

Computer Communication Review

acm  sigcomm

Special Interest
Group on
Data
Communication

The
SIGCOMM
Quarterly
Publication

Volume 25
Number 5
October, 1995

Problems with DCE Security Services; G. White and U. Pooch	5
An Architectural Approach for Integrated Network and Systems Management; R. Boutaba and S. Znaty	13
The (Un)Revised OSI Reference Model; J. Day	39
Practical Experiences in Interconnecting LANs Via Satellite; N. Celandroni, E. Ferro, F. Portorti, A. Bellini, and F. Pirri	56
Performance Problems in BSD4.4 TCP; L. S. Brakmo and L. L. Peterson	69



COMPUTER COMMUNICATION REVIEW

A Quarterly Publication of ACM SIGCOMM

Volume 25, Number 5
ISSN #: 0146-4833

October, 1995

Contents

Chairman's Message 4

Papers

Problems with DCE Security Services; G. White and U. Pooch 5
An Architectural Approach for Integrated Network and Systems Management; R. Boutaba and S. Znaty 13
The (Un)Revised OSI Reference Model; J. Day 39
Practical Experiences in Interconnecting LANs Via Satellite; N. Celandroni, E. Ferro, F. Potorti, A. Bellini, and F. Pirri 56
Performance Problems in BSD4.4 TCP; L. S. Brakmo and L.L. Peterson 69

Features

Bibliography of Recent Publications in Computer Communication 87
SIGCOMM Calendar of Events 97

Calls for Papers, Conferences, Journals, and other Announcements

SIGCOMM Award Nominations 106
ACM and SIGCOMM Membership Application Form 107

PRACTICAL EXPERIENCES IN INTERCONNECTING LANs VIA SATELLITE

Nedo Celandroni*, Erina Ferro*, Francesco Potorti*
Alessandro Bellini[†], Franco Pirri[†]

*CNUCE, Institute of National Research Council
Via S. Maria 36 - 56126 Pisa - Italy

[†]Depart. of Electr. Engineering, University of Florence
Via S. Marta 3, 50100 Firenze - Italy

ABSTRACT

We present an experiment in interconnecting LANs via a satellite link and describe the individual components involved in the experiment. The project was developed in two phases: a) design and realisation of a satellite access scheme that supports real-time and non real-time traffic with a signal fading countermeasure, called FODA/IBEA-TDMA¹; b) interconnection of LANs where real-time and non real-time applications run. The experiment was presented the first time in June 1994 as a demo in which the Eutelsat satellite was used (in the 12/14 GHz band) to exchange data between two LANs in Pisa and Florence, while video and audio applications running on PCs connected the two sites. The demo was repeated a few weeks later and the Intelsat satellite was used in the 20/30 GHz band.

KEYWORDS: satellite, TDMA fade countermeasure, satellite LAN interconnection, satellite videoconference

1. INTRODUCTION

Much work is in progress in the field of communication systems for real-time multimedia applications over local area networks (LANs). The goal of our research was to design and develop a heterogeneous network capable of simultaneously handling video, voice and data. The network is heterogeneous in many senses: it is composed of various types of transmission media (LANs and satellite links); it carries different types of traffic

(video, voice and computer generated data); and it uses several protocols, such as TCP/IP over Ethernet for the connection to external LANs, an internal Token Ring LAN used for data, audio and video applications, FODA/IBEA for the satellite network, and GAFO² for the communications between the terrestrial gateway (where the data coming from the various LANs converge) and the satellite network.

Satellite links are more problematic than terrestrial links for two main reasons. The satellite round trip delay is quite significant when transmitting real-time data (stream data). Moreover, the HF signal may be considerably attenuated by bad atmospheric conditions, with a consequent degradation in the quality of the data transmitted. The latter is currently one of the major problems for communication satellites operating at frequencies above 10 GHz. As the K_U band (12/14 GHz) is shortly going to be saturated, we considered the K_a band (20/30 GHz) for the satellite network design. FODA/IBEA [11], the satellite access scheme we have developed, operates in TDMA. It allows the simultaneous transmission of real-time and non real-time (datagram) data, while maintaining the BER³ required by applications, even in conditions of signal fading, by varying the information bit energy (fade countermeasure technique). At the same time, the University of Florence developed a gateway which takes data coming from both Ethernet and Token Ring LANs and sends them to the satellite network using GAFO, a protocol developed to communicate with the satellite earth stations.

¹ FIFO Ordered Demand Assignment / Information Bit Energy Adaptive - Time Division Multiple Access [11]

² Gateway-Foda protocol, a protocol developed ad-hoc for the communications with the FODA/IBEA stations.

³ Bit Error Rate.

