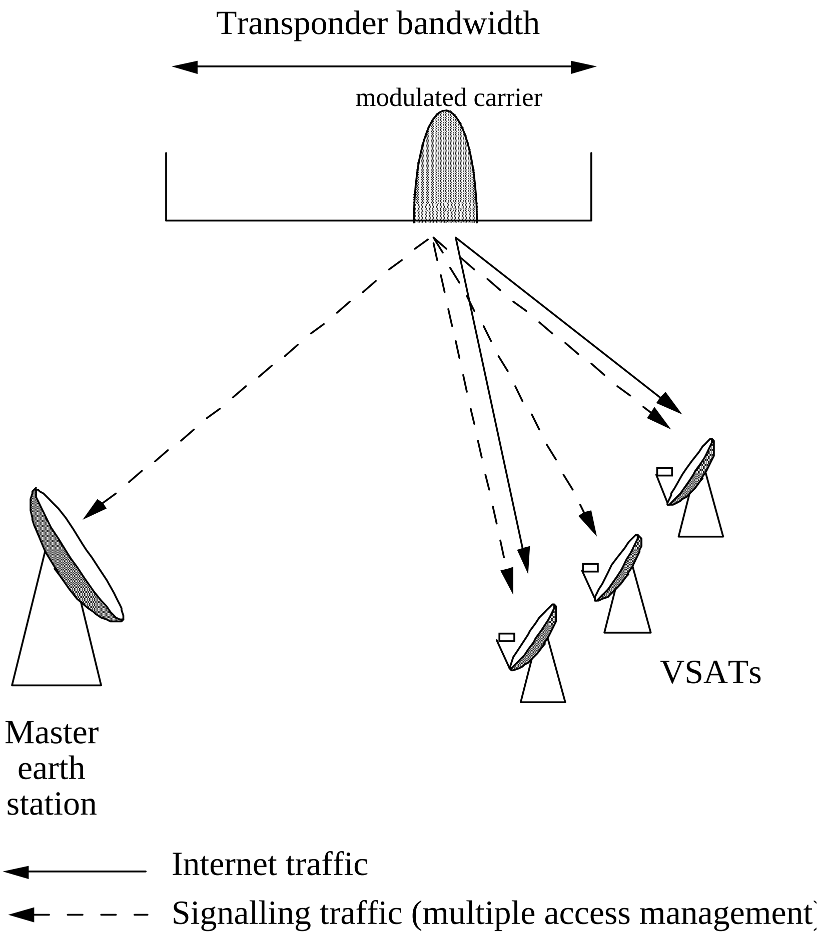

Fig. 1. Network organisation

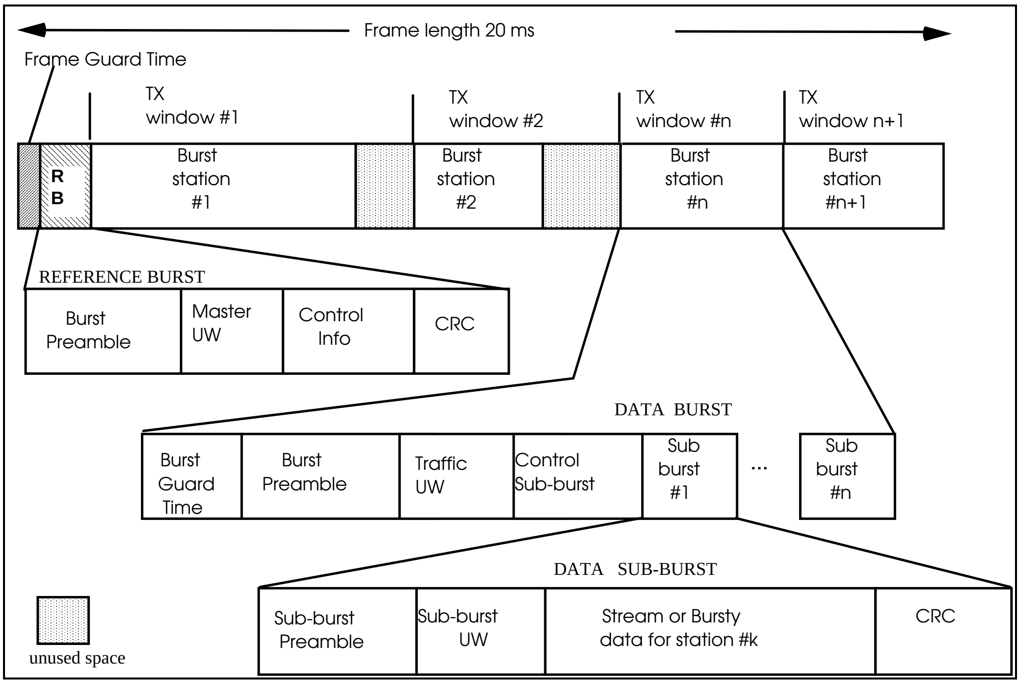

Fig. 2. Frame structure in FODA/IBEA

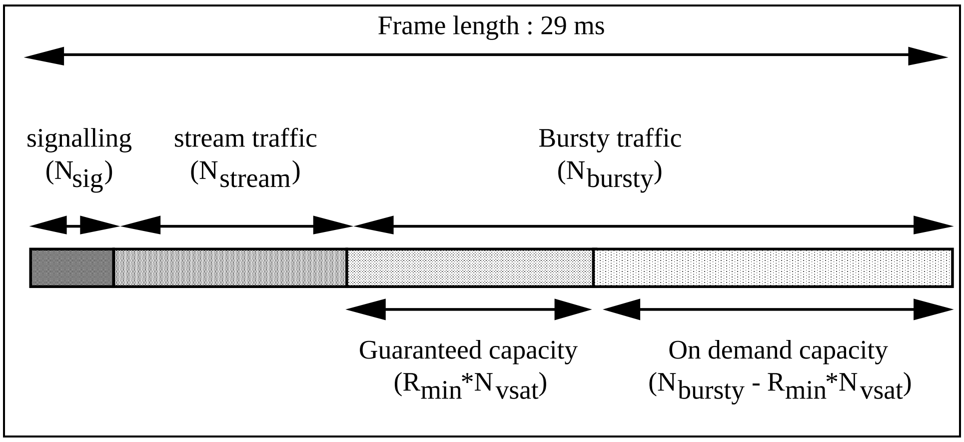

Fig. 3. Frame structure in CFRA

The different complexities of the two frame formats are mainly due to:

- the preambles and postambles that are implemented in FODA/IBEA and are not required in CFRA, which considers a preamble-less modem. In addition to the standard overheads of any burst-mode modem (preambles and unique words-UW-), the special features of the modem used by FODA/IBEA need sub-burst unique words and preambles, and control sub-bursts information, in order to manage at the correct bit and coding rates the data sub-bursts;
- the fade countermeasure technique, which is implemented in FODA/IBEA and does not exist in CFRA. The fade countermeasure feature of FODA/IBEA has not been taken into account in the comparison between the two protocols, but it must be underlined that this feature is, above all, responsible for the complexity of the whole FODA/IBEA system.

Another difference between the two schemes is that while FODA/IBEA can manage packets of any length, CFRA was thought to support constant size cells. For this reason, a traffic organised into ATM cells was considered during the simulation for the comparison of two schemes.

The rest of this section will present both protocols in more detail. In this description the parameters are given with values that correspond to optimised values for each protocol, considering the retained frame duration.

2.1 FODA/IBEA

FODA/IBEA is a protocol that incorporates fade countermeasure techniques. These features were not used in our simulations so the behaviour of FODA/IBEA in clear sky conditions will be described first. How it copes with rain attenuation will be dealt with later.

The FODA/IBEA frame, which is 20 ms long, can accommodate stream transmissions up to a threshold *NSUB* (Normal Stream Upper Boundary). The remaining space in the frame is reserved for bursty data. If no stream traffic is present, bursty data can occupy the whole frame, decremented by the space needed for signalling. The assignment for stream data is in the same quantity as the request. Once assigned, the allocation is maintained until the application returns the bandwidth to the master station or the VSAT ceases to transmit. A sending VSAT transmits the stream data in the same transmission window together with its bursty data, when present, in order to save the channel overheads due to burst preambles. The request for bursty data is $r = q + H\ i$ where q is the amount of traffic waiting in the queue for transmission in bytes (backlog), i is the incoming traffic, in bytes per frame, and H has the dimension of a time. The product of H by the traffic entering the station is a sort of prediction of the backlog for the future time H. The prediction is made considering the input traffic only, without decrementing it by the output traffic. This feature makes the system faster in reacting to positive variations of the traffic level with respect to negative variations. Therefore, a too high value of H reduces considerably the efficiency of the system. The best value of H has been chosen simulating a system close to saturation with Poisson traffic generators and looking at the average delay of all the stations. The best value resulted to be 0.4 s, and it is practically independent of the number of active stations. This value, in fact, gives the best performance at high channel loads, while it does not significantly increase the delay, at low-medium loads, with respect to lower values [28].

