A SATELLITE ACCESS SCHEME FOR TRANSMISSION OF ATM CELLS

Internal Report C95-13

March 1995

Nedo Celandroni
Erina Ferro
Francesco Potortì
Rupert Temel

Table of Contents

1. Introduction

In this report the design of a system for transmitting ATM cells over a satellite channel is presented. The proposed access scheme, named *SACDA-TDMA* (Satellite ATM Cells Demand Assignment) is a modified version of the FODA-TDMA system, described in [1].

ATM is a connection-oriented transport technology designed to provide a wide variety of services. ATM networks can achieve bandwidth efficiency through statistical multiplexing of cells from different connections, each of which may have widely differing traffic characteristics. ATM networks must manage traffic in such a way as to maximise bandwidth utilisation while minimising the probability of losses due to congestion.

The presented architecture is based on the following assumptions, which are defined by the existing prototype of hardware, realised on transputers and described in [4].

- No fade countermeasure technique is implemented.
- The control is centralised.
- Real time data (stream) and non-real-time data (datagram) must be transmitted simultaneously.
- The data coding rate is fixed to 1/2.
- The data bit rate is fixed to 2.048 MHz.
- The channel overhead must be reduced as much as possible; therefore stream and datagram ATM cells from the same station must be transmitted in the same transmission window.
- 2 unique words are used: the master unique word (MUW) and the slave unique word (SUW).
- A burst is a stream of data packets. Different destinations can be specified in each sub-burst.

All the aspects of the access scheme are discussed in a general way. Different solutions are given for some points. The implementation of a particular version depends on the defined system criteria.

2. System design criteria

Before an access scheme can be designed, it is necessary to answer and to clarify some important points, which determine the system design.

The system design depends on:

(1) *How is the end product defined ?*
The effort in the implementation should be adapted to the desired aim. It makes no sense to implement a very complex and high sophisticated access scheme if only three bridges will be available for experimental tests only. If the end product is a commercial device, more work must be done to implement a reliable system.

In any case, the software must be written in a modular way; in such a way the programmer can quite easily be able to extend the software from a simple system to a complex one.

(2) *Which is the maximum BER (Bit Error Rate) ?*
This is important because it influences the use of the CRC.

	BER $<= 10^{-10}$	CRC not necessary (in theory)
10^{-10}	$<=$ BER $<= 10^{-8}$	CRC can be used or not
10^{-8}	$<=$ BER	CRC necessary

In the reality, no commercial application is thought without CRC. Ethernet, for example, normally has a BER less than 10^{-10}; nevertheless it uses a 32-bit CRC for packets 1500 bytes long.
Surely the **CRC** is used for the burst control block (BCB), the reference burst (RB) and for the cell headers (in the following described).

For the following discussions, the BER is assumed to be greater than 10^{-8}, so a CRC is necessary.

(3) *How is the performance of the modem ?*
The length of the preamble necessary for the modem acquisition is important because the preamble is part of the frame overhead.

(4) *What is required to the system?*

It is important to understand that it is not possible to make an access scheme which is optimal for all the requirements, which are:

- system robustness
- minimum end-to-end delay
- preferably, no time jitter
- maximum bandwidth usage
- simple implementation.

Some requirements are in contrast to other requirements. For example, if the system is designed to be very robust, it is not possible to maximise the bandwidth use.

(5) *The types of the data sources.*

By *stream* we mean connection oriented applications that cannot tolerate out-of-order delivery of packets. In order to open and close a virtual channel, stream applications need to send and receive some control packets before sending the data.

We furtherly distinguish the stream traffic in Constant Bit Rate (*CBR*) traffic and Variable Bit Rate (*VBR*) traffic.

By *CBR* we mean stream real-time application data that need guaranteed bandwidth and fairly constant transmission delay. CBR traffic is generated by applications that have fairly constant throughput in the short term scale, i.e. on the scale of tenths of second. Such applications include voice, slow-scan TV, fixed rate videoconference, measurement data and so on. CBR applications declare their throughput (i.e. the requested bandwidth) to the system when they open the virtual channel.

A subclass of the CBR traffic is the Reducible Bit Rate (*RBR*) traffic. An RBR application generates a fixed bit rate traffic, but can accept a request to temporarily reduce its throughput. An example is the H261 video codec, whose throughput is constant, but reducible on command. RBR applications declare to the system both their normal and reduced throughputs when they open the virtual channel.