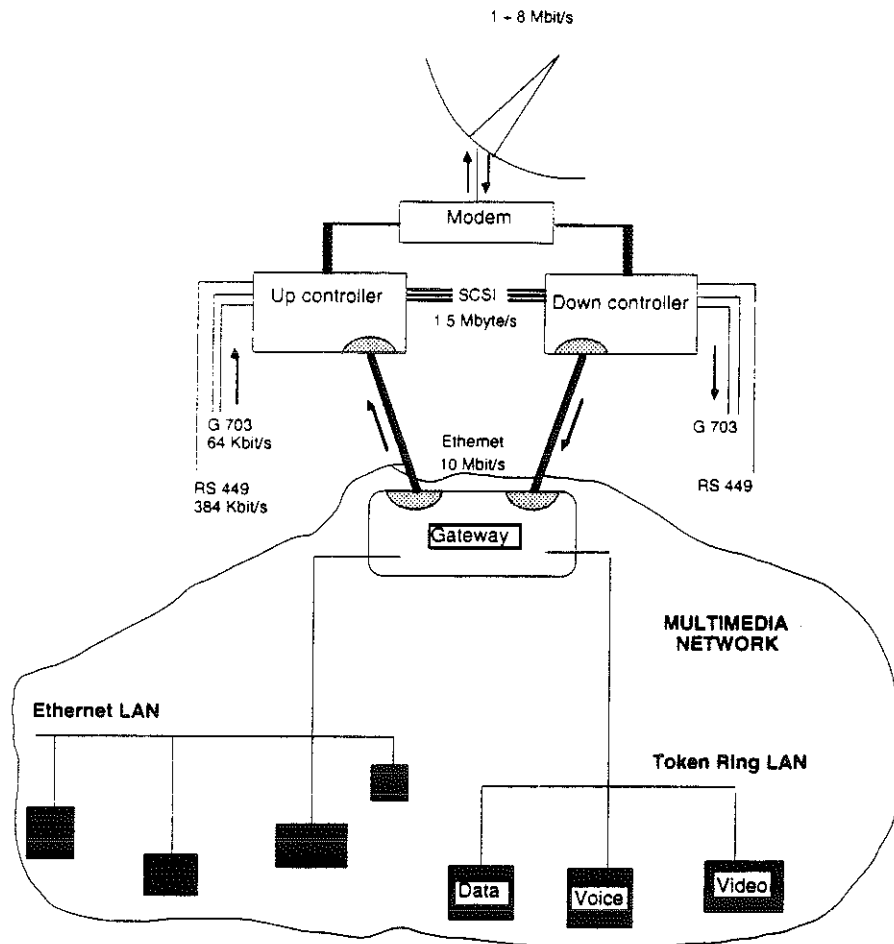


Fig. 1. The network scenario

The paper is organised as follows. In Section 2 the network scenario is presented, while in Section 3 the earth station hardware is briefly described. Section 4 describes the satellite access scheme, and Section 5 highlights the features of the GAFO protocol. Section 6 describes the LAN interconnection and outlines the demonstration of the whole interconnected network.

2. THE NETWORK SCENARIO

Figure 1 depicts the network scenario used in the demo of the project. The three earth stations had the same configuration; two of them were located in Pisa, while the third one was located in Florence (about 100 Km away). Although for practical reasons the demo only involved three stations, this configuration can easily be extended to a larger number of stations, up to the hardware limit of 112. Each station needs a gateway that connects it to an Ethernet LAN and to stream applications (data, voice, video), or else the

stations can themselves be Ethernet LAN hosts, provided that the applications that use the satellite link are equipped with the GAFO protocol interface.

In the configuration used in the experiment, the stream applications run on hosts connected to a Token Ring network, while the datagram applications run on hosts connected via Ethernet. Data from both the LANs converges at the gateway which communicates with the FODA/IBEA satellite access system by using the GAFO protocol. The set of applications running on Ethernet plus those running on the Token Ring network plus the gateway are collectively called a Multimedia Network. This part was developed by the University of Florence.

FODA/IBEA runs on hardware designed by CNUCE, in collaboration with Telespazio, and developed by Marconi R.C. (U.K.). It consists of a TDMA controller (which includes a variable coding rate codec), and a variable bit rate burst modem. The TDMA controller is split into two parts: the Up

Traffic type	Throughput	Packet Length	Tolerated Delay	BER Required	FODA/IBEA COS
interactive	low	small, variable	low	$< 3 \cdot 10^{-7}$	2
control	low	small, variable	medium	$< 10^{-8}$	1
mail	low/medium	medium, variable	high	$< 3 \cdot 10^{-7}$	2
bulk	high	medium/high, constant	medium/high	$< 10^{-8}$	1
fax	medium	big, constant	medium	$< 3 \cdot 10^{-5}$	3
VOICE	medium	small, constant	low, constant	$< 3 \cdot 10^{-7}$	2
special				$< 3 \cdot 10^{-5}$	3
standard				$< 10^{-3}$	4
degraded				$< 3 \cdot 10^{-5}$	3
VIDEO	medium	big, constant	low, constant	$< 3 \cdot 10^{-5}$	3
slow scan TV					
video-conference	high	big, constant	low, constant	$< 10^{-8}$	1

Table 1. Traffic characteristics

controller, for transmissions towards the satellite (up-link), and the Down controller, to receive data from the satellite (down-link). The communication between the Down and the Up controller is performed via an SCSI link at 1.5 Mbytes/s. Dedicated Ethernet lines connect the gateway to the Up controller and the Down controller, respectively.

Stream traffic, such as voice, slow-scan TV, videoconference, etc., is characterised by a constant packet arrival rate. Stream applications usually require short and fairly constant delays at the receiving end, cannot tolerate out-of-order delivery of packets, but can tolerate occasional bit errors and dropped packets. In practice, stream traffic needs a fixed amount of bandwidth on a regular basis and the satellite network should not alter the inter-arrival time of the packets.

Datagram traffic is sub-divided into *bulk* and *interactive* types. Bulk traffic applications (typically file transfer protocols) usually transmit a large number of packets; however the delay caused by the network crossing is not critical, so this type of traffic does not have the rigid delay requirements of stream traffic. The satellite network can deliver the packets out of order, since higher level protocols take care of reordering them. However, corrupted or lost packets may severely impair the end-to-end throughput of such traffic, especially on a high delay medium like the satellite network.

Interactive traffic (terminal access to computers, database enquiries, operator message exchanges, etc.) demands error free, reliable delivery of packets and short end-to-end delay to guarantee acceptable

response times. It often consists of short packets (64 bytes on Ethernet is typical) with very irregular inter-arrival times.

Table 1 outlines different types of traffic in terms of the key parameters of delay, throughput and BER range which are commonly required. The last column in the table (FODA/IBEA Class Of Service (COS)) represents the mapping of the BER required by the associated type of traffic onto the satellite network. Four classes of service are envisaged, and COS number 1 is associated with the maximum protection from errors of the data transmitted. The COS number must be specified when using the GAFO protocol to pass data from the terrestrial to the satellite network. For stream data, the COS value is specified only once, when the virtual channel is set up; for datagram data the relevant COS value is associated with each packet.

Table Legend:

Throughput (T):
 low ($T < 1$ Kbit/s)
 medium (1 Kbit/s $< T < 100$ Kbit/s)
 high ($T > 100$ Kbit/s)

Packet length (L):
 small ($L < 128$ bytes)
 medium (128 bytes $< L < 1$ Kbit)
 big ($L > 1$ Kbyte)

Delay (D):
 low ($D < 0.5$ s)
 medium (0.5 s $< D < 2$ s)
 high ($D > 2$ s)

3. THE EARTH STATION HARDWARE

Figure 2 shows the hardware of the TDMA controller and the variable bit rate burst modem.

As some boards contained in the Down controller are the same as in the Up controller, some explanations are given for the Up controller alone.