Bursty requests are organised by the master into a ring, and cyclically scanned to compute the amount of capacity for each assignment. New bursty requests are put at the current head of the ring in order to be scanned first. This reduces the delay time between the first request and the assignment after a period with no transmission. The capacity of the bursty assignment is proportional to the request in a range of values between a minimum and a maximum threshold (T_{min} and T_{max}, respectively). The coefficient of proportionality (*f* in Table 2) is equal to the number of active stations N divided by 100, with 5% as minimum and 50% maximum. After each assignment, the bursty request is decreased by the assignment itself and the next request is analysed, if capacity is still available in the frame. The first assignment that does not fit entirely into the current frame will be re-considered as the first assignment in the next frame, where the rest of the computed amount is assigned. The ring is scanned at most once per frame. If capacity is still available in the frame after a complete scan of the ring, the available space is shared among the active VSATs. Each VSAT is granted at least one assignment every eight frames, though usually every station gets an allocation in every frame.

The influence of the H and f values on a transient due to a step of traffic is analysed in detail in [30].

FODA/IBEA is also designed for operation under fade conditions due to rain. This feature makes the system particularly suitable for working in the 20/30 GHz band, where the signal attenuation due to bad atmospheric conditions is quite high. In order to cope with different levels of the signal attenuation, the system varies the energy contained in an information bit. This is done by dynamically varying the transmission power when possible, along with reducing the data coding rate (uncoded, 4/5, 2/3, and 1/2 are the possible coding rates) and the data bit rate. According to the fade level detected, the modem is capable of dynamically adjusting its transmission rate within a data burst. This allows the individual data sub-bursts of a data burst to have different symbol rates (and, hence, different energies) as required. The symbol rates available are 512, 1024, 2048 and 4096 kbaud, using either BPSK or QPSK modulation schemes. A bit rate range of 512-8192 kbit/s is thus available for the system.

The fade countermeasure technique implemented in FODA/IBEA is quite different from other fade countermeasures, such as site diversity [25] and frequency diversity [26]. Those two methods allow a very high level of link availability but they are very complex and expensive, as they require or a duplication or some additional hardware. With respect to those, the fade countermeasure implemented in FODA/IBEA is cheaper (it does not require any additional hardware, being implemented in software) and it can support heavy fade levels. With respect to other fade countermeasure techniques operating in TDMA, such as the burst length control technique [27], it results in a better optimisation of the satellite channel use and can compensate for higher levels of attenuations, by achieving a gain also by reducing the data bit rate in addition to the coding gain.

2.2 CFRA

Unlike FODA/IBEA, CFRA handles cells of constant size and makes use of a preamble-less modem. The frame duration is 29 ms (Figure 3). It consists of 132 time slots with 55 bytes per slot, and 10 bits of guard times between bursts. Each time slot can contain an ATM cell whose 5 byte header is extended to 7 bytes using a block coding for header error protection.

The frame is organised into three sub-frames. The first sub-frame begins with a number, N_{sig}, = 4 time slots where signalling cells are transmitted from all VSATs to the master earth station. The second sub-frame contains N_{stream} slots for transmitting stream traffic cells.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

The last sub-frame has N_{bursty} slots for transmitting bursty traffic cells. These latter slots are divided into two parts: one with a number of fixed capacity slots equal to $R_{min}\,N_{vsat}$ (where N_{vsat} is the number of VSAT stations), and the other consisting of the slots available for assignments on demand. Stream traffic is handled by means of fixed assignments (with a capacity equal to the one requested) as this type of traffic usually requires a fixed delay, which cannot be guaranteed by demand-assignment techniques.

The bursty traffic is handled according to a combination of a fixed assignment, for *short bursts*, and a demand-assignment, for *long bursts*. The distinction between short and long bursts is made by the L_o parameter, expressed in *numbers of cells*. A burst consisting of less than L_o cells is *short*, and *long* if it consists of more than L_o cells. Two bursts are considered as distinct and consecutive when separated by an entire frame without any cell delivered to the sending VSAT by the terminal connected to it. The short burst, when stored in the station input buffer before transmission, does not trigger the demand for more capacity than the standard R_{min} value. When a long burst is recognised, a demand in the form of a *start of burst* (S) message is sent by the VSAT to the master to get the larger capacity R_{max}. The VSAT continues to transmit the burst with the R_{min} capacity until it receives the assignment for the larger allocation. When the long burst ends, the VSAT transmits to the master the *end of burst* (E) message, which initiates the de-allocation of the capacity $R_{max} - R_{min}$. The R_{min} and R_{max} constants are set to 1 and 30 slots per frame, respectively, while the constant L_o is set to a burst length of 10 ATM cells. These seem to be the best values for bursty traffic, as the results in [5] show.

Table 2 highlights the most important characteristics of the two protocols.

3. Harmonisation problems

FODA/IBEA and CFRA are two very different systems for many reasons. They are conceived to provide both real-time fixed bandwidth and allocation on demand to requesting applications, but they rely on different hardware and have different implementation statuses.

FODA/IBEA uses a 5 MHz channel bandwidth to provide up to an 8 Mbit/s gross information rate. It is based on stable technology using burst preambles, which constitute a significant overhead. Input packets can be of any size. They are assembled together when sent on the satellite channel, and disassembled when received from the satellite. The fade countermeasure method is the most remarkable feature of the system, allowing it to survive up-link plus down-link fades of up to 25 dB [31], while maintaining the stated BER, at the

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

expense of the usable bandwidth. Any of the stations can assume the role of master for synchronisation and capacity allocation purposes. Prototypes of the system have been working for some years.

	CFRA	**FODA/IBEA**
control	centralised	centralised
frame length	29 ms	20 ms
modem	preamble-less	needs preamble
channel bit rate	4 Mbit/s	1÷8 Mbit/s
FEC coding rate	1/2	none, 4/5, 2/3, 1/2
stream allocation	as requested	as requested
bursty request	S message when burst longer than L_o; E message to release the extra capacity	$r = q + H\, i$, where r is the request q is the input queue length H is a temporal constant of proportionality i is the instantaneous traffic
bursty allocation	R_{min} fixed; R_{max} on receiving the S message	min (T_{max}, max ($f\, r$, T_{min})) f is the coefficient of proportionality r is the request

Table 2. Comparison between the CFRA and FODA/IBEA protocols

CFRA uses a 3 MHz channel bandwidth to provide a 2 Mbit/s gross information rate. Input packets are 53 bytes long for ATM compatibility. CFRA is based on advanced technology using a preambleless modem, which operates virtually without any overheads. A fixed 1/2 rate convolutional encoding is used, and the first 5 bytes of each input cell are further protected with block coding.

Since the two systems are so different, we compare the performance of the allocation schemes alone, assuming that a net information bandwidth of 2 Mbit/s is available for both systems. This approach avoids all the problems of comparing the overheads of the two systems, allowing one to concentrate on the relative merits of the “pure” allocation schemes. However, there is a significant drawback to this simple comparison scheme, in that FODA/IBEA’s overhead is not fixed, but increases with the number of stations transmitting in a frame. Since the higher the traffic load, the lower the number of stations in a frame is— and hence the overhead- this comparison puts FODA/IBEA at an advantage with respect to CFRA in situations of very high traffic loads. This advantage can be considered as if the

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

total bandwidth available to the system were greater for FODA/IBEA at high loads by a quantity that is always less than 5%. This 5% is an upper boundary which is only reached very rarely (when the channel is completely saturated with traffic). Our results highlight the general behaviour of the two schemes, and are accurate for each individual scheme. However, for each simulated point in the FODA/IBEA graphs, the load is not necessarily exact, especially at high traffic loads, unless a safety error of [0; -5%] is applied to it. In practice, this error detracts nothing from the conclusions that can be obtained from the simulations.