By *VBR* we mean a stream traffic whose throughput changes on the sub-second time scale, like that one produced by an MPEG video codec. The packet delay is generally variable, usually higher for packets generated during the throughput peaks. Therefore, some buffering at the destination end is required.

By *datagram* traffic we mean a connection-less type of traffic without particular delay requirements, that can tolerate out-of-order delivery of packets and a high jitter, but that usually expects a low bit error rate. Nothing is assumed about the kind of throughput a datagram source can have, in terms of both volume and variability; no virtual channel is opened, no throughput declarations are required. Each packet has an individual history.

Two priority levels are envisaged in this scheme for the datagram traffic. The high priority traffic is called *interactive*, because its natural application is conceived for interactive computer sessions. The low priority level is called *bulk*, because it can conveniently be used for applications like file transfer.

Datagram packets delivery is not guaranteed. It means that the system may ask the gateway to momentarily stop the datagram transmissions (for ex. when congestion is detected), and that packets exceeding the system buffering capacity may be dropped. Moreover, the delay a datagram packet experiences can occasionally be very high, in the order of seconds. Anyway, a minimum datagram bandwidth is always granted to each station, so at least the sending of control information is guaranteed.

3. The data queues

It is assumed that all the data interfaces which are connected to the satellite switch ("*bridge*") deliver the data in standard ATM cells format.

All the data queues are considered in the bridge. If the queues were in the data interfaces, a special protocol would be necessary to get ATM cells from the interface. These data interfaces then cannot be used as a standard ATM data sources in ATM networks.

One data queue is considered for each stream application.

One only data queue is considered for all the datagram applications. If there are datagram priorities, like between bulk and interactive traffic, one datagram queue is considered for each priority.

4. ATM cells transmission over the satellite

4.1 Standard ATM cells

Standard ATM cells come from the data interfaces. Two types of ATM cells are defined: the UNI cell (User Network Interface) and the NNI cell (Network Node Interface). The satellite switch is part of an ATM network, so the NNI format is relevant.

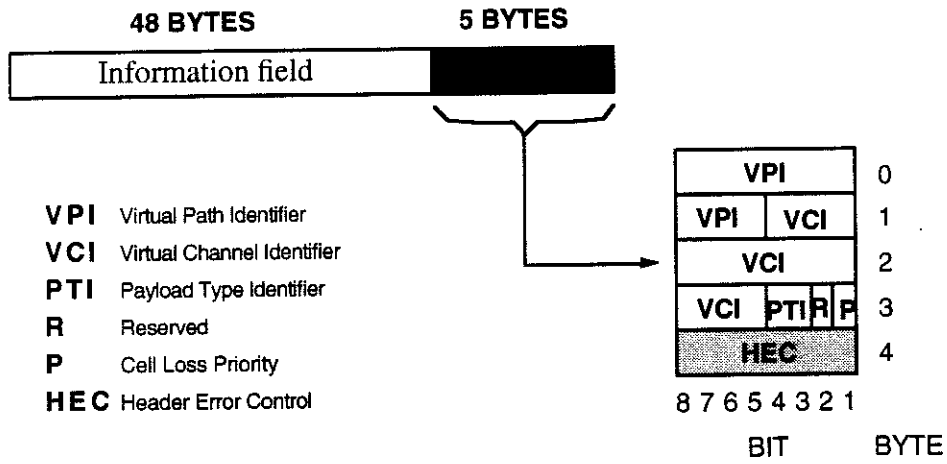

Fig. 1. Standard ATM Cell (NNI)

The VPI field (12 bits) and the VCI field (16 bits) contain routing information used in the ATM network. The HEC byte is an error detection/correction mechanism for the cell header contents used only to avoid miss-routing of cells.

The P bit is used to indicate whether the cell can be discarded or not by the network for congestion control.

When a transmission window is assigned to a slave station, ATM cells of both stream and datagram applications can be transmitted together. In absence of stream traffic, like during the silence periods of the voice transmission, datagram ATM cells can be transmitted in the space devoted to the stream transmission. Therefore, it is important that stream and datagram data have the same format when transmitted over the satellite. The data addressing is a problem that must be dealt with in two different ways, according to stream data or datagram data, respectively. Stream data does not need the destination address repeated each time, because a virtual connection between the sender and the destination must be opened before the real transmission. Datagram data is autonomous at each transmission, i.e. the destination and the source address must be repeated each time.