The Up controller consists of the following boards:

- Standard industrial board with a 25 MHz 68030 processor, 4 Mbyte DRAM memory, bidirectional SCSI-I interface, VMEbus interface, a LANCE chip for Ethernet, 4 RS-232 ports (only one is actually used), timers, real-time clock, and a small non volatile RAM. It runs VMExec, a PSOS based real time operating system. The FODA/IBEA software—which controls and monitors the functions of the transmit control unit—uses the services offered by the real-time kernel, namely the priority preemptive task scheduler and inter process communication manager. Events, message queues, semaphores, task suspension and resumption are supported.
- Transmit serial interface This is a wire wrapped VMEbus double Eurocard board. It acts as a buffer between the VMEbus and the three serial ports of the transmit controller: a high speed RS-449/RS-422 port (384 Kbit/s), and two CCITT G.703 ports (64 Kbit/s). This board is not currently used by FODA/IBEA, although the necessary software hooks to use it are in place.

The processor interfaces to the controller via a number of VMEbus memory mapped ports which allow transmit data to be written and control and monitoring data to be exchanged. The interfacing function is performed by two boards: the transmit modem interface and the line interface.

- Transmit modem interface This is a wire wrapped VME double Eurocard board. It contains all the channel coding, framing and symbol rate selection functions. It processes data presented to it through the 32-bit VME data bus and also provides status information for circuit and modem monitoring purposes. Transmit and control data are directed to a Data FIFO (First In First Out) which can store up to 512 bytes of data and control information at a time. The control information sets the channel coding hardware to the appropriate mode and

the burst transmission timing. The burst format is depicted in Fig. 3. An interrupt is sent through the VMEbus when the FIFO reaches less than half its capacity. Control data is also written to an event memory, which is used to synchronise burst transmissions to the transmit frame. This board is connected to the modem by a large shielded cable.

- Line interface board This board, which has no software interface with the VMEbus, contains the line drivers for interfacing to the modem and also provides the 16.384 MHz reference clock, derived from a 32.768 MHz oven controlled crystal oscillator situated on the line interface board. A coaxial cable carries the clock signal to the Down controller.
- Channel encoder This board has no software interface with the VMEbus. It performs the data coding functions.

The Down controller hardware consists of the following boards:

- Receive control processor Same as in the Up controller.
- Receive serial interface Same as in the Up controller.

The interfacing function with the demodulator is performed by three boards: the receive modem interface, a line interface board, and the channel decoder.

- Receive modem interface This is a wire wrapped VME double Eurocard board. It is controlled via a number of ports mapped onto the VME memory space which allow both reading and writing of status, control and data information by the processor. The circuit also provides the interface to the modem, for data reception and fault monitoring through a shielded cable. The decoded data stream is routed to the receive modem interface. This interface uses an event memory to perform real-time reception and decoding of the received bursts. The event memory is filled with information on the initial code rate and symbol rate of the bursts within the frame, together with their expected timing.
- Line interface Same as the Up controller.
- Channel decoder Same as the Up controller.

The most interesting feature of the burst modem is that its symbol rate, coding rate and transmit power level can be dynamically adjusted within a data burst. The transmit power of the individual data sub-bursts (see Fig. 3) is tuneable in steps of 0.5 dB in a 20 dB range. Each sub-burst can be transmitted at a different symbol rate (and, hence, different symbol energy) as required. The available symbol rates are 512, 1024, 2048 and 4096 Kbaud in either BPSK or QPSK modulation formats. In order to meet the dynamic rate requirements the modem is implemented using digital signal processing (DSP) techniques. The codec supports variable coding rates: *uncoded PSK* and $1/2$, $2/3$, $4/5$ punctured codes, derived from a $1/2$ convolutional encoder, thus providing one more way to adapt the information bit energy to the channel quality. Coupled with the symbol rate agility, this allows an overall information rate of 512-8192 Kbit/s in QPSK or 256-4096 Kbit/s in BPSK.

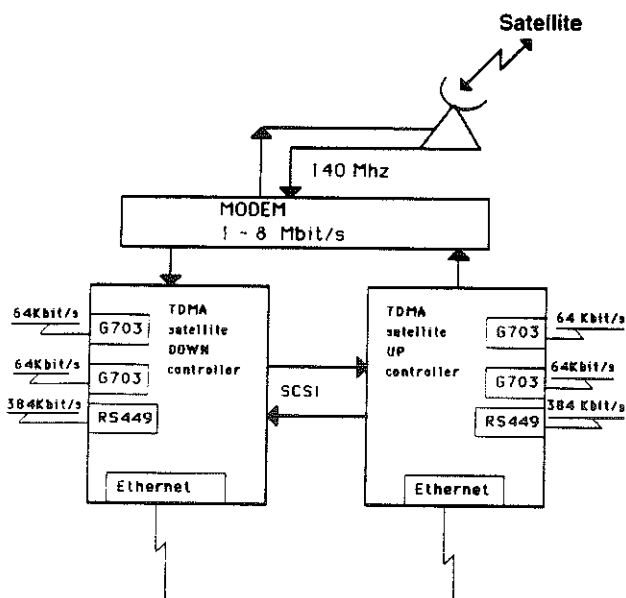


Fig. 2. The hardware configuration

The decoder operates asynchronously with 1 bit sign/3 bits magnitude soft decisions.

Depending on the BER requested by the application, the FODA/IBEA system software guarantees error rates better than 10^{-3} , $3 \cdot 10^{-5}$, $3 \cdot 10^{-7}$ and 10^{-8} , by adapting the power level, the symbol rate and the code rate to the signal/noise conditions.

A detailed description of the hardware can be found in [1], [2], [3], [4] and [5].

4. THE SATELLITE ACCESS SCHEME

4.1 Master/slave functions

FODA/IBEA-TDMA control is centralised. Any station can be the control station (*master*), while concurrently acting as a slave station. If faults arise, the master station may be replaced by any other station.

The issue of centralised versus distributed control is in constant debate. A centralised algorithm is more robust than a distributed one in most sorts of contingencies, because one station alone (the most powerful or the least faded one) is responsible for the consistency of the burst time plan. In the distributed case all stations are responsible, so all stations have to listen to each other, thus increasing the chances that some control information is lost or misinterpreted and the burst time plan corrupted. Distributed control, on the other hand, is more responsive to traffic variations, because only one round trip time is needed for the network to find out the traffic situation of its components, so highly bursty traffic sources are more efficiently dealt with, thereby increasing the global channel utilisation. Distributed control systems may compensate for the lack of robustness with complex recovery algorithms, which usually impose some additional channel overhead.