4. The simulator

Both FODA/IBEA and CFRA were simulated using *FRACAS*, a C language specialised emulator developed at CNUCE in order to evaluate the performance of TDMA satellite access schemes [4]. The bandwidth request and allocation policies can be chosen from a set of predefined ones, included in the *FRACAS* library, and all the relevant network configuration parameters are tuneable. Each station in the network can be configured with a number of different deterministic or stochastic traffic generators, which produce traffic that can be classified in either of four classes, namely *cbr* (constant bit rate traffic), *vbr* (variable bit rate traffic), *interactive* , and *bulk*. For each run, various statistics can be collected for subsequent analysis. *FRACAS* can be extended to include more traffic generators, access schemes, and statistics collectors. It is written in C, so it is very efficient, which allowed us a simulated run time of 4000 seconds for each of the points in our graphs.

FRACAS is a *discrete time* emulator with a granularity of one frame. Since all the timings are multiples of a frame, no delay shorter than this quantization unit can be resolved. This implies, for example, that if in the real system the delays are less than one frame long, they are rounded down to zero. In general, the delays obtained with *FRACAS* have an error in the interval [0; -1 frame] but, unless the distribution of the delay is concentrated over a few frame’s time, it can be approximated with an error of -1/2 frame.

As far as the comparison between FODA/IBEA and CFRA is concerned, this means that the difference (FODA/IBEA delay) - (CFRA delay) has an error in the interval [-20; +29] ms, because 20 and 29 ms are the frame lengths of the two access schemes, respectively. The mean error that derives from this effect is therefore +4.5 ms.

The reliability of *FRACAS* was tested using traces obtained from measures on the prototypal implementation of FODA/IBEA. The output of the emulation differs from the measures only because of the above discretisation errors. This result makes us confident about the reliability of the emulation of the FODA/IBEA system. As far as the CFRA system is concerned, some preliminary tests were conducted with the use of two emulators:

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

OPNET and FRACAS. The models of the CFRA protocol were refined and debugged until no difference greater than the statistical uncertainty was found between the results obtained with the two different emulators.

The 95% confidence interval of the simulation results described in the rest of the paper is less than ±5% for all the values resulting from the simulation.

5. Traffic models

In references [6, 7, 8, 10, 11, 24] the most important results of the ongoing debate on the self-similar nature of internet traffic are analysed and discussed. Aggregates of traffic from a large number of sources whose traffics exhibit long range dependency appear to be self-similar, and can be approximated well by Fractional Gaussian Noise. Therefore, a Fractional Gaussian Noise process can be considered as a reasonable approximation of LAN traffic generated by many hosts [7, 8, 10, 11]. Indeed, if one observes the amount of traffic on the time scale of a frame, the distribution density of LAN traffic resembles a Gaussian function, reflecting the combination of traffic generated by many sources.

A different model has to be adopted, however, for internetwork traffic and single host traffic which differs from LAN traffic [8, 11, 24]. Here, the traffic is generated by a small number of sources and is far from being Gaussian. This traffic is much burstier, in the sense that its peak/average ratio is higher and most often no traffic at all is observed on the time scale of a frame, so the distribution of the number of packets generated per frame peaks off at zero.

On the basis of the above considerations, we chose two basic traffic types generators: the first one, named *composite traffic generator* (CTG), simulates the traffic generated by only one or by a small number of sources; the second one, the *fractal traffic generator* (FTG), simulates the traffic generated by a large number of sources. If the satellite network is used to bridge many LANs with high traffic, the latter modelling is probably a good fit for the real traffic.

The composite traffic generator is described in Figure 4.

CTG incorporates a bulk traffic generator and an interactive traffic generator. These generators were designed and implemented in the simulator so as to match the characteristics of wide-area TCP/IP conversations, as reported in [12] and [13]. The bulk traffic generator in Figure 4 delivers files with an exponential interarrival distribution, whose mean is chosen, for each simulation run, in order to get the desired mean bulk traffic load. The file length distribution is made by a *constant plus an exponential term,* where the *constant* is 100 bytes, and the mean *constant plus exponential* is 5000 bytes. In order to model the exchange of data, the file transmit process (Fig. 4) formats the file into 512 byte packets, which are then delivered at two different rates scheduled by a two state-Markovian machine (Fig. 5).

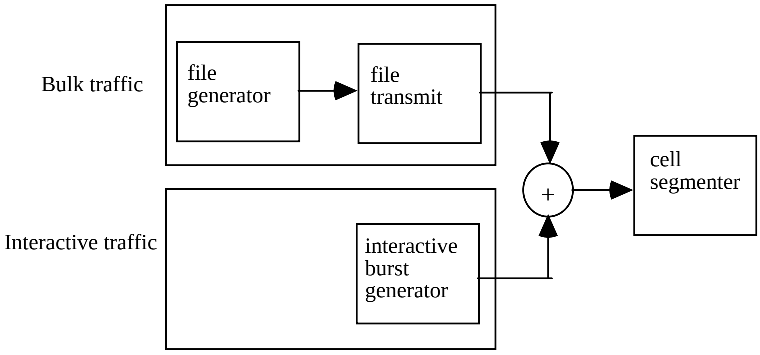

Fig. 4. The composite traffic generator generating LAN-like traffic

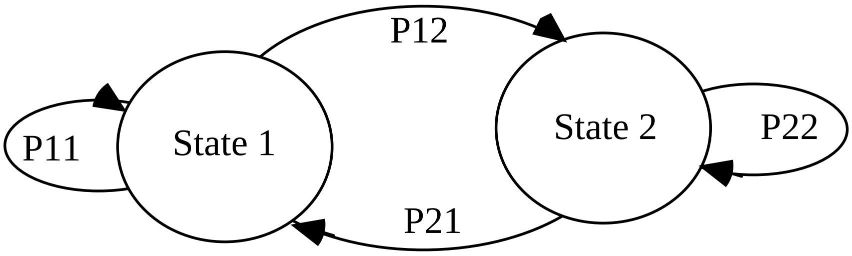

State 1 inter-arrival time: constant (5ms) + exponential; mean = 50 pkt/s
State 2 inter-arrival time: constant (5ms); mean = 200 pkt/s
packet length: constant; 512 bytes
P12=P21=0.2; P22=P11=0.8

Fig. 5. Two-rate model of a file transfer

The interactive traffic generator in Figure 4 delivers packets with a constant (10 ms) plus exponential inter-arrival distribution, with a mean chosen in order to get the desired interactive traffic load. The packet length is generated according to the histogram reported in Table 3, which gives a mean value of 165.5 bytes.

cumulative distribution	packet length
0.8	10
0.85	50
0.90	100
0.95	1000
1	2000

Table 3. Histogram for interactive traffic packet length

The fractal traffic generator used is a Fractional Gaussian Noise generator implemented using a Random Midpoint Displacement algorithm. An initial analysis of the statistical properties of this algorithm is given in [6], with results that we considered to be good enough for our purposes. This generator exhibits relatively low burstiness and a long-term correlation, which we truncated to about 10 minutes of simulation time. Our input traffic traces are thus made up of consecutive 10 minute long independent batches. Fractional Gaussian noise is defined by three parameters: the mean, the peakedness, and the Hurst parameter. For both interactive and bulk traffic we used a Hurst parameter equal to 0.85, in line with the findings in [7]. The peakedness, defined as the ratio between the variance and the traffic distribution mean values, was set at the maximum that would not make the generated trace become negative too often, that is, that would not make the aggregated traffic of all the stations noticeably different from a normal distribution. When the traffic trace generated by the algorithm becomes negative, it is truncated at zero, thus distorting the distribution of the traffic generated. The peakedness was set equal for all stations, and proportional to the traffic load, in order to obtain a constant *form factor*, which is defined as the ratio of the standard deviation and the mean of a distribution. The only difference in generating interactive or bulk traffic with the fractal traffic generator is the data generation granularity: interactive traffic is generated in 1-byte units, while burst traffic is generated in 32-byte units before packetisation.

Figure 6 shows the probability density of the traffic generated by CTG. The distribution includes 10% of interactive traffic and 90% of bulk traffic. The peaks are caused by the bulk generator, which produces bursts of 512 byte packets. The rest of the distribution is the effect of the interactive traffic superimposed on the bulk traffic. The value at the origin represents the probability of the traffic being null on the time scale of a frame, which is here equal to 66%.

Figure 7 shows the probability density of the traffic generated by FTG. As before, the distribution includes 10% of interactive traffic and 90% of bulk traffic. The distribution density is the sum of two Gaussian functions, truncated at zero. All the negative values of traffic are set to zero, so the value at the origin is the area of the negative tails of the Gaussian functions. The probability of the traffic being null on the time scale of a frame is low, only 1.8%.