There are three possibilities to transmit ATM cells over the satellite:

(1) To add a satellite header to each cell, which contains the destination address. This is very simple, but it wastes bandwidth.

(2) To compress the standard ATM header. The 5 Byte header of the standard ATM cell is replaced by a 3 Byte header just for the satellite transmission. This is possible, because parts of the header are constant for an ATM connection. This solution is shown in detail in the next chapter.

(3) To multiplex cells. Cells with the same header can be grouped together by using one only header, longer than 5 bytes, for all these cells. Unfortunately, there are some disadvantages:

- Fragmentation/reassembling techniques are required;
- The allocation algorithm is not easy;
- More work is required by the receiver (which is critical) because it must rebuilt the original formats.

4.2 Internal cell format

The transmission of standard ATM cells over the satellite would be inefficient, because parts of the headers are constant for an ATM connection (fields VPI and VCI). The HEC byte is also not necessary if data is coded.

Therefore, a format for transferring ATM cells over the satellite is defined, named SATM (Satellite ATM format), for satellite switch internal use only. On entering the satellite switch, all the standard ATM cells are converted to the SATM format, shown in Fig. 2.

The VPI and the VCI fields of Fig. 1 are replaced by a 12-bit Virtual Link Identifier (VLI) field. The VLI is similar to the Virtual Link in the Transputer Systems, which defines a point to point communication.

When an ATM connection is established, a VLI is dynamically assigned to the new connection. The receiver of the destination station gets a new VLI entry in a RX table, which identifies the link where the cells must be addressed to be received by that station.

With this structure, the Multicast and the Broadcast facilities can be easily realised.

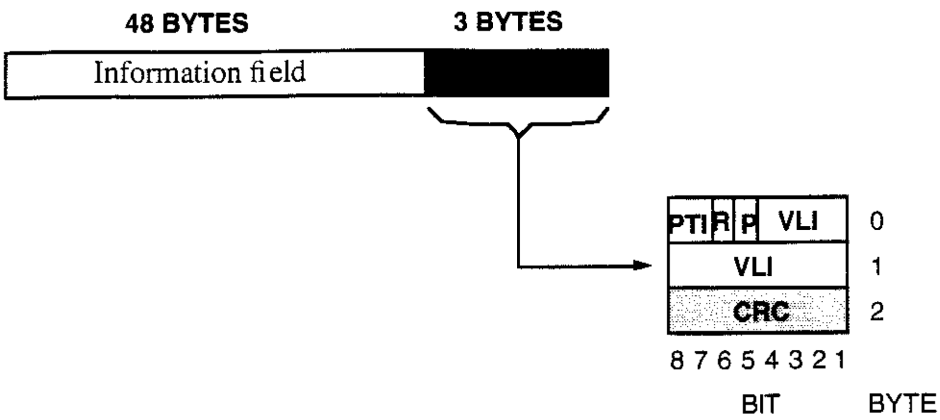

VLI Virtual Link Identifier
PTI Payload Type Identifier
R Reserved
P Cell Loss Priority

Fig. 2. SATM Cell format for internal use of the satellite switch

Remark:

It is not necessary to detect an error in the information field of the SATM cell because the end-user protocol is assumed to handle it. Error correction for reliable data transfer is generally performed through coding or retransmission protocol. For data connections using ATM, existing protocols such as TCP can be encapsulated into ATM cells.

On the other hand, it is very important to detect an error in the cell header, to avoid misrouting of cells. If the BER of the system is very low ($< 10^{-10}$), the header may be not checked. Otherwise a checking should be done and the CRC used. Therefore, the decision about the adoption or not of the CRC byte in the SATM cell header depends on the satellite quality. If not adopted, the SATM cell header must be considered 2 bytes long and the whole SATM cell 50 bytes, otherwise 1 more byte must be added.

5. Burst Formats

5.1 The Master Burst

As already said, the control is assumed to be centralised. The centralised control is handled by a *master* station which also acts as a normal slave station.

The master:

a) allocates the time within the TDMA frame for the stream and the datagram transmissions of the slaves, including itself;

b) communicates the allocations to the slave stations by sending a reference burst at the beginning of each frame, which is even used for synchronisation purposes;

c) uses the frame space allocated to itself to send data over the network in the same way as any other slave station. The only exception is that the data transmission of the master station, if any, always follows the reference burst, in order to save overhead on the channel. This burst, composed by reference burst and data burst, is called *Master Burst.*

The slave:

a) sends to the master station the allocation requests for stream and/or datagram transmissions;

b) uses its allocated frame space to send data over the satellite network.