We chose centralised control because of its robustness, letting the complex datagram assignment algorithm (discussed below) compensate for the responsiveness deficiencies caused by the double round trip delay.

The functions of the master and slave stations are reported below.

The master:

- allocates time within the TDMA frame for the stream and the datagram transmissions of the slaves on the basis of their requests,
- broadcasts the fade level declared by each station,
- maintains the synchronisation of all the stations within the network,
- is responsible for maintaining the optimum transmit power level, which is used by all the slaves as a reference level,
- uses the frame space allocated to itself to send data over the network in the same way as any other slave station.

The slave:

- sends allocation requests for stream and/or datagram transmissions to the master.

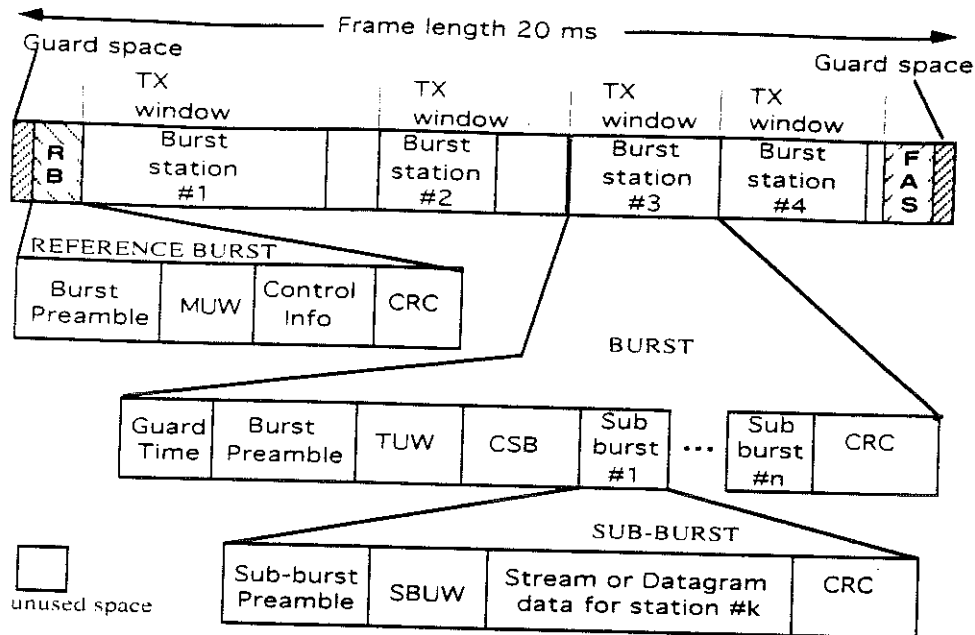


Fig. 3. Frame, Burst and Sub-burst Formats

- b) communicates its fade level to the master station.
- c) adjusts the transmit power level keeping the master as a reference.
- d) uses its allocated frame space to send data throughout the satellite network.
- e) chooses the symbol rate and the code rate for the single data sub-burst to send inside its transmission window, depending on the fade level of each link.

4.2 The frame format

The FODA/IBEA frame is 20 ms long: the format is shown in Fig. 3.

Picture legend:

- RB Reference burst.
- Burst Preamble For modem acquisition.
- Sub-bursts preamble For modem acquisition. Shorter than the burst preamble.
- MUW Master Unique Word. For modem acquisition. Relevant to the reference burst only.
- TUW Traffic Unique word. For modem acquisition. Relevant to each data burst
- SBUW Sub-burst Unique Word. Half the length of the MUW or TUW.
- Control info This part of the RB contains the stream and datagram assignments of each station, the fade level measured by each station, the frame number, the address of the next master station (in case of master recovery procedure), etc.

- CSB Control Sub-Burst. It describes the relevant data burst. It contains information describing the individual sub-bursts contained in the burst, such as the data destination, length, coding and bit rate, information relative to the disassembling and reassembling of the data, the datagram and/or the stream requests, the measured fade levels, etc.
- CRC Cyclic Redundancy Check.
- TX window Time for data transmission assigned to station.
- FAS First Access Slot.

As in all TDMA systems, the choice of the frame length is the result of a compromise between delay and jitter issues, which call for a short frame length, and overhead issues, which call for a long one.

The minimum overhead per frame (due to the reference burst and the guard times, in Fig. 3) is about 8.6% of the available bandwidth of 8 Mbit/s with the timings of the current prototypal hardware and the parameters chosen for the reference burst. The rest of the channel is available for burst transmissions, where every burst and every sub-burst within a burst needs some overhead.

The real-time data jitter issues have been studied in [10] which shows that the worst case jitter added by the FODA/IBEA system is two times the length of the frame, and the average delay added to the round trip delay is equal to the frame length. Our prototypal

software adds one more frame length delay, mainly because of CPU speed problems.

The *reference burst* is sent by the master station in order to assign the time allocations (transmission window times) for the transmissions of the requesting slave stations. The reference burst is always sent at a bit rate of 2 Mbit/s and a coding rate of 2/3, thus allowing a 10^{-5} probability of correct reception for down link fades up to 14 dB.

A station uses its assigned transmission window to send its stream and datagram data together. Data can be multiplexed inside a transmission window, allowing the datagram data to be transmitted in the portion of the stream assignment when possible, such as during the silence periods of the voice applications. Due to the assignment algorithms (described in 4.4), the part of a transmission window assigned to the stream data transmission maintains its size in every frame (if no explicit updates are requested), while the part for the datagram data may change size at each frame. Therefore, the transmission window sizes must be recalculated by the master at each frame.

Space for a First Access Slot (FAS) is reserved every 32 frames, by reducing the space available for data transmission, to allow a new station to enter the network. This slot is used in contention among all the new stations. The FAS is sized for sending requests, control information and a small amount of data, plus the uncertainty due to the current satellite position with respect to the nominal satellite position. This safety space is currently set to $\pm 150\mu\text{s}$. The FAS has a fixed position inside the frame (before the end of the frame). It must be accessed at 1 Mbit/s, 2/3 coded. This bit rate is lower than the bit rate of the RB because even heavily faded stations must have the opportunity to enter the satellite network. In this case, the faded station receives the RB affected by the down link fade only, but it receives its own FAS affected by both up and down link fades. The lower FAS bit rate is intended to cope with this situation.