The two types of traffic used in the simulation can be viewed as the two extremes of traffic multiplexing from different sources: the composite traffic generator represents well the traffic produced by a few stations, the fractal traffic generator reproduces the traffic behaviour resulting from the aggregation of a high number of stations.

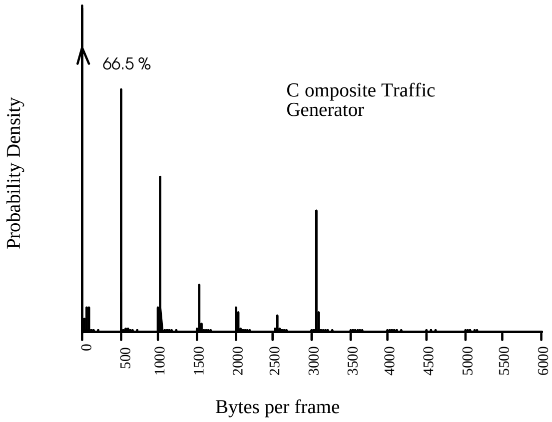

Fig. 6. Traffic distribution for the composite traffic generator (10% interactive traffic, 90% bulk traffic). High peaks at n*512 bytes account for bulk traffic, while low peaks for interactive traffic.

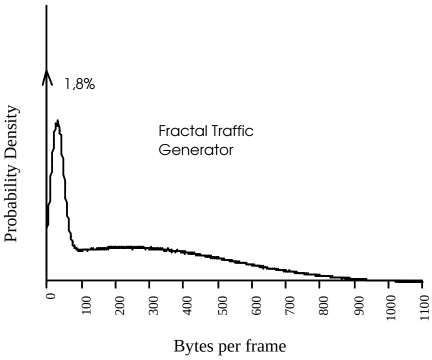

Fig. 7. Traffic distribution for the fractal traffic generator

6. Simulation results

To compare different access schemes and the influence of a type of traffic, a user of the simulation software can set up the mix of components in the aggregated traffic by choosing, as shown in Fig. 8, the percentage of the stream traffic (alpha) in the total traffic, the percentage of the interactive traffic (% inter) in the bursty capacity used, and the percentage of the bulk traffic (100 - %inter).

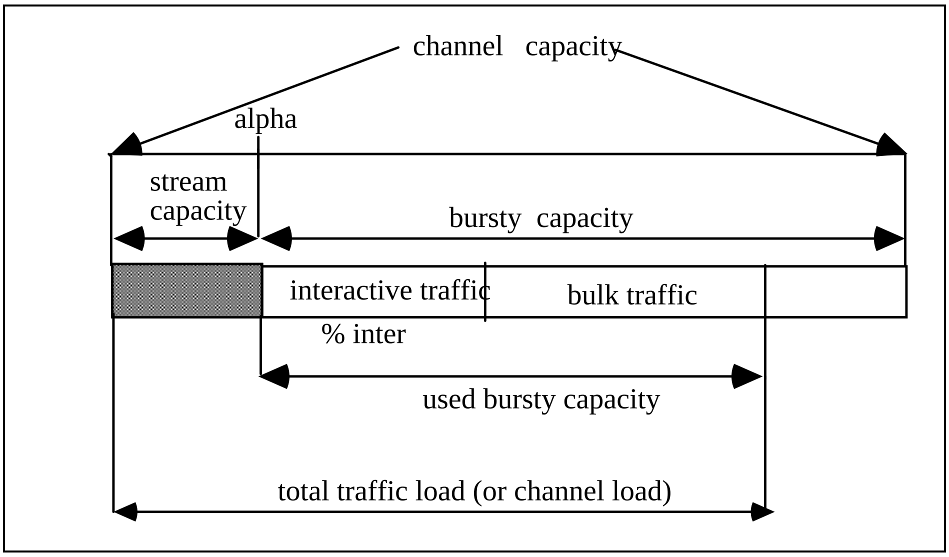

Fig. 8. The mix of components in the aggregated traffic

While both protocols allow us to allocate guaranteed bandwidth channels for stream applications, we did not use this feature in our tests (alpha=0, in Fig. 8). In fact, it would not have added much to the results, as the only effect of allocating a stream channel is, in both allocation methods, a corresponding reduction in the channel capacity available for demand assignment. All simulations are aimed at measuring the performance of the systems for bursty best-effort traffic.

In all the tests we used a network of 16 stations, but the results obtained can be scaled for a larger number of stations. In [29] results can be found of the simulation of FODA/IBEA considering a network with up to 48 stations, using Poisson, two-states Markov-modulated Poisson and Fractional Gaussian Noise traffic generators. As for CFRA, no simulation has been performed up to now for networks of varying size.

The 16 stations here considered generate bursty traffic alone, of the type usually found on LANs. A net information bandwidth of 2 Mbit/s is available for both access schemes. The evaluated performance is the delay experienced by the individual traffic cells when crossing the satellite network. The delay is thus the sum of the satellite round-trip delay, which we set to 254 ms, the processing delay, which we set to 1 frame time, and the input queueing delay. Our experience in the implementation of the FODA/IBEA access scheme shows that a processing delay of one frame is the minimum achievable unless most of the access scheme is implemented in hardware, which would be impractical for schemes of this complexity. Since the frame length of CFRA is greater than that of FODA/IBEA, CFRA's minimum delay is 283 ms, when no queueing occurs, while the minimum for FODA/IBEA is 274 ms. In the legends of all the figures F/I stands for FODA/IBEA.

Figures 9 and 10 show the mean delay of the cells crossing the network when all the stations are loaded with a mix of bulk and interactive data with equal average rates. The total traffic load varies from 10% to 90% of the channel capacity, while the quotas of interactive data (% inter) are 5, 10, 15, and 20% of the total traffic load, respectively. The mean cell delay averaged over all the stations is plotted separately for interactive and bulk traffic. When the load used is generated by the composite traffic generator (Figs. 9a, 10a) CFRA performs better than FODA/IBEA for low and medium traffic loads (in the range of about 30÷50% of the channel capacity), getting delays that are about 25÷50 ms lower, for both interactive and bulk traffic. At higher traffic loads the opposite situation occurs. FODA/IBEA performs better when loaded with fractional Gaussian noise-type traffic at all the load intensities (Figs. 9b, 10b). Note how the interactive traffic is privileged in both systems, since it has a mean delay lower than bulk in all situations.

Figures 11 and 12 show the length of the tails of the delay distributions. All the stations have the same load intensity, as in the previous case. The quota of the bulk traffic is 90% of the channel load, while the remaining 10% is devoted to the interactive traffic (%inter=10%). Four percentile values are represented, namely 90, 95, 99, and 99.9. As in the previous case, CFRA performs better at channel loads lower than 50% of the channel capacity when loaded with the composite traffic (Figs. 11a, 12a), while FODA/IBEA performs better at higher loads or, at all loads, when the channel is loaded with fractional Gaussian noise traffic (Figs. 11b, 12b).

Figures 13 and 14 show the ability of the two access schemes to handle unbalanced traffic. All the stations are equally loaded with 10% of interactive and 90% of bulk traffic (%inter=10%). Let us call this loading condition the *base load*. Three stations are loaded with an additional bulk traffic, which we name *overload*, adding a traffic load that is *ovl* times the base bulk load. Thus, for *ovl* = 0, all the stations have an equal load, while for ovl = 15, stations 1, 2 and 3 have a bulk load 16 (15+1) times greater than that of the other 13 stations. This arrangement highlights what happens in the network when three stations experience a sudden and possibly temporary exchange of data. These figures show the delays of both the overloaded and non-overloaded stations, for three different values of *ovl*: 3, 7, 15. The highest value of *ovl* emulates a situation in which the traffic of the three overloaded stations adds up to 79% of the total system load.

The results are essentially the same as in the previous tests: CFRA is better at handling the traffic of the composite generator for low and medium traffic loads, while FODA/IBEA performs better with high traffic loads and, at all traffic loads, with the fractional Gaussian noise traffic.

7. Conclusions

A common simulation tool, FRACAS, has been used to compare two different DA-TDMA satellite access schemes, FODA/IBEA and CFRA, which are considered as two good candidates for providing a flexible interconnection of local area networks over satellite links. The two protocols were tested under similar conditions of traffic and for a given satellite network.

The reliability of the emulator used was assessed by comparing its results for FODA/IBEA with measurements obtained by using FODA/IBEA on real satellite links. Two different traffic patterns were used: a composite traffic generator, which represents the traffic generated by only one or by a small number of sources, and a fractal traffic generator, which represents the traffic generated by a large number of sources.