For implementation simplicity it has been chosen that the reference burst has a fixed length, dimensioned on the number of stations the system is designed for.

The Reference Burst is always sent at the beginning of the frame. It contains the stream allocation S_i and the datagram allocation D_i for each station *i*. The allocations are valid for the transmission in the frame successive to that one where the reference burst has been received.

The allocations are expressed in multiples of the SATM cell. The amplitude of the transmission window (W) of a station *i* is the sum of both the allocations:

$$W_i = S_i + D_i$$

For any slave station the transmission overhead (preamble P, unique word UW and burst control block BCB) is included in the stream assignment or, if no stream is present, in the datagram assignment. In the master assignment, only the overhead due to the BCB is included,

as the preamble and the UW used for the data transmission are those ones used for the RB transmission.
The (S_i, D_i) values are used by the <u>slave</u> station i to compute its transmission start time:

$$S_i = P + MUW + RB + \sum_{j=1}^{i-1} W_j \qquad \text{if master is the station } k < i,$$

$$S_i = P + MUW + RB + \sum_{j=1}^{i-1} W_j + S_k \qquad \text{if master is the station } k > i.$$

Figure 3 shows the format of the master burst.

P	MUW	RB	BCB	Data

P	Preamble
MUW	Master Unique Word
RB	Reference Block
BCB	Burst Control Block
Data	Number of SATM-cells

Fig. 3. The master burst format

5.1.1 The Reference Burst (RB)

The RB contains the allocation information (S_i, D_i) for all the active stations which have been scheduled for transmission in the next frame. The RB format is shown in Fig. 4.

RBS	K	S1,D1	S2,D2	....	Sn,Dn	

Fig. 4. The reference burst format

RBS reference burst status. 1 byte where the bits are used as flags. Currently only the less significative bit is used. If ON, data from the master station follows the reference burst.

K number of the master station. 1 byte.

(Si, Di) Stream and datagram allocations for station i. Each allocation value is 1 byte long, therefore each (S_i, D_i) pair is 2 bytes long.

fnr Frame number. 2 Bytes.

CRC 2 Bytes. If the CRC of the reference burst is not correct, the RB is not valid and the station which received the errored RB is not allowed to transmit in the next frame.

The length of the RB is then $L_{RB} = 6 + 2*N$ bytes.

While, once obtained, the stream allocation is valid for all the frames up to when the sending station explicitly releases the bandwidth by means of a stream request for zero, the datagram allocation is not guaranteed and, even if present, it is generally different from the allocation in the previous frame. Therefore, being stream and datagram data transmitted in the same window (datagram always following the stream data), the presence (or not) of the datagram at the station *i* influences the starting point of all the following transmission windows, which generally vary frame by frame. The union of the stream and of the datagram data in the same transmission window induces two problems:

- a jitter is added to the stream data,
- if a station does not receive the reference burst, it is not allowed to transmit anything in the next transmission frame.

On the other hand, the bandwidth is very well utilised because the transmission overhead is reduced to the minimum.

Both the problems could be avoided if each station could allocate two separate transmission windows (one for stream and one for datagram) and if all the stream allocations were concentrated in the first part of the frame, followed by the datagram allocations. In this case the loss of the reference burst would inhibit the datagram transmissions only, not the stream transmissions which have the same amplitude and position frame by frame. On the contrary, the transmission overhead due to the preambles is, of course, doubled.

5.2 The burst control block (BCB)

The BCB is sent at the beginning of each data transmission. The format is described in Fig. 5.

id	nc	RDi	[redacted]

Fig. 5. The Burst Control Block (BCB) format

id Identifier of the sending station. 1 Byte.

nc Number of SATM cells in the data part of the burst. 1 Byte.

RD_i Datagram request of slave station *i*.. 2 Bytes.

The request is sent in each burst as a multiple of SATM cells. It must be a 16 bit value, because the request can be much larger then the assignment.

Remark: The stream request is sent by using the control slot, when assigned. In fact, the stream request must be issued only once and, if accepted, it is considered valid in each frame till when the sending station releases the bandwidth (with a request for zero) or it is declared dead (see Chapter 6).

CRC The CRC of the BCB. If not correct, the whole burst must be thrown away. 2 Bytes.

