# DEMAND ASSIGNMENT TDMA SATELLITE SCHEMES: DISTRIBUTED AND CENTRALISED SOLUTIONS

Nedo Celandroni, Erina Ferro, Francesco Potorti'
CNUCE, Institute of National Research Council
Via S. Maria 36 - 56126 Pisa - Italy
Fax: +39-50-904052 / 589354
Phone: +39-50-593111
E-mail: n.celandroni@cnuce.cnr.it
e.ferro@cnuce.cnr.it
pot@cnuce.cnr.it

**ABSTRACT:** This paper presents the performance measurements from a comparison between FODA/IBEA and two other TDMA satellite access schemes: FEEDERS and DRIFS. The three schemes differ in terms of the scheduling of the channel capacity: i.e. centralised control in FODA/IBEA and distributed in the other two. All these access schemes have been designed at CNUCE, where the simulation tool used for the comparison was also developed. The two distributed control access schemes have been studied and simulated to analyse the behaviour of the capacity assignment algorithm, while the recovery procedures needed for the network stability are presented in [5]. All the access schemes support an *aggregated* traffic, and the quality of the data transmission is guaranteed even when the transmitting signal experiences a severe attenuation due to bad atmospheric conditions.

## 1. INTRODUCTION

FODA/IBEA (Fifo Ordered Demand Assignment/Information Bit Energy Adapter) is a centralised demand assignment satellite access scheme, operating in TDMA and developed on a prototype of hardware manufactured by Marconi R.C. (U.K.). It was tested on several different satellites (Olympus, Italsat, Eutelsat) in order to evaluate its performance [1, 2]. The goal of our study was to transform FODA/IBEA from a centralised into a distributed control scheme, to save one round trip time (RTT) between the times of the request and the assignment of the capacity. On the other hand, note that although centralised control is expensive in terms of end-to-end delay, it is nevertheless more robust than a distributed one, and the simplicity of its implementation and robustness may compensate for the longer allocation delay. DRIFS (Distributed allocation with Request In Fixed Slots) and FEEDERS (Faded Environments Effective Distributed Engineering Redundant Signalling) derive from FODA/IBEA, but they have distributed control. The main characteristics of FODA/IBEA are maintained, i.e. the ability to support aggregated traffic (real-time or *stream* traffic, and non real-time or *bursty* traffic) and the capability to maintain the quality of service requested by the applications in any weather conditions, even when the transmitting signal is faded. The fade countermeasure technique adopted adapts the energy per information bit to each individual link status, which depends on atmospheric conditions. The total attenuation of each link (up-link plus down-link) is compensated by varying the transmitting power, which is used to compensate the whole or part of the up-link attenuation. On the other hand, the total compensation is completed by varying the data coding and bit rates, according to the level of the fade detected both in the sending and in the destination stations. The redundancy is obtained by increasing the data coding rate first[1], then reducing the data bit rate[2] as well, if the data coding redundancy is insufficient to counter the fade. This countermeasure technique was developed in the FODA/IBEA system by exploiting a very powerful feature of the modem prototype, which is able to dynamically change the data bit rates on a sub-burst basis[3], according to the fade levels of the destination and source stations. The price of this feature is the need for a short preamble between sub-bursts, besides the rather long preamble needed in front of each burst to allow carrier and bit timing synchronisation. All these preambles considerably reduce system efficiency. The three access schemes are in principle independent of the hardware used. In any case, we will refer to a modem which uses preambles for acquisition (as used in FODA/IBEA), though all the three access schemes proposed would obviously benefit from a more efficient preambleless modem.

Section 2 presents a short overview of the three access schemes, addressing the reader to the literature where each access scheme is detailed. In Section 3 some comparison results are presented and discussed, while the complete results can be found in [8]. Finally, our conclusions are reported in Section 4.

## 2. SHORT ACCESS SCHEMES OVERVIEW

### 2.1 FODA/IBEA

Some concepts are hereafter reported in relation to their use in FODA/IBEA, but where necessary we will explain how DRIFS and FEEDERS make use of these concepts too.