The results show that CFRA displays lower delays at low and medium traffic loads (approximately less than 2/3 of the channel capacity) when the composite traffic generator is used. On the other hand, FODA/IBEA performs better when the channel is heavily loaded with the composite traffic generator and, for all the traffic loads, when a fractional Gaussian noise generator is used.

It can be concluded that CFRA is more suitable for connecting the clusters of only a few stations and even individual ones, via a satellite network, where the traffic load is light. FODA/IBEA, on the other hand, is better suited to interconnecting networks, with heavy traffic in between, and within networks with many hosts.

The behaviour of the two access schemes can be explained in terms of their respective mechanisms: FODA/IBEA privileges full exploitation of the channel capacity, at the expense of the queueing delay. CFRA performs best when bursts of traffic in input to a satellite terminal are frequently intermixed with periods of silence, as its allocation request algorithm privileges fast response times to rapid input traffic variations. With continuous traffic, it becomes less efficient. This difference between the access schemes is also reflected in the relative complexity of the algorithms used in the two allocation schemes, CFRA being much simpler to implement than FODA/IBEA.

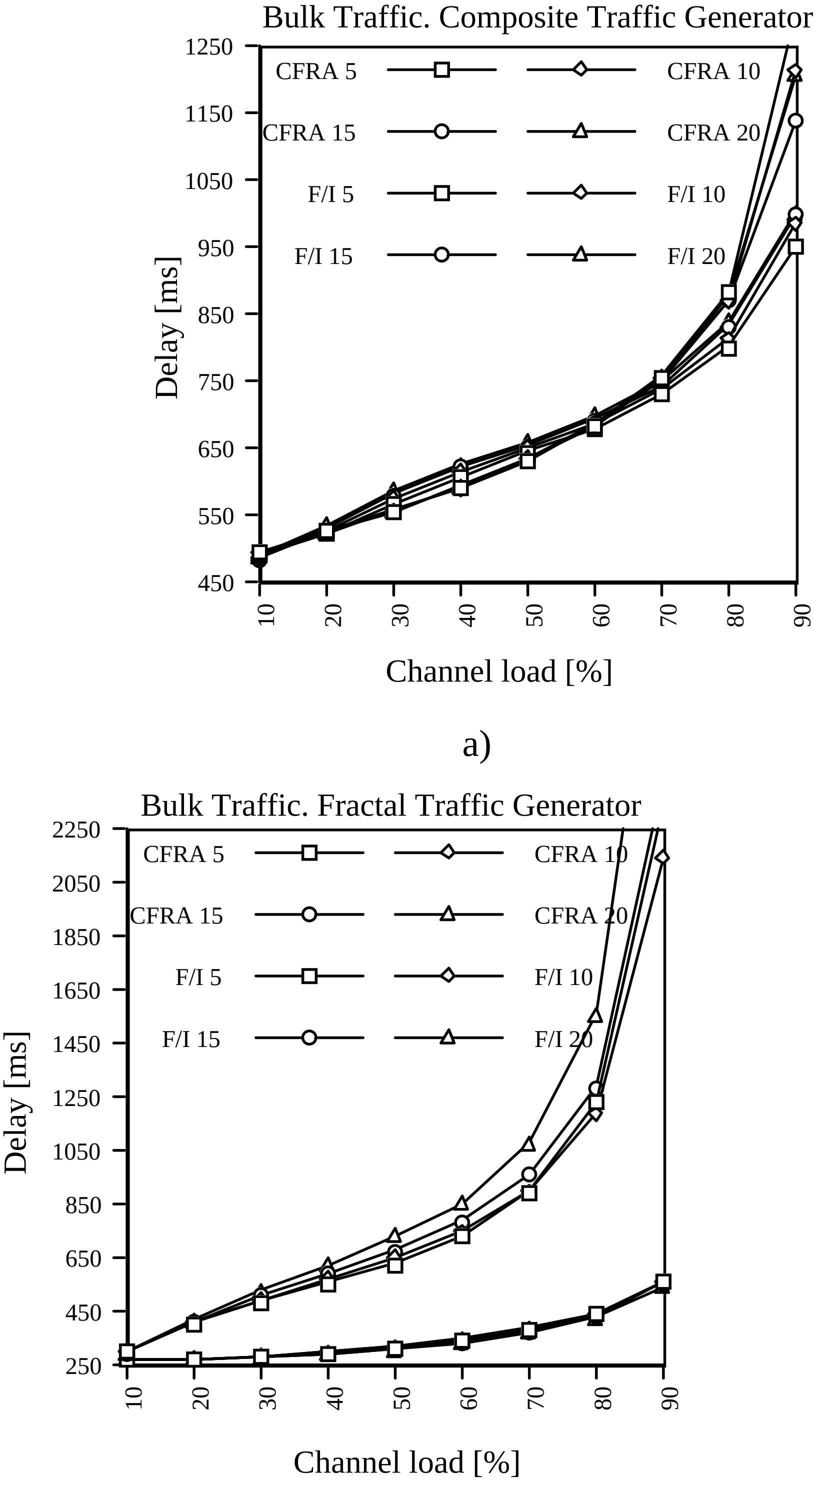

a)

b)

Fig. 9. Bulk traffic mean cell delay versus channel load for different quotas of interactive traffic (%inter). The load is distributed evenly among all the stations.

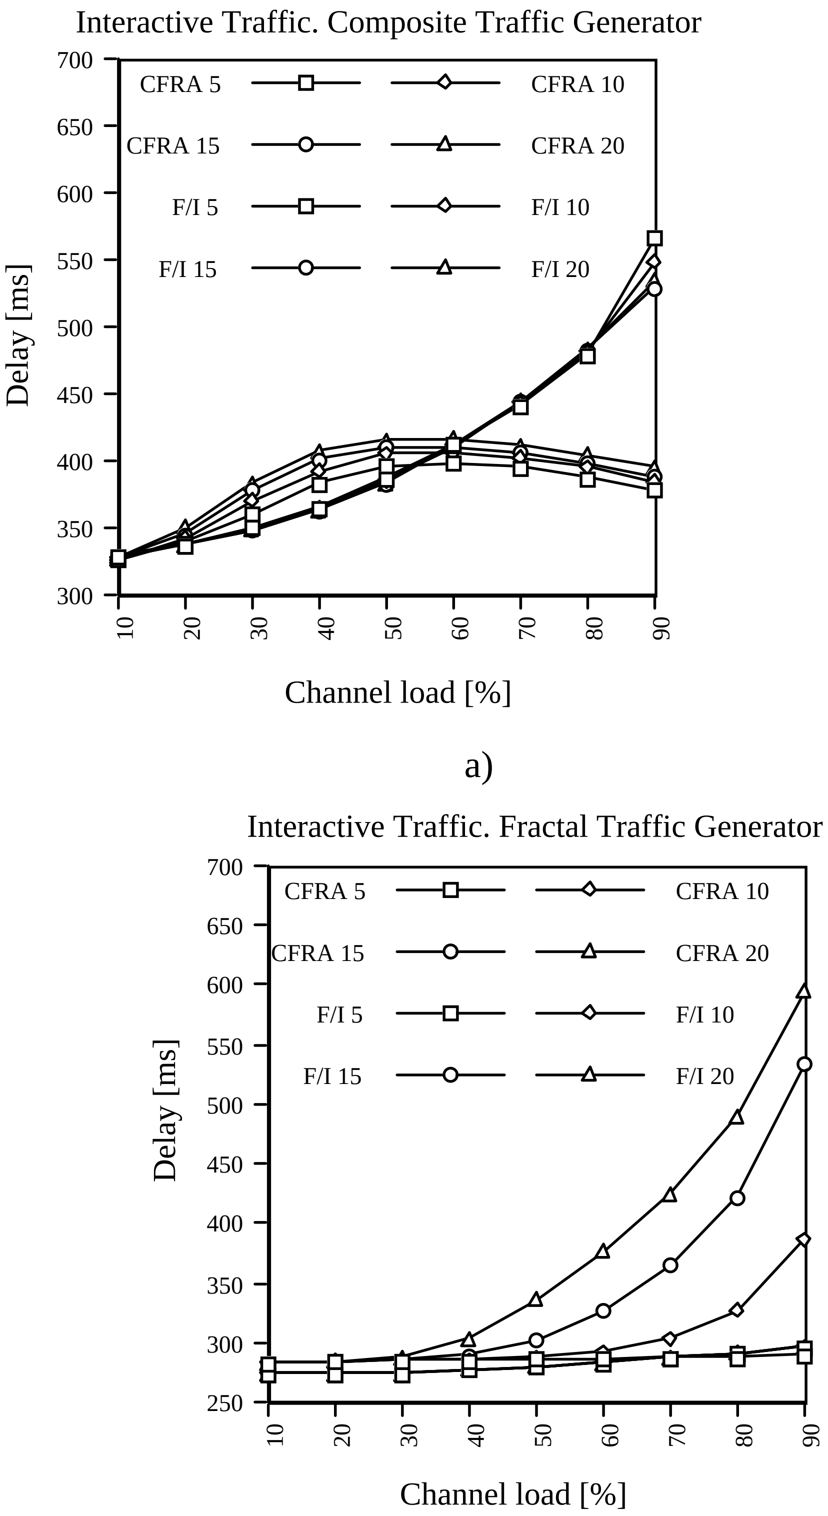

a)

b)