5.3 The slave Burst

Figure 6 shows the format of the slave burst.

P	SUW	BCB	Data

Fig. 6. Format of the slave burst

The BCB has been already defined in 5.2. The data part is a number of SATM cells, at least one.

6. Stream Request and Allocation

Stream connections require that the requested amount of bandwidth be guaranteed. But some stream devices exist which do not need the whole guaranteed bandwidth at any time, like, for example, video interfaces which have implemented a compression algorithm. If the picture is stable, less cells are sent as in case of a moving picture. The unused bandwidth of a station can be used to transmit its datagram traffic.

6.1 The request

Once declared active by the master station (see chapter 9), a control slot is assigned to a station with a fixed frequency. The requests for virtual connections (connection requests) between the sending station and one or more destination stations is assumed to be issued by means of the control slot. After a virtual connection has been established, the request for the stream bandwidth must be sent to the master, again by using the control slot. The format of the transmitted burst is as shown in Fig. 6.
The stream request must be issued only once and, if accepted, it is considered valid until the slave station decides to close the transmission or the station is declared dead by the master. The stream request is put in the Data field of Fig. 6, where Data is one SATM cell. The SATM cell format for the stream request has the format shown in Fig. 7.

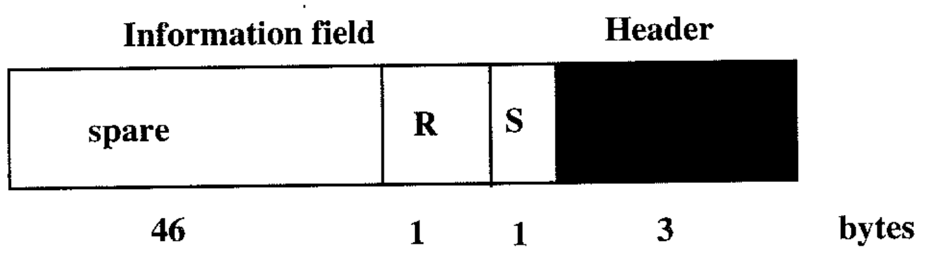

Fig. 7. The stream request format

S status. 1 byte. One bit must be set-up indicate that a stream request is present.

R stream request value. Number of SATM cells per frame required. 1 byte.

The unicity of the stream request implies that no space be wasted every time in the BCB (which is sent at each transmission!) to contain the stream request.
A station can send the stream request with two different meanings:

(1) *Global Request*

Each time a new stream application needs to enter the system, the station sends to the master a request which is the sum of all the stream requests, i.e. the new value is added to the previous value already assigned. In the master table of the stream requests, the new value replaces the existing one, if the increase is granted.

(2) *Single Request*

For each new stream connection the station sends the request relevant only to the new application. The new value, if granted is added by the master to the allocation of that station.

6.2 The assignment

On receiving a connection request, the master decides whether or not the requested bandwidth can be guaranteed. If the available bandwidth is less than the requested one there are two possible strategies:

(1) The master assigns the available. The slave station can decide whether to use the assigned bandwidth or to release it.

(2) The request is granted only if the whole request can be satisfied. Otherwise the request is refused.

The minimum assignment for a stream channel is 1 SATM cell / frame, plus the burst overhead (Guard time, P, UW, BCB) which must be added only if this is the first assignment of the station.

7. Datagram Request and Allocation

A datagram connection has no bandwidth guaranteed.

7.1 The request

The datagram request is sent by the slave station in the BCB at each transmit burst, apart from the first datagram request which must be issued in the control slot. Because there is (at most) one burst per frame of a station, the datagram request is sent once in a frame. If a station has no assignment, the control slot can be used to send the request.

The datagram request is a multiple of a SATM cells.

The request (R) is computed as R = B + H T, where:

B the station backlog, i.e. the number of SATM cells which are lying in the datagram queue, waiting for transmission to the satellite.

T number of SATM cells per frame. Number of cells arrived during the last frame.

H temporal constant. The product of H T is a sort of prediction of the backlog for the future time H. The optimum value for the H parameter is different for each system and depends on the system overhead (guard time, CCR, UW, BCB, RB). In the FODA system, the optimal value was 0.4 sec.