- The *frame* is the interval of time between two consecutive reference bursts sent by the master station to allow system synchronisation. In each frame a station may transmit a data burst if an allocation is assigned (Fig. 1). The frame length is fixed to 20 ms. The reference burst contains the burst *time plan* (BTP) for the assignments of the transmission times (*transmission windows*) of the stations. The BTP must be known by each station, so as to transmit and to receive all the incoming bursts as well, since the incoming bursts bit rates must also be known in advance by the modem.
- Requests for fixed-bit-rate (FBR) or for variable-bit-rate (VBR) stream data are sent only at the opening of the stream sessions and, if accepted, the relevant stream allocation is kept until it is explicitly relinquished or the requesting station is declared dead.

(1) Available values are: 1/2, 2/3, 4/5 and uncoded.

(2) Available values are: 1, 2, 4 and 8 Mbit/s.

(3) The burst is the amount of data transmitted by a station in its transmission window. It may be formed by several sub-bursts, each one addressed to a different station.

0-7803-3250-4/96$5.00©1996 IEEE

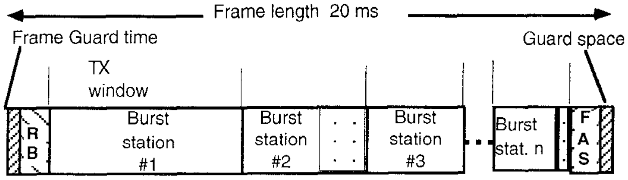


Fig. 1. Frame format in FODA/IBEA

• Bursty requests are sent as frequently as possible, computed as $r = q + H\,i$, where $q$ is the station backlog, i.e. the volume of data already waiting for transmission to the satellite, $i$ is the amount of traffic coming into the station, and $H$ is a temporal constant. The requests are sent mostly piggy-backed with the data or by using the *control slot*, a portion of the frame assigned, on a round robin basis, to the stations which have no assignment for that frame. In FODA/IBEA there is one control slot for every group of eight stations. The master organises the bursty requests of all the slaves into a *ring,* which it scans cyclically to compute the assignments. The length of the bursty assignment $(a)$ is proportional to the request in a range of values between a minimum ($T_{min}$) and a maximum ($T_{max}$) thresholds: $T_{min} \leq a = f\,r \leq T_{max}$, where $f$ is the coefficient of proportionality in the assignment, set equal to the number of active stations N divided by 100, with 5% as minimum and 50% maximum. $T_{min}$ was introduced to avoid too small allocations when the transmission overheads are too big with respect to the information data. $T_{max}$ prevents an overloaded station from removing too much capacity from the other stations. After each bursty assignment, the relevant request in the ring is decreased by the assignment itself and the next request is analysed, if space is still available in the frame. The ring is not scanned more than once in a frame. We call *assignment cycle* a complete scan of the ring. After an assignment cycle, any unused space in the frame is redistributed among all the stations [2].

• In unfaded conditions, the stream traffic cannot go over a fixed limit in the frame, leaving the remaining portion of the frame to the bursty traffic, which can temporarily expand in the frame if the stream traffic is insufficient to reach the boundary. The stream traffic can go beyond this boundary only in faded conditions, due to the redundancy of the data. In this situation the bursty traffic may momentarily be suspended and the whole frame devoted to the stream traffic. No new stream requests are accepted during the boundary overflow time. Fading conditions also cause an increase in the backlog and in the instantaneous traffic of the faded station as well as of the stations transmitting to it, and this automatically increases the bursty request. The possible squeezing of the bursty capacity generally requires a congestion control scheme, which FODA/IBEA implements by blocking the growth of the backlog for a while when the internal queues length is such that the queuing time estimated goes beyond a threshold.

• A new station is given the opportunity to enter the satellite network by using the FAS (First Access Slot) space in the frame. FAS has a fixed position before the end of the frame, and its frequency is every 32 frames. Since it is accessed in contention mode, if there is a collision with another station entering the system, the colliding stations wait for a random number of frames before repeating the operation. When the maximum number of active stations is reached, the FAS space is temporarily deallocated; it is allocated again when at least one more entry is possible in the system.

### 2.2. DRIFS

Each frame consists of a *control sub-frame* (CSF), used in Fixed-TDMA, followed by a *data sub-frame* (DSF), used in Demand Assignment-TDMA, and finally by an FAS. No reference burst is supported.