Fig. 10. Interactive traffic mean cell delay versus channel load for different quotas of interactive traffic (%inter). The load is distributed evenly among all the stations.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

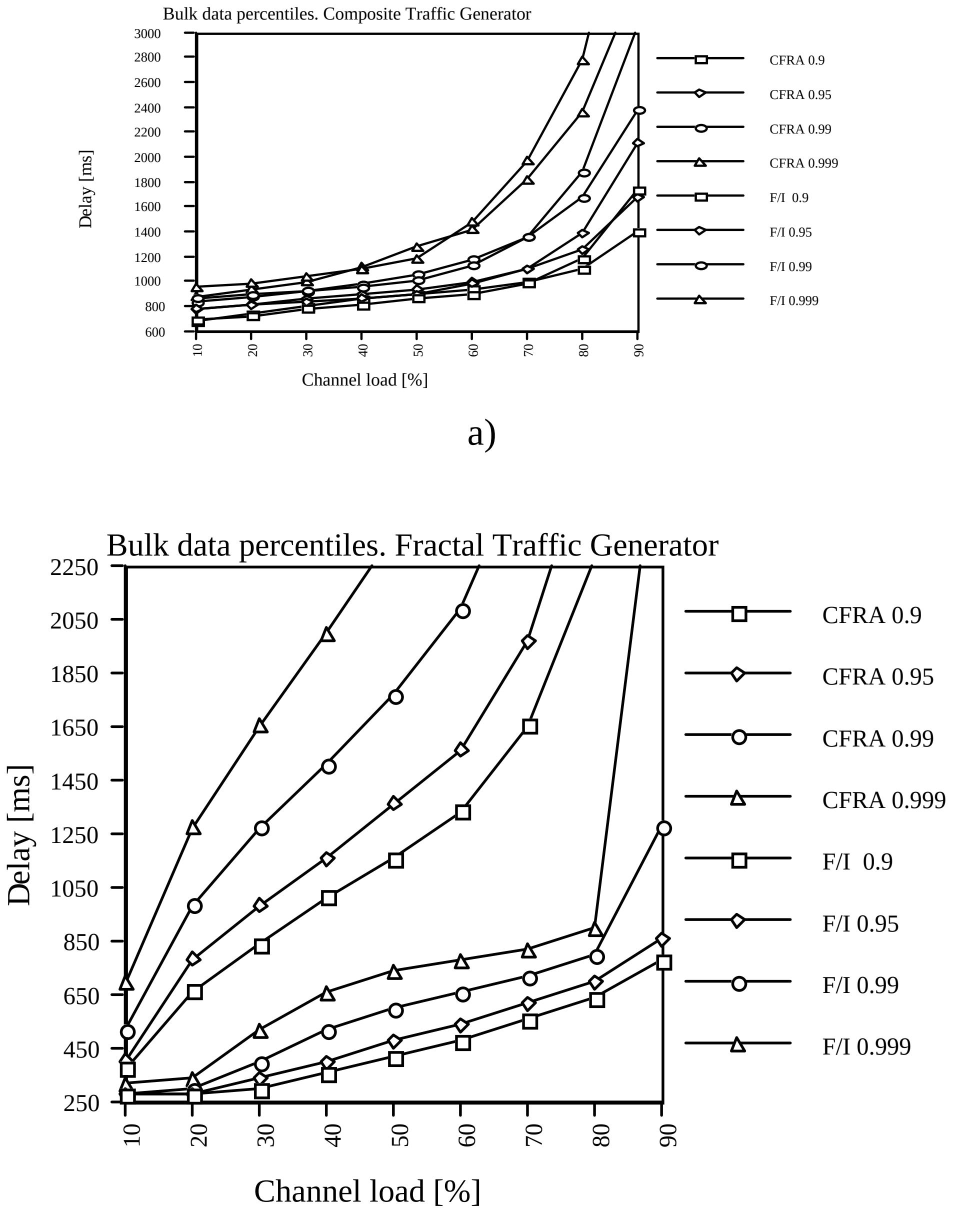

a)

b)

Fig. 11. Cell delay percentiles versus channel load when the interactive traffic is 10% of the channel load (%inter=10%). The load is distributed evenly among all the stations. Bulk data generated by composite traffic generator and fractal traffic generator.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

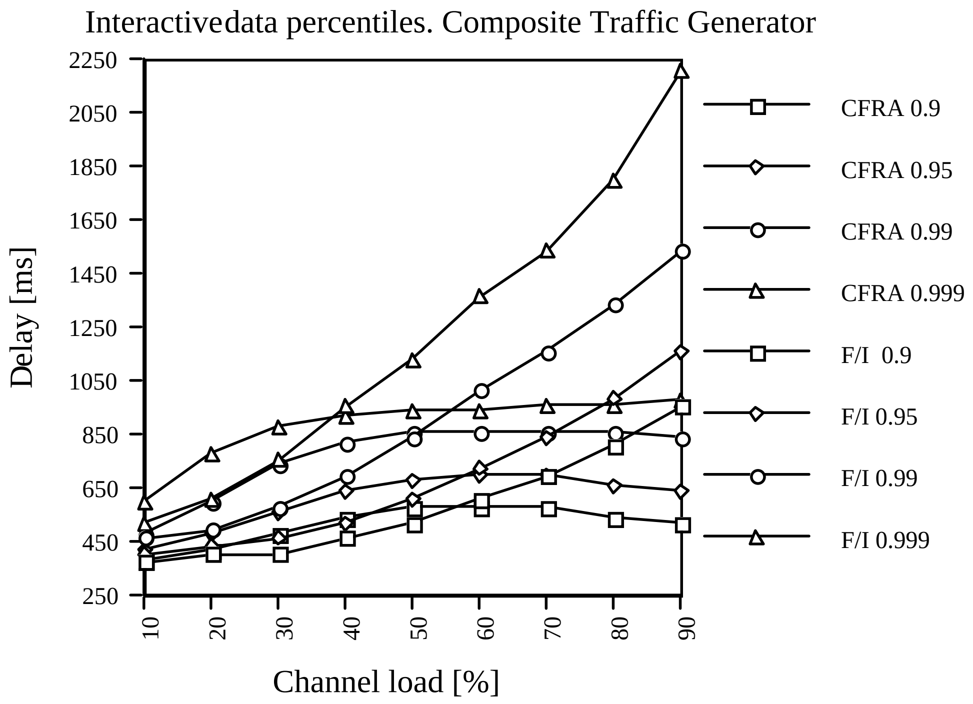

a)

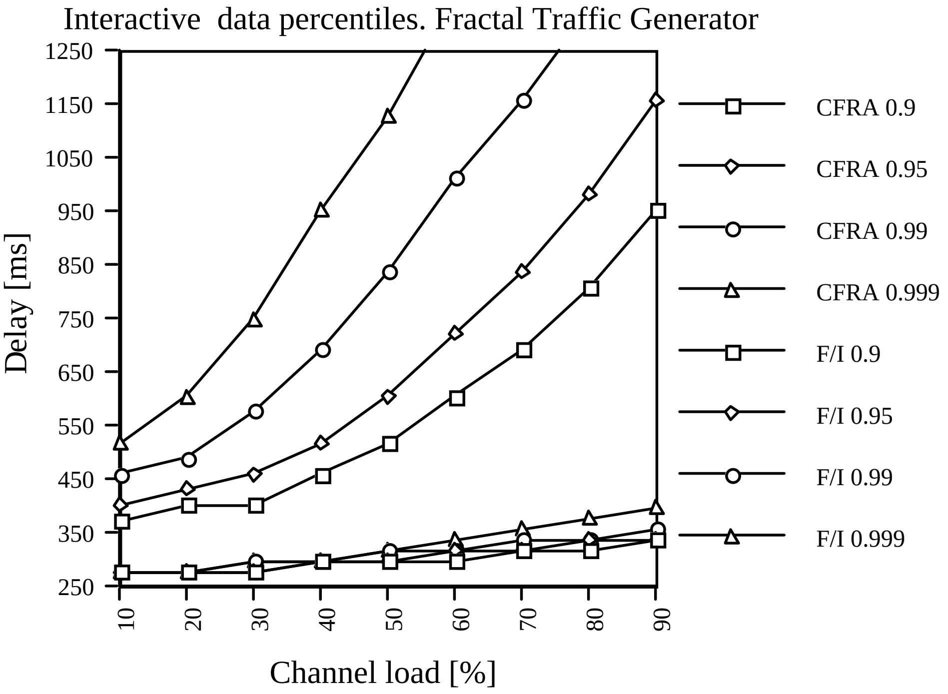

b)

Fig. 12. Cell delay percentiles versus channel load when the interactive traffic is 10% of the channel load (%inter=10%). The load is distributed evenly among all the stations. Interactive data generated by composite traffic generator and fractal traffic generator.