The term *stream sub-frame* will be used hereafter to indicate the overall amount of frame devoted to the stream allocations. Likewise, *datagram sub-frame* indicates the amount of frame devoted to the datagram allocations. In both cases, it does not mean that all the stream assignments are concentrated in the first part of the frame and the datagram assignments in the second part. The terms merely indicate the amount of frame available to support stream and datagram traffic, respectively. In the absence of stream traffic,

the datagram sub-frame can be extended to the whole frame (apart from the reference burst space).

4.3 The Fade Countermeasure Technique

In unfaded conditions, the stream sub-frame is sized up to an NSUB (Normal Stream Upper Boundary) value. When fade occurs the stream transmissions can expand by increasing, if necessary, the stream sub-frame size up to a value called ESUB (Extended Stream Upper Boundary), and the datagram sub-frame is correspondingly squeezed, thus reducing the bandwidth available for datagram traffic. The data enlargement is due to the adaptation of the information bit energy to the channel conditions. In order to maintain the requested bit error rate, the code rate is changed; if changing the coding is insufficient, the data rate is reduced too. The fade level of the sending station and of the receiving station in the worst conditions (if more than one receiving station) are kept into account. The information relevant to the fade level of the stations in the satellite network is maintained in the reference burst.

Reducing the data rate by one octave means that the bit energy is doubled and that the E_b/N_0 ⁴ ratio increases by 3 dB. The coding rate can be reduced in smaller steps, and moreover the corresponding E_b/N_0 gain is greater than the bit energy increase, because of the Viterbi decoding technique. That is why reducing the coding rate is preferred rather than reducing the bit rate, as long as it is possible. Therefore, there is a relationship between the E_b/N_0 values of the sending and the receiving stations, the class of service which must be maintained, and the choice of the adopted coding rate and bit rate. In the FODA/IBEA software there is a complex table linking all these values in such a way that by using any of this information as input the other ones can easily be obtained. Data redundancy is achieved by increasing either the coding rate (R_c), or the bit rate (R_s), or both (R_i). The redundancy information factor R_i indicates the ratio between the energy per information bit of the data sent at certain bit and coding rates and the energy per information bit of the same data sent uncoded at 8 Mbit/s. R_i can be expressed as the product between the code redundancy R_c and the bit rate (speed) redundancy R_s . R_i expresses the data enlarging due to the increase in the code rate and to the reduction in the bit rate. When the signal attenuation increases, we compensate by reducing the coding rate first. We

⁴Bit energy-to-noise density ratio.

halve the bit rate only when E_b/N_0 drops under the value required for the burst acquisition.

The relationships between R_i and the fade level supported by the relevant R_i are shown in Fig. 4. The fade level represents the number of dBs of the available C/N_0 (carrier power over noise density) ratio below the value of 81 dB assumed as fade level 0 (reference level or clear sky condition). When operating in clear sky conditions at the bit rate of 8 Mbit/s, the E_b/N_0 ratio is 12 dB. With QPSK modulation, this value of E_b/N_0 allows the reception of uncoded data with a BER lower than 10^{-8} , which is commonly taken as an equivalent to a "no errors" condition for satellite transmissions.

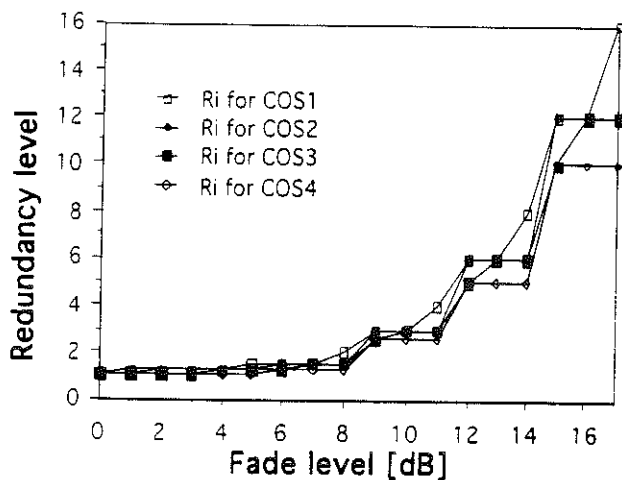


Fig. 4. Redundancy level behaviour

The choice of 8 dB as the minimum E_b/N_0 value necessary for modem acquisition comes from an experimentally measured burst acquisition miss rate of about 10^{-5} with our hardware. The non acquisition is measured by counting the number of non acquisitions, which account for missed carrier acquisition, missed clock synchronism and missed UW⁵ due to too many errors. The false alarm condition can be measured separately and its occurrence can be made as small as desired with respect to the other contingencies by choosing the appropriate UW pattern and length.

4.4 The requests and the assignments

Once active, a slave station can make requests to send stream data and datagram data during its transmission window. When the channel utilisation is low, the

stations will usually have a chance to transmit at every frame, even if they have no data to send. Even when the channel is very busy, every station is guaranteed to get at least one allocation every eight frames, as described in 4.4.2.

4.4.1 Real-time data

In clear sky conditions, to get a channel allocation, an application makes a request to a station via the GAFO protocol, indicating the requested channel space C_r , the minimum acceptable space C_m and the requested class of service (see Table 1). The value C_m is lower than C_r if the application can respond to a bandwidth compression request from the station; otherwise they are equal. When a new stream application needs channel space, the slave sends to the master a request which is the sum of its currently active stream application sizes plus the new one. The request includes the transmission overhead, and has a NEW flag, meaning that it is made because of a new application. The master grants a NEW stream request if the stream sub-frame size does not exceed the NSUB value. The rest of the channel is used for datagram transmissions. When a request for sending stream data issued by a slave station is accepted, the master maintains it until it is updated by the slave station.

During fading contingencies, a slave may need to increase the data redundancy of some of its stream applications in order to maintain the requested BER. In this case, the slave asks the master to enlarge its stream allocation, making a request with the EXTRA flag set. The master station acknowledges requests with the EXTRA flag if the stream sub-frame size does not exceed the ESUB value, which is greater than NSUB. In such a situation, the total channel space for datagram (the datagram sub-frame) is reduced for the duration of the fading event. The master can also "partially acknowledge" the request, so that the applications will run in degraded BER conditions.