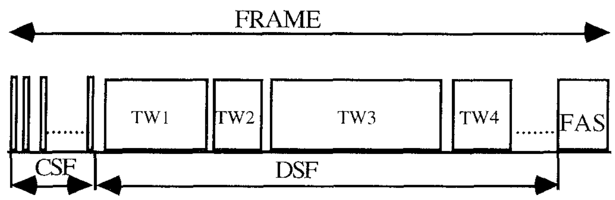


Fig. 2. Frame structure in the DRIFS scheme.
TW = Transmit Window

The control sub-frame is used for information such as requests and fade levels, while the application data is sent in the data sub-frame. Like in FODA/IBEA, the FAS is used by those stations willing to enter the system, but its frequency is every frame, if the maximum number of stations has not been reached. Each active station has its own control slot in the CSF, whose size is thus proportional to the number of active stations. The position of the control slot of a station is fixed. It is only shifted back when a preceding station leaves the system. The first control slot is preceded by a special unique word, to allow system frame synchronisation. The bursts sent in the control slots following the first one are preceded by the control unique word, to distinguish them from the data bursts. Each active station that has an allocation for transmitting data sends one burst in the data sub-frame. Since the CSF length is not negligible with respect to the frame length (at least using traditional modems which 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, 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 $x$. The control data must be strongly protected against either loss of acquisition or transmission errors. Assuming we use the same modem that we have at present, we devise a bit rate of 2 Mbit/s and a coding rate of 4/5. The request and allocation algorithms both for stream and for bursty data are the same as in FODA/IBEA. The difference is in the assignment cycle, which in FODA/IBEA can be any length, while in DRIFS it is always equal to $C_c$ frames. Part I of [5] details the DRIFS scheme, together with the estimation of wrong reception and non transmission probabilities.

### 2.3 FEEDERS

In FEEDERS the concept of a "master" station still exists, but only as a reference in case of errors. The frame is the interval of time in which all the active stations transmit a data burst which contains control information which may or may not be followed by application data. The reference burst is sent every $n_a$ frames (Fig. 3) with a special unique word, allowing all the traffic stations to synchronise with the network. Inside the reference burst, a reference BTP is also sent to allow any new station that wants 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.

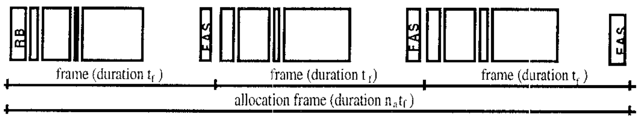


Fig. 3. Frame structure in the FEEDERS scheme

The *allocation frame* (AF) is the time between two consecutive BTP applications. The AF coincides with the time between two consecutive reference bursts sent by the reference station. 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 and $t_f$ is the time frame length. 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. This redundancy dramatically reduces the probability that none of the stations understand the control information. The BTP is applied to all the frames in an AF.

The allocation request and assignment for stream capacity is computed similarly as in FODA/IBEA. The allocation request for bursty traffic is computed by each station as $r = \min[f\ (q + H\ i),\ r_{max}]$, where $f$ (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 the allocations are computed in a distributed way. They can vary on an AF 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 the stations.

The bursty allocation is such that for each station an allocation equal to the request, increased by the preamble overhead, is preliminary allocated in every 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 ones which sent a null request. If $C_r$ is negative, all the requests are reduced by a suitable factor, to make the allocations occupy exactly the capacity $C_d$ reserved for datagram. In this case the stations which sent a null request receive an allocation which is only enough to send one control message per frame. No control slot is used. Part II of [5] details the FEEDERS scheme, together with the estimation of the wrong reception and non transmission probabilities.

## 3. SOME COMPARISON RESULTS

Our aim is to compare the three methods in terms of bursty data capacity assignment, since the stream capacity allocation does not present particular problems and the algorithm is the same for all three methods. Many simulation runs we thus made to compare the three methods with Poisson, Two-states Markov-modulated Poisson and Fractional Gaussian Noise traffic generators. Although Poisson traffic is not good for modelling the Ethernet LAN traffic [6], it does have attractive theoretical properties, and for wide area traffic it is still valid for modelling the arrival of user sessions [7]. Moreover, Poisson traffic was used to load the FODA/IBEA system for measurement tests with four earth stations on the Italsat satellite; therefore the comparison between the simulation results and the measurements on satellite needs the Poisson load of the channel. Poisson can be considered as a case of particularly well-behaved traffic, from a burstiness point of view. In addition, it does not exhibit long range dependence, i.e. its autocorrelation function decreases exponentially. The Two-state Markov-modulated Poisson traffic we used still does not exhibit long range dependence, but does exhibit high burstiness. On average, the high traffic state lasts for 0.5 seconds and the low traffic state for 2.5 seconds. The mean traffic generation rates are in 17:1. The parameters of this generator, which are equal for all the stations, were chosen in order to represent a worst-case load for satellite protocols from a burstiness point of view, yet to prevent any single station from generating a peak traffic greater than half the capacity of the channel.