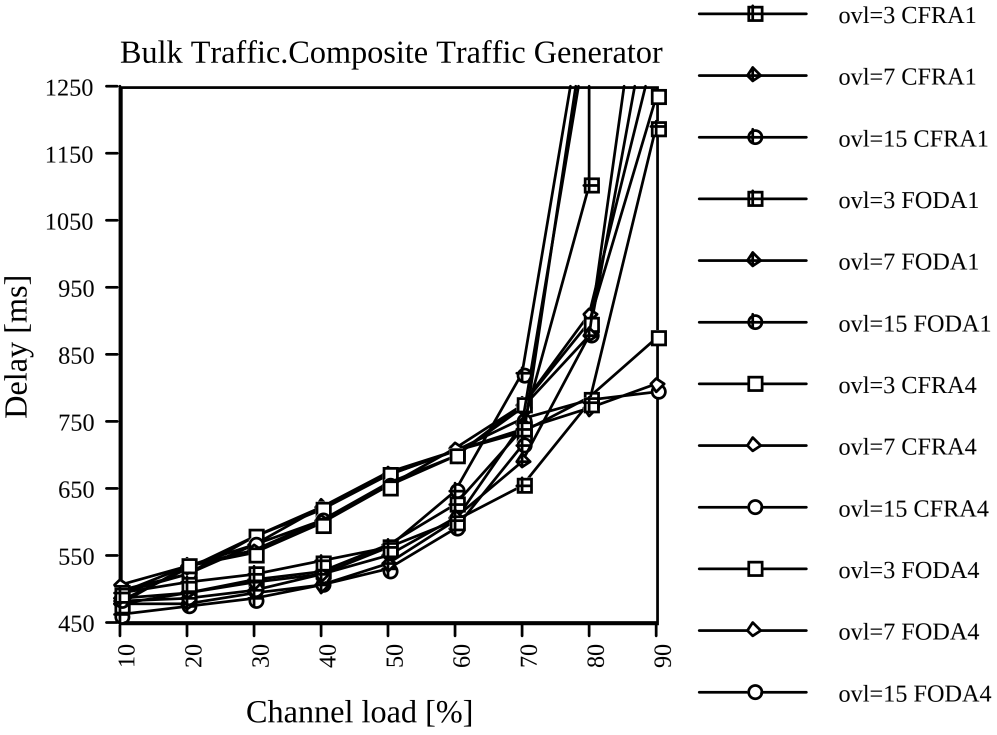
Bulk Traffic.Composite Traffic Generator
Delay [ms]
1250
1150
1050
950
850
750
650
550
450
10
20
30
40
50
60
70
80
90
Channel load [%]
ovl=3 CFRA1
ovl=7 CFRA1
ovl=15 CFRA1
ovl=3 FODA1
ovl=7 FODA1
ovl=15 FODA1
ovl=3 CFRA4
ovl=7 CFRA4
ovl=15 CFRA4
ovl=3 FODA4
ovl=7 FODA4
ovl=15 FODA4

a)

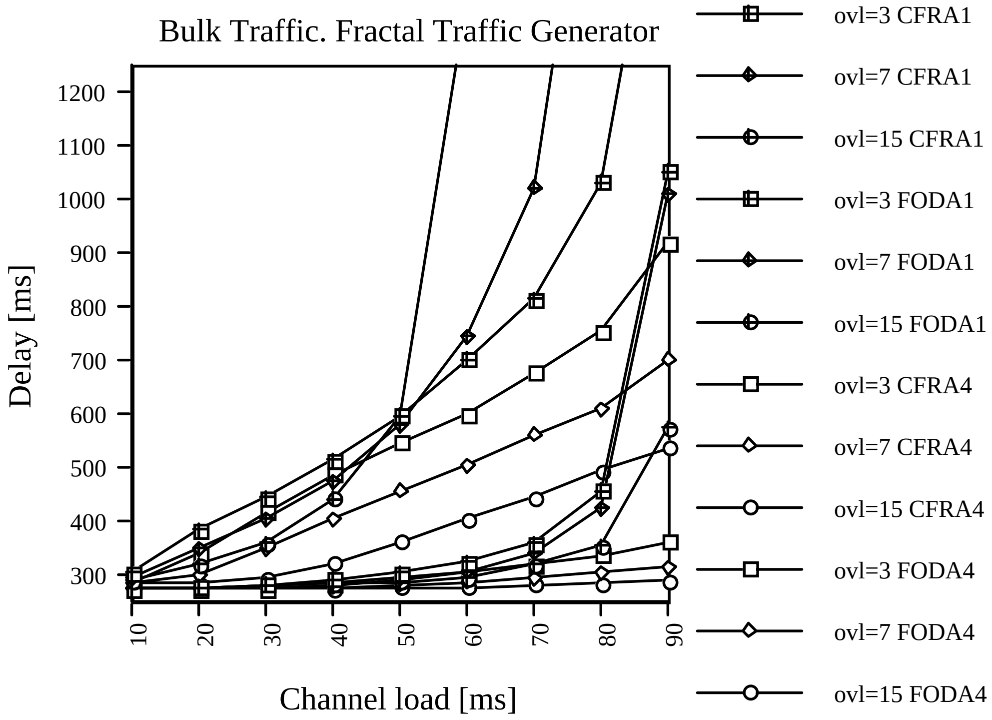
Bulk Traffic. Fractal Traffic Generator
Delay [ms]
1200
1100
1000
900
800
700
600
500
400
300
10
20
30
40
50
60
70
80
90
Channel load [ms]
ovl=3 CFRA1
ovl=7 CFRA1
ovl=15 CFRA1
ovl=3 FODA1
ovl=7 FODA1
ovl=15 FODA1
ovl=3 CFRA4
ovl=7 CFRA4
ovl=15 CFRA4
ovl=3 FODA4
ovl=7 FODA4
ovl=15 FODA4

b)

Fig. 13. Bulk traffic delay versus channel load for bulk overloads 3, 7 and 15 of the first three stations. FODA1 stands for FODA/IBEA stations 1-3; FODA4 stands for stations 4-16. The same for CFRA.

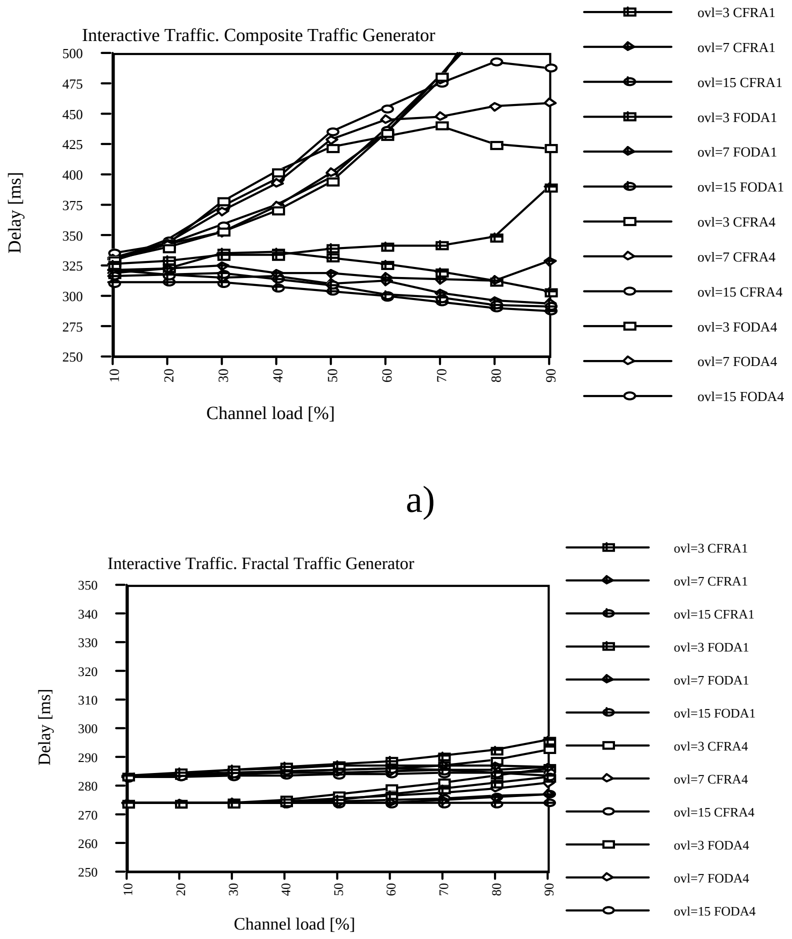

a)

b)