7.2 The assignment algorithm

The master organises the requests in a circular queue and cyclically scanned to compute the amount of time of each assignment. New datagram requests are put at the current beginning of the list in order to be scanned as first ones. This allows to reduce the delay time between the first request and the assignment after a period of no transmissions. Any further datagram request received from the same station (besides the first one) is considered as an update and replaces the previous value.
After each assignment, the datagram request is decreased by the assignment itself and the next request is analysed, if space is still available in the frame. The first assignment that does not

entirely fit the current frame will be analysed again as first assignment in the next frame, where the rest of the computed amount will be finally assigned. All the space up to the end of the frame (if insufficient for a minimum assignment) is given as an over-assignment to the last processed station.

If space is still available in the frame after an entire assignment cycle, that space is shared among all the active stations, even including those stations which had no datagram assignment in that frame.

7.3 The assignment quantity

The length of the assignment (a) is proportional to the request in a range of values between a minimum (T_{min}) and a maximum (T_{max}) thresholds. We have:

$$T_{min} \le a = f\, r \le T_{max}$$

where f is the coefficient of proportionality in the assignment. One suggestion from FODA is to choose f equal to the number of active stations N divided by 100, with 5% as minimum and 50% as maximum. T_{min} has been introduced for efficiency purposes. It avoids too small allocations when the transmission overheads -due to preambles and headers- are too big with respect to the information data. T_{max} prevents an overloaded station from removing too much capacity from the other stations.

8. Station Stop/Fail

- **Slave dead**

A slave station is declared dead when the master does not receive anything from that station for a certain number of frames. To avoid to be declared dead, a slave which has no traffic must send at least a stream request for zero channels or a datagram request for zero at the frequency of the control slot assignment. When declared dead, the stream and the datagram allocations of the station are cleared by the master. No control slot is assigned to an inactive station.

- **Slave stops**

When a slave station decides to stop the transmission, the allocated stream channels must be released by sending a request for zero. If a datagram request is still active (> 0), the datagram allocation procedure must be interrupted by sending a request for zero which will update the current value.

- **Master stops**

If the master wants to stop, it informs the new master station about it. The replacement must be done without interrupting at least the stream transmissions of the other stations.

9. Synchronisation

A small part of the bandwidth per frame is allocated for synchronisation purposes. It includes the *Control Slot* (CS) or the *First Access Slot* (FAS).
Fig. 8 shows the use of this slot.

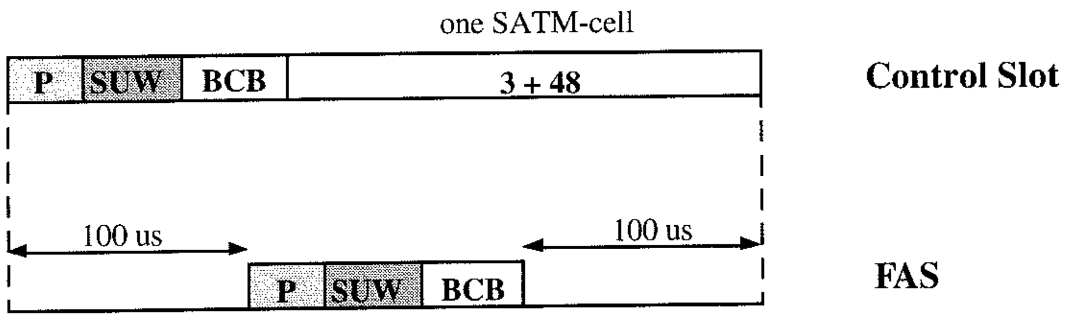

Fig. 8. Control Slot or First Access Slot usage

9.1 The Control Slot (CS)

The number of control slots present in the system depends on the number S of stations the system is designed for. If $2 \leq S \leq 4$, 1 control slot is sufficient; if $4 \leq S \leq 8$, 2 control slots are required and so on. In each frame the control slots are assigned on a Round Robin basis to those stations which had no assignment in the current frame, to give the active stations the possibility to send their datagram and stream requests.

The control slots are assigned by the master station after that even the datagram assignments are completed, in order to know which are the stations with no assignment. The assignment of the control slot, in the reference burst, has always the couple of values (0, 1+overhead). The position of the control slots in the frame can be anywhere after the master burst, as it is treated as a normal datagram assignment.

9.2 The First Access Slot (FAS)

The FAS is used to allow an inactive station to come up. The station sends a short burst, which only contains the Preamble, the UW and the BCB in the centre of the synchronisation slot of Fig. 8. Once declared active by the master station, the new station will use the control slot, at its first assignment, to enter the first stream or datagram request.
The FAS slot is present only once every 32 frames. In the frame where the FAS slot is present, no control slots are allocated. The FAS slot is sized as the sum of the control slots plus

the uncertainty due to the current satellite position with respect to the nominal satellite position. The FAS has a fixed position inside the frame (before the end of the frame).