The fractal generator is an approximated Fractional Gaussian Noise generator implemented with a simple Random Midpoint Displacement algorithm. This generator exhibits relatively low burstiness and a long-term correlation, which we have truncated to about 10 minutes simulation time. We used a Hurst parameter equal to 0.85, following the findings published in [6]. The peakedness, defined as the ratio between the variance and the traffic distribution mean values, has been set equal for all stations. Its value is 1/10 of the total channel mean load of each simulation run; this proportionality aims at obtaining a constant form factor (standard deviation / mean value) for the total load in all the simulation runs.

Figure 4 shows the reliability of the simulation tool used. The mean end-to-end delay, averaged over 30 s, is shown as a function of each of the four stations' load. Each quasi-horizontal line represents a run characterised by an overall channel load level, while each quasi-vertical line represents the delay of one of the four stations. The results obtained by simulating FODA/IBEA sufficiently agree with the ones obtained with the real tests on the Italsat satellite. In this case, the simulation tool was adjusted to introduce a minimum delay of 292 ms, considering the pure Round Trip Time (252 ms) plus two frames. This is because an extra frame is introduced by the present system implementation, due to insufficient CPU power in the satellite controllers. This additional frame could easily be saved by using a more powerful CPU board, which is now on the market. In the legends of the following figures -P stands for Poisson traffic, -I for Impulsive or Two-states Markov-modulated Poisson, and -F for Fractal traffic. All the tests were averaged over 30,000 frames and refer to the case in which the whole channel capacity is devoted to the bursty data. Figures 5 and 6 report the mean end-to-end delay of the overall channel as a function of the traffic load for the three traffic types, and for 12 and 32 stations, respectively. In [8] results for 4 and 48 stations are also presented, and show that for a low number of stations FEEDERS performs much better than the other schemes, principally due to its lower overhead. The distributed methods generally perform better than the centralised one as the traffic burstiness increases for all the loading levels. With 32 stations FEEDERS is penalised by the higher number of preambles per frame, which considerably reduce the channel efficiency. This method performs even better than FODA/IBEA for impulsive traffic, but is the worst in the other two traffic cases. DRIFS is the best with 32 stations for high loads, while it is penalised by having a fixed allocation cycle (4 frames) at low-medium loads. The 12 station case was investigated in more detail. Figures 7÷10 refer to this case. Figure 7 shows the 95, 99 and 99.8 percentiles of the channel delay versus the channel load for the three schemes and for the Fractal traffic model. In [8] results for the other two traffic generators can be found. Figures 8÷10 show the mean end-to-end delay of each station as a function of the station load for all the three schemes, using Fractal traffic. Again, results can be found in [8] when using Poisson and Impulsive traffics, and the following considerations derive from the complete set of results. Stations 7÷9 are equally loaded, so they are represented by only one line. The same is for stations 10÷12. The original goal was to make all the stations experience the same delay when the channel is not saturated, independently of the loading condition of each station. This aim has roughly been reached with Poisson and Fractal traffics by all three schemes, though the distributed schemes perform better than the centralised one. When the channel approaches saturation, in all cases the most loaded stations are penalised, due to the maximum allocation limit present in all the schemes. When the system is loaded with Impulsive traffic, the delay tends to increase almost linearly with the load of the stations in all three schemes. Note, however, that the Impulsive traffic pattern used represents the most severe loading condition. The best performance of FEEDERS is evident, as already seen by looking at the channel delay characteristics.