If the master only partially acknowledges the increased EXTRA request of a station, the slave asks the stream applications (by using the GAFO protocol) to reduce their throughput to the declared configurations. If the applications cannot be sufficiently compressed, they will work in degraded BER conditions until some applications end and release the channel or the fade becomes lighter. It is up to the applications to decide whether to continue the session

⁵Unique Word

at a degraded performance or to give up, leaving their stream allocation to other users.

Both when some applications are compressed and when some of them work in degraded conditions, the station is said to be *suffering*, and it will periodically ask the master for more stream allocation space. The space it may obtain at each request will be distributed among the suffering applications.

4.4.2 Non real-time data

Each station manages stream and datagram data by making different requests for each type of traffic. The policies used by the slaves for requesting datagram space and those used by the master to allocate the datagram bandwidth to the stations are completely different from those used for stream traffic. In particular, stream allocations that are accepted are guaranteed until they are released, while datagram is never guaranteed, and the available bandwidth for each station generally varies on a frame-by-frame basis.

The request a slave makes to the master to send datagram data is proportional to the datagram traffic coming into the station (*traffic*) plus the volume of data already waiting for transmission to the satellite (*backlog*).

We have:

$$\text{request} = \text{backlog} + h \text{ traffic}$$

where h is a temporal constant of proportionality. Simulation results, obtained by loading the channel with Poisson generators of datagram traffic for 10 stations indicated a value of 0.4 s for the h parameter as a good trade-off between the average transmission delay at high and low traffic loads.

The slaves issue their datagram requests as frequently as possible, in order to give the master an up-to-date situation of the traffic entering them. The master organises the requests of all the slaves into a *datagram ring*, which it scans cyclically to compute the assignments. Any datagram request received from a slave already in the datagram ring is considered an update and replaces the previous value. The length of the assigned transmission window (a) is proportional to the request in a range of values between a minimum (T_{min}) and a maximum (T_{max}) threshold. We have:

$$T_{min} \leq a = f r \leq T_{max}$$

where f is the coefficient of proportionality in the assignment. In the current implementation f was set equal to the number of active stations divided by 100, with 5% as minimum and 50% as maximum. T_{min} was

introduced for efficiency purposes. This avoids 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 assignment, the relevant datagram 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 a complete scan of the datagram ring an *assignment cycle*. So, no more than one assignment cycle is made in a frame.

In each frame the master performs the following actions to complete the assignment procedure.

- It accommodates as many requests as possible according to the space available in the datagram sub-frame. Some stations then will have got an assignment, others will not, either because they did not make any request, or because there was not enough space available to satisfy all the requests.
- At this point a fixed assignment (previously reserved), called a *control slot*, is given to a station that has no assignment. The station is chosen with a Round-Robin criterion. If all stations already have an assignment, the control slot space is added to the unassigned space (which may even be null). In the frames where space for an FAS is reserved (once every 32 frames) a control slot is not assigned. After the control slot assignment, the unassigned space (if any) is divided among the stations as described in the following points.
- If all stations have a datagram assignment, the available space is divided among those stations which are still in the datagram ring, i.e. those whose request has not been exhausted.
- Otherwise, the available space is divided among all the stations. This is only possible if the resulting quotient is greater than the minimum datagram assignment.
- If this is not the case, the space is divided only among all the stations which did not receive any assignment. Again, the quotient must be greater than the minimum assignment.
- If this is not the case, the space is divided only among all the stations which have already received an assignment. In this case, it is not necessary to satisfy the constraint on the minimum because this assignment only is an enlargement of the space assigned in the first step of this procedure.

One control slot is assigned every eight stations, thus every station is guaranteed to get at least one allocation every eight frames. The control slot assignment coupled with this complex set of rules

allows a quick response time to traffic spikes at the slaves' inputs. When the system is lightly loaded, the system behaviour approaches that of a fixed TDMA (F-TDMA). We call this situation *pre-assignment* mode. The system gradually migrates towards a "pure" FODA/IBEA assignment scheme when the channel load exceeds a threshold (*pre-assignment limit*, [9]). A moderately loaded system can absorb abrupt traffic variations without appreciable delays, because each slave station has some spare capacity [13].

When the load is low (pre-assignment mode) the assignment cycle is always one frame long, and some spare capacity is always available to be shared among the stations. When the assignment cycle consistently exceeds one frame, no capacity is ever available to be shared. In this state the control slot is cyclically assigned, while the other rules for distributing the spare space are not applied because there is never a residual capacity after the datagram ring has been scanned.

A study of the performance of the datagram assignment algorithm is presented in [9].

Fading conditions cause an increase in the backlog and in the instantaneous traffic of the faded station and of the stations transmitting to it, because the same information needs more channel space to be transmitted at the same BER. This automatically increases the datagram request of the slaves to the master. Due to the increase in the datagram requests and to the possible simultaneous expansion of the stream sub-frame, which reduces the datagram sub-frame size, the overall datagram capacity of the system may become significantly smaller under heavy load conditions.

This scenario is not unlikely, and such a squeezing of the datagram capacity can be quick enough to call for a congestion control scheme. FODA/IBEA incorporates a simple backpressure method to relieve congestion, by blocking the growth of the backlog for a while when the internal queues grow too much. Since datagram data is collected from high speed networks, the only effect of this procedure is to slow down the datagram traffic coming from the remote hosts for a convenient interval of time. High and low water marks on the backlog queue of each station are used to detect the "danger" and "out of danger" conditions, and trigger the sending of the appropriate GAFO control messages to the gateway.

In [14] a procedure to make a rough estimation of the queuing delay is indicated.

5. THE PROTOCOL TO ACCESS THE SATELLITE NETWORK

An ad-hoc protocol (GAFO protocol [6]) has been developed to handle communications between the satellite network and the gateway interfacing it. The protocol provides for control messages and data messages. Specifically, the control messages allow the system to:

- set-up/close a connection-oriented session for stream applications;
- build/cancel/modify a sub-set of users allowed to receive data;
- exchange control information about the network addresses of the gateway, the Up controller and the Down controller;
- stop/resume sending datagram data (congestion control);
- exchange information about the status of other stations.

The data messages support :

- sending/receiving stream data with/without CRC;
- sending/receiving datagram data with/without CRC;
- sending/receiving test data with/without CRC.