Fig. 14. Interactive traffic delay versus channel load for bulk overloads 3, 7 and 15 of the first three stations. FODA1 stands for FODA/IBEA stations 1-3; FODA4 stands for stations 4-16. The same for CFRA.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>

References

1. N. Celandroni, E. Ferro, N. James, F. Potortì: "FODA/IBEA: a flexible fade countermeasure system in user oriented networks", Inter. Journal of Satellite Communications, 10, No. 6, pp. 309-323, 1992.
2. N. Celandroni, E. Ferro, F. Potortì: "The performance the FODA/IBEA satellite access scheme measured on the Italsat satellite", proceedings of the ICDSC-10 Conference, pp. 332-328, Vol. 1, Brighton 15-19 May 1995, UK.
3. N. Celandroni, E. Ferro, F. Potortì, A. Bellini, F. Pirri: "Practical experiences in interconnecting LANs via satellite", ACM SIGCOMM Computer Communication Review, Vol. 25, No. 5, October 1995.
4. N. Celandroni, E. Ferro, F. Potortì: "FRACAS: FRAmed Channel Access Simulator. User Manual", CNUCE Report C95-27, September 1995.
5. T. Zein, G. Maral, T. Brefort, M. Tondriaux: "Performance of the Combined/Fixed Reservation Assignment (CFRA) scheme for aggregated traffic", Proceedings of the COST-226 Final Symposium, pp.183-198, Budapest (H), 10-12 May 1995.
6. Wing-Cheong Lau, Ashok Erramilli, Jonathan L.Wang, Walter Willinger: “Self-Similar Traffic Generation: The Random Midpoint displacement Algorithm and Its Properties”, ICC ‘95, Seattle, Washington USA, June 18-22 1995, pp. 466-472.
7. Will E. Leland, Murad S. Taqqu, Walter Willinger, Daniel V. Wilson: "On the self-similar nature of Ethernet Traffic (extended version)", IEEE/ACM Transactions on Networking, Vol. 2, No. 1, February 1994.
8. Samples of external Ethernet traffic measured at Bellcore during October 1989 are available with anonymous FTP from flash.bellcore.com, directory: pub/lan_traffic/OctExt.TL.
9. Celandroni N., Ferro E.**:** "The FODA-TDMA satellite access scheme: presentation, study of the system and results", IEEE Transaction on Communication, V. 39, N. 12, pp.1823-1831, December 1991.
10. Vern Paxon, Sally Floyd: "Wide Area Traffic: the failure of Poisson Modelling", IEEE/ACM Transactions on Networking, Vol. 3, No. 3, June 1995.
11. Ikka Norros: “The management of Large Flows of Connectionless Traffic on the Basis of Self-Similar Modelling”, Proceedings of ICC ‘95, Seattle, Washington USA, June 18-22 1995, pp. 451-455.
12. Caceres R., Danzig P., Jamin S., Mitzel D.J.: “Characteristics of wide area TCP/IP conversations”, ACM Computer Communications Review, Vol. 11, No. 4, 1991.

13. Danzig P., Jamin S., Caceres R., Mitzel D.J., Estrin D.: “An empirical workload model for driving wide area TCP/IP network simulations”, Internetworking Research and Experience, Vol. 3, pp. 1-26, 1992.
14. N. Celandroni, E. Ferro, F. Potorti’: “FODA/IBEA-TDMA Satellite access scheme for mixed traffic at variable bit and coding rates", CNUCE REPORT C92-07, December 1992.
15. N. Celandroni, E. Ferro, F. Potorti’: “Outage probability of an adaptative TDMA satellite access scheme", ICC’ 93, Geneva (CH), May 23-26, 1993, pp. 14449-11454, Vol. 3.
16. N. Celandroni, E. Ferro, F. Potorti’: “FODA/IBEA-TDMA: A flexible fade countermeasure system for integrated services in user oriented networks", International Journal on Satellite Communications, Vol. 10, N. 6, pp. 309-323, Nov.-Dec. 1992.
17. Dornier Deutche Aerospace: "Advanced Business Communications via satellite", System Description, July 1992.
18. N. Celandroni, E. Ferro, N. James, F. Potorti’: Design and implementation of a flexible software based TDMA system", proceedings of Globecom' 92, Orlando December 6-9 1992, pp. 693-699, Vol. 2.
19. Jahangir I Mohammed, Tho Le-Noc: “Performance anlysis of Combined Free/Demand Assignment Multiple Access (CFDAMA) Protocol for Packet Satellite", ICC 94, pp. 869-873.
20. N. Celandroni, E. Ferro, F. Potorti’: “A traffic generator for the performance evaluation of a satellite network", Proceedings of 1992 International Conference on Communication Technology, Beijing, China, Sept. 16-18 1992, pp. 9.08.01-9.08.05, Vol. 1.
21. Tho Le-Noc, Jahangir I Mohammed: "Combined Free/Demand Assignment Multiple Access (CFDAMA) Protocol for Packet Satellite Communications", 2nd IEEE International conference on Universal Personal Communications, ICUPC' 93, Ottawa 12-15, 1993.
22. Tho Le-Gnoc, S.V. Krishnamurthy: “Performance of Combined Free/Demand Assignment Multi-Access (CFDAMA) Protocol with Pre-Assigned Request Slots in Integrated Voice/Data Satellite Communications", Proceedings of the International Conference on Communications, ICC' 95, pp. 1572-1576, 1995.
23. Tho Le-Noc and V. Krishnamurthy: "Performance of Combined FREE/DEMAND Assignment Multiple-Access Schemes in satellite communications", Internationnal Journal of Satellite Communications, Vol. 14, 11-21 (1996).

24. Parag Pruthi, Ashok Erramilli: “Heavy-tailed ON/OFF source behaviour and self-similar traffic”, Proceedings of ICC’95, Seattle, pp. 445-450, 1995.
25. Julius Goldhirsh: “ Space diversity performance prediction for earth-satellite paths using radar modelling techniques”, Radio Science, Vol. 17, No. 6, Nov. Dec. 1982, pp. 1400-1410
26. F. Carassa, G. Tartara, E. Matricciani: “Frequency diversity and its applications”, International Journal on Satellite Communications, Vol. 6, July-September 1988, pp. 313-322.
27. Antony S. Acampora: “ The use of the resource sharing and coding to increase the capacity of digital satellites”, IEEE On Selected Areas In Communications, Vol. JSAC-1, No. 1, January 1983.
28. N. Celandroni, E. Ferro, F. Potorti' : "Experimental results of a demand-assignment thin route TDMA system", International Journal on Satellite Communications, Vol. 14, No. 2, pp. 113-126, March-April 1996.
29. N. Celandroni, E. Ferro, F. Potorti' : "Comparison between distributed and centralised demand assignment TDMA satellite access schemes", International Journal on Satellite Communications, Vol. 14, No. 2, pp. 95-112, March-April 1996
30. N. Celandroni: “Analysis of transient of the datagram traffic for a demand assignment TDMA satellite access scheme”, to appear on Performance Evaluation.
31. N. Celandroni, F. Potorti' , S. Rizzo: "An inexpensive rain fade countermeasure technique for DA-TDMA satellite systems”, GLOBECOM ‘96, 18-22 Nov. 1996, London, pp.1000-1004.

Copyright 1997 by John Wiley & Sons, Ltd. <URL: http://www3.interscience.wiley.com/>