## 4. CONCLUSIONS

The simulation results highlight that the FEEDERS access scheme works much better than FODA/IBEA and DRIFS with any type of traffic if the number of stations is not very high (less than 32). From 32 on, DRIFS and FODA/IBEA perform better, since FEEDERS cannot accommodate too many stations in one frame. The optimal distributed algorithm is thus a compromise between FEEDERS and DRIFS. This scheme should work like FEEDERS until the number of stations is less than 32, then, according to the increasing number of stations, it should enlarge the allocation frame to $C_c$ frames, as designed for DRIFS. This is an interesting result, and the performance of such a scheme could be investigated in a future work. Another interesting result is that FODA/IBEA is competitive with DRIFS, as pointed out also by the real experiment of LAN interconnection (*Thin route TDMA for LAN interconnection*) where it was used accessing the Italsat satellite.

## REFERENCES


1. N. Celandroni, E. Ferro, N. James, F. Potorti', "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. Potorti', "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. Potorti', 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, M. Conti, E. Ferro, E. Gregori, F. Potorti',"A bandwidth assignment algorithm on a satellite channel for VBR traffic", to appear on the International Journal on Satellite Communication.
5. N. Celandroni, E. Ferro, F. Potorti', "Study of distributed algorithms for satellite channel capacity assignment in a mixed traffic and faded environment", CNUCE Report C95-28, September 1995.
6. W. E. Leland, M. S. Taqqu, W. Willinger, D. V. Wilson, "On the self-similar nature of Ethernet Traffic (extended version)", IEEE/ACM Transactions on Networking, Vol. 2, No. 1, February 1994.
7. V. Paxon, S. Floyd, "Wide Area Traffic: the failure of Poisson Modeling", IEEE/ACM Transactions on Networking, Vol. 3, No. 3, June 1995.
8. N. Celandroni, E. Ferro, F. Potorti', "Comparison between distributed and centralised demand assignment TDMA satellite access schemes", submitted to the International Journal on Satellite Communications, September 1995.


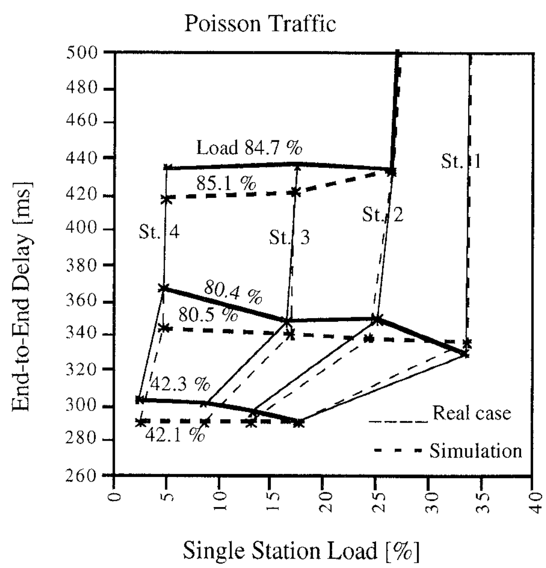


Fig. 4. Comparison between the FODA/IBEA real and simulated cases. Mean end-to-end delay versus single station load. 4 stations. Poisson Traffic. 30s run

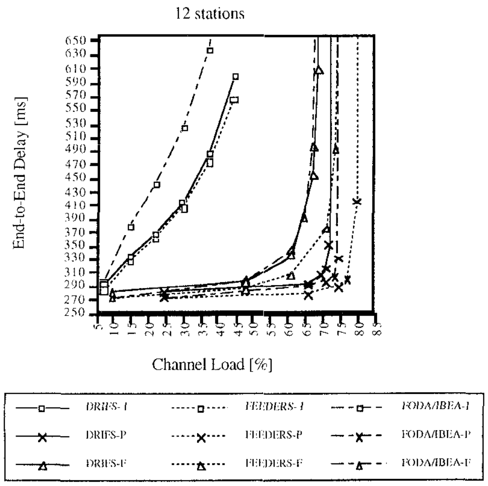


Fig. 5. Comparison between FODA/IBEA, DRIFS and FEEDERS for Fractal, Poisson and Impulsive traffic. Mean end-to-end delay versus single station load. 12 stations.