The protocol was implemented in the C language on the Up controller, the Down controller and the gateway. The control packets are 12 bytes long, while the overhead for the data packets is 6 bytes. This information is not transmitted to the satellite link.

6. LANs INTERCONNECTION




Fig. 5. MALAN Logical Structure

While the CNUCE Institute was developing the satellite access scheme and the relevant hardware of the earth station, the Department of Electronic Engineering of the University of Florence was carrying out an experiment based on current

multimedia applications, to prove that satellite interconnected LANs can support voice, video and data transmissions [12]. This experiment, named MALAN (Multimedia Applications on LANs) was developed in a scenario characterised by two fundamental components:

- the communication link, consisting of earth stations and satellite.
- the local infrastructure, connected by the communication link and consisting of a collection of LANs (Ethernet and Token Ring).

Two fundamental prerequisites had to be fulfilled:

- to develop a set of significant applications involving voice, video and data;
- to package those applications in a portable system.

The logical configuration of the experiment is shown in Fig. 5, where each box labelled Multimedia Network represents a portable system together with the applications.

A Multimedia Network consists of several elements (Fig. 6):

- the *Gateway*, which connects the whole system to the Communication Link;
- the *Applications*, i.e.:
 - *Data* a file transfer application based on TCP/IP.
 - *Video* a real-time videoconference application.
 - *Voice* a real-time telephone link application.
- the internal Token Ring LAN.
- the external Ethernet LAN.

Each Multimedia Network is connected to the Communication Link through the gateway. It introduces a simple priority mechanism in favour of the stream traffic, as it can distinguish between information coming from the Data application and information from Voice and Video applications. The latter always has precedence over the former. In addition, the gateway measures the total traffic flowing back and forth from the LAN and the Communication Link.

Applications were selected in order to represent the most significant media: data, voice and video. *Data* is a TCP/IP based application and transfers files. The files being transferred are static images which are displayed as received. Each image is 153 Kbytes and the image format is VGA bitmap. *Data* could thus be defined as an image transfer application. In this way two goals are achieved:

- to significantly load the LAN,
- to visually control the results of the application.

Images are transferred iteratively, without any action by the user.

Voice provides a real-time telephone-like connection. Using a simple interface, users call a remote destination and start a conversation. Two options are available: voice sampled at 8 bits and voice sampled at 14 bits extended to 16, without any encoding. In the first option, traffic of 64 Kbit/s is generated; in the second, traffic of 128 Kbit/s is generated.

Video provides the core of a video-conference service, i.e. the real-time transmission of voice and video. Making use of the same interface applied in *Voice*, a remote user can be called and a remote conversation can start. Two communication options are available: 384 Kbit/s and 1664 Kbit/s.

The active elements in the Multimedia Network are connected by means of two LANs: a 10 Mbit/s Ethernet and a 16 Mbit/s Token Ring. Ethernet and Token Ring are selected alternately, so as to get an empirical comparison.

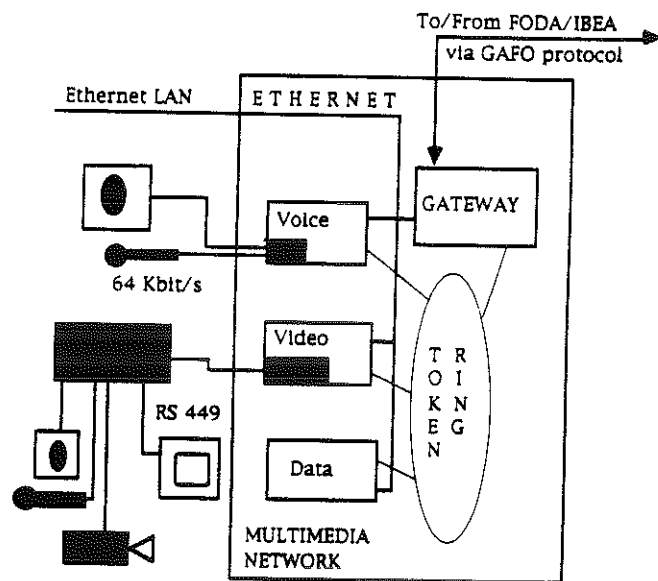


Fig. 6. The Multimedia Network structure

6.1 Physical description of the multimedia network

From a physical point of view, the Multimedia Network is a rack packaging four computers connected by two separate networks, an Ethernet and a Token Ring, plus a complex set of peripheral equipment for voice and video acquisition. The physical structure of the Multimedia Network is depicted in Fig. 6.

The boxes within the Multimedia Network are PC based on Intel 486 DX2, 33MHz clock, equipped with 8 Mbyte RAM and a 130 Mbyte hard disk. Each computer is equipped with two network boards: Ethernet and Token Ring. In addition, the gateway is equipped with an extra Ethernet to interface the Communication Link. *Voice* and *Video* are equipped with a board, named XIL-DSP, designed and developed by DEE. XIL-DSP contains a Field Programmable Array (FPGA) and a Digital Signal Processor (DSP), memory and logic to interface to the external world. The programmability of both the FPGA and the DSP allows the same board to be used for both voice and video.

A microphone and a loudspeaker are connected to *Voice* by means of an XIL-DSP board. A Rembrandt video-codec, made by Compression Labs Inc. (CLI), is connected through an RS-449 link to video. A monitor, a microphone, a loudspeaker and a camera are also connected to the videocodec. Each computer can be connected to a monitor and a keyboard. An MS DOS operating system is installed on each computer, and the driver for the communication adapter is a Packet Driver rel. 10 from Clarkson University. The drivers for XIL-DSP were specially developed by DEE. C is used as the programming language, with a small Assembler portion.

6.2 The experiment

The Olympus satellite was used for many hours to set up and test the TDMA controller hardware and the satellite access scheme. When Olympus became inoperative, the Italsat satellite was used as a test bed. A videoconference demonstration was held in Pisa on June 10 for the 13th Management Committee Meeting of the COST-226 project. In this case the Eutelsat satellite was used (12/14 GHz frequency band). The same demo was repeated few days later in the presence of representatives from several Italian Research Institutes, and the Italsat satellite was used (20/30 GHz frequency band). In both cases, the satellite was accessed by means of 3.5 m mobile antennas.

Three stations were involved, one in Florence and two in Pisa. There was a non-stop videoconference working between two stations. Each one sent and received traffic from two stream applications running on PCs, exchanging compressed digital video and sampled digital audio data through the FODA/IBEA system. Both local (Pisa-Pisa) and remote (Pisa-Florence) connections were established. The total traffic on the satellite channel always exceeded 2.5Mbit/s. Datagram connections were also tested

with TCP/IP packets, both alone and with stream traffic, involving all three stations in several configurations composed by different numbers of FTP and TELNET sessions plus the *data*, *video* and *voice* application described above. Data was transmitted at 8 Mbit/s, uncoded.

In this experiment *Voice* and *Video* were tested running in the presence of *Data*. In other words, *Data* was inserted to represent the background traffic sustained by the network during normal activities. The data transfer made use of all the spare capacity left by the stream application (video and/or audio). The quality of human interaction, both for voice and video, was evaluated on three levels:

- poor i.e., poor intelligible sound and video;
- avg i.e., average intelligible sound and video;
- high i.e., fully intelligible sound and video.

The results of the MALAN experiment were encouraging. Even the worst quality video, at 384 Kbit/s, was intelligible except for rare sequences characterised by rapid motion (as in cartoons). In any case, it is difficult to establish where the fault lies and 384 Kbit/s are probably intrinsically insufficient for quality video applications.

The most important results can be summarised as follows:

- Voice sampled at 8 bits is not good for all configurations. On the other hand, voice sampled at 14 bits, extended to 16, is always good, even in the presence of heavy background traffic.
- The same applies for video. Video at 384 Kbit/s is not good and video at 1664 Kbit/s is good.
- The maximum traffic available for the data application in the case of Token Ring is less than the one available for Ethernet.

With regard to the last point, it is surprising that a deterministic protocol, such as Token Ring, performs worse than a random access protocol, such as Ethernet, in the transmission of isochronous traffic. Three main accidental causes account for this experimental anomaly:

- 1) the priority mechanisms available in the Token Ring were not exploited in the experiment;
- 2) stream and datagram applications use short packets, up to a few hundred bytes long; this may favour Ethernet;
- 3) software drivers available in the public domain may be not optimised for the Token Ring.

CONCLUSIONS

In the first phase of the project the FODA/IBEA satellite network (software plus hardware) was developed which, up to now, is a prototype unique in the world of TDMA satellite networks which can adapt the robustness of the data transmission to the fade conditions of both the sending and the receiving stations. Inside the same transmission window data are addressed with different redundancy information values according to the atmospheric conditions of the stations. The performance of the FODA/IBEA satellite access scheme has been studied in depth [8, 9, 10, 13, 15].

As far as the second phase of the project is concerned, important conclusions can be drawn from the results of this project:

- the satellite link can be divided among many users, like any other terrestrial link;
- the already existing infrastructures (LANs) can be used for multimedia applications.

As far as the first point is concerned, only one satellite channel can interconnect several sites with a complete net. This is not easy to achieve with terrestrial links. Only one channel multiplexes stream and datagram data together, reserving the available bandwidth for the datagram after stream handling. It is possible to define high-level software protocols which reserve a minimum space for the datagram data and handle the stream sessions in an adaptive way (by reducing the performance if too many requests have to be handled).

As for the second point, we have demonstrated that the satellite can effectively carry multimedia applications between existing LAN technologies. This has advantages both in terms of costs and in the local diffusion of services.

REFERENCES

- [1] GEC-MARCONI Research Centre: "Working specification for Olympus TDMA equipment", Issue 2.0, Y/212/9883, August 1990.
- [2] GEC-MARCONI Research Centre: "Olympus TDMA equipment. Software Manual. Issue 2", MR 91/60A, Y/212/10441, September 1991.
- [3] GEC-MARCONI Research Centre: "Olympus TDMA equipment. User Manual. Issue 2", MTR 91/60A, Y/212/10440, September 1991.
- [4] GEC-MARCONI Research Centre: "Olympus TDMA equipment. Transmit and receive controllers. Hardware manual. Issue 2", MTR 91/60A, Y/212/10442, September 1991.
- [5] GEC-MARCONI Research Centre: "Olympus TDMA equipment. Burst mode modem equipment manual", MTR 91/60A, Y/212/10443, September 1991.
- [6] Celandroni N., Ferro E. "The GAFO protocol. The protocol to access the FODA/IBEA system", CNUCE Report C94-03, January 1994.
- [7] Celandroni N., Ferro E., Potortì F. "MTG. Multi-applications Traffic Generator. Presentation and Use", CNUCE Report C94-04, January 1994.
- [8] Celandroni N., Ferro E., Potortì F. "Outage probability of an adaptive TDMA satellite access scheme", proceedings of the IEEE International Conference on Communications ICC'93, Geneva (CH), May 23-26, 1993, pp. 1449-1454, Vol.3.
- [9] Celandroni N., Ferro E., Potortì F. "The performance of a capacity assignment algorithm for non real time traffic in a TDMA satellite system", submitted to the International Journal on Satellite Communications, September 1994.
- [10] Celandroni N., Ferro E., Potortì F. "Jitter removal for real-time applications in a satellite TDMA scheme", submitted to the International Journal on Satellite Communications, June 1994.
- [11] Celandroni N., Ferro E., James N., Potortì F. "FODA/IBEA-TDMA: a flexible fade countermeasure system for integrated services in user oriented networks", International Journal on Satellite Communications, Vol. 10, pp. 309-323, December 1992.
- [12] Bellini A., Pirri F., Guarducci M. "Multimedia Applications on Local Area Networks: a practical experience", 3rd Int. Conference on Broadband Islands, Hamburg, Germany, 7-9 June 1994.
- [13] Celandroni Nedo "Linear analysis of the transient for a datagram demand assignment satellite access scheme in TDMA", CNUCE Report C94-29, December 1994.
- [14] Celandroni N., Ferro E. "The FODA-TDMA satellite access scheme: presentation, study of the system and results", IEEE Transactions on Communications, Vol. 39, N. 12, pp. 1823-1831, December 1991.
- [15] Celandroni N., Ferro E., Potortì F. "The performance of the FODA-TDMA satellite access scheme measured on the ITALSAT satellite", ICDS-10, 15-19 May 1996, Brighton, UK, pp. 332-338.