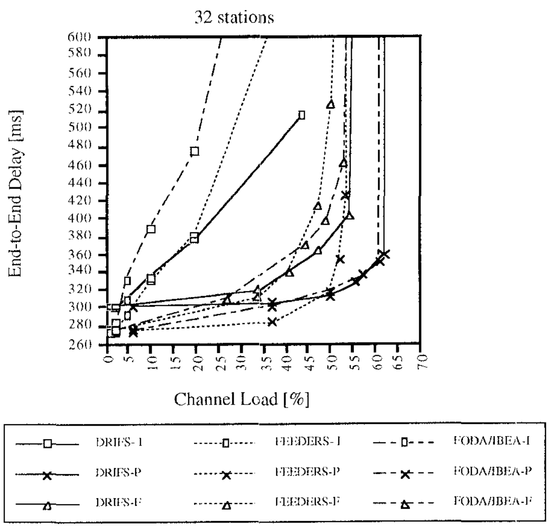


Fig. 6. Comparison between FODA/IBEA, DRIFS and FEEDERS for Fractal, Poisson and Impulsive traffic. Mean end-to-end delay versus single station load. 32 stations.

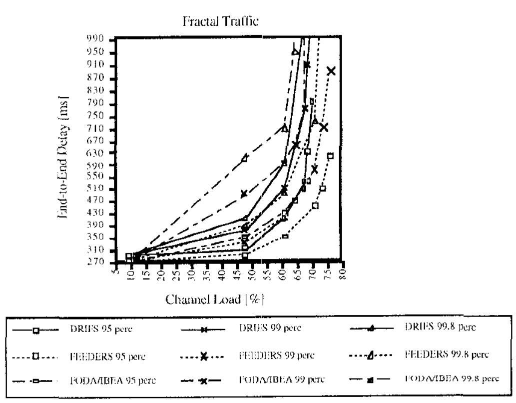


Fig. 7. 12 stations. Fractal Traffic. Percentiles of the channel end-to-end delay versus channel load

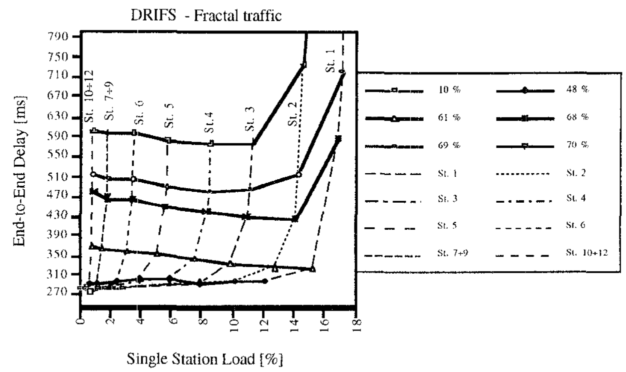


Fig. 8. 12 Stations. Fractal Traffic. DRIFS scheme. Mean end-to-end delay versus single station load

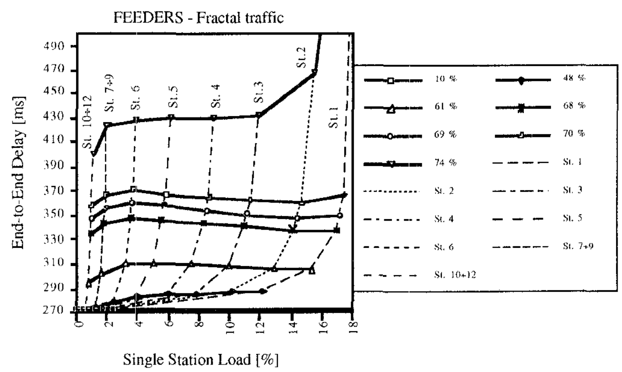


Fig. 9. 12 Stations. Fractal Traffic. FEEDERS scheme. Mean end-to-end delay versus single station load

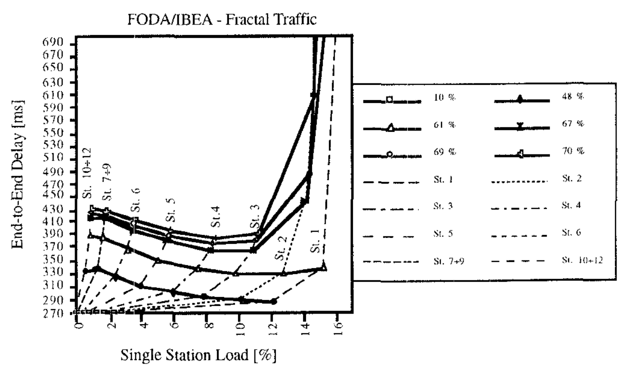


Fig. 10. 12 Stations. Fractal Traffic. FODA/IBEA scheme. Mean end-to-end delay versus single station load