10. System Architecture

The system architecture is proposed in Fig. 9, with the assumption that the input to the satellite controller are standard ATM cells. All the data interfaces deliver standard ATM cells. All the data queues are in the satellite controller.

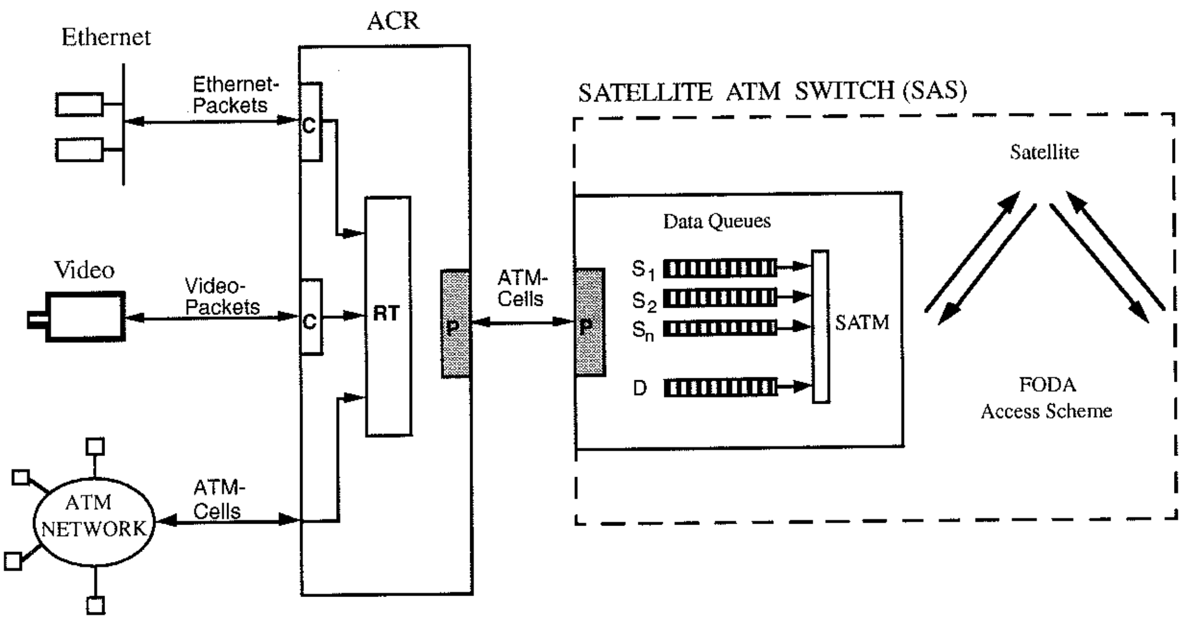

ACR	ATM - Converter & Router
C	Conversion to standard ATM cells
RT	Routing
P	Protocol between ACR and SAS
SAS	Satellite ATM Switch
SATM	Satellite ATM cell (3+48)
S_i	i-th Stream queue
D	Datagram queue

Fig. 9. The System Architecture

11. The Congestion Control

Congestion is defined as the state in which the network is not able to meet the required quality of service for already established connections and/or for new connection requests. A robust set of traffic and congestion control procedures is necessary to maintain a high network efficiency. Traffic control refers to the set of actions taken by the network to avoid congestion conditions, while congestion control refers to the set of actions taken by the network to minimise the intensity and the duration of the congestion.

In theory, congestion can be avoided by ensuring adequate capacity to support all connections. Then, congestion can only occur if more traffic is permitted than the network can handle. This is applicable to the fibre optic networks not to the satellite networks where the satellite capacity is a very expensive resource and bandwidth-limited compared to optical-fibre links.

11.1 Cell Discard

SATM cells with low priority can be discarded if the necessity arises. The user may set to 1 the P bit (cell loss priority bit) in the original ATM cell format (Fig. 1) to indicate cells which must be discarded with higher priority than others. When a congestion situation is detected in the station, the SATM cells with the P bit set to 1 (Fig. 2) are discarded so preserving, as long as possible, the P=0 cells flow. This procedure is quite simple from a control standpoint but the dropping of data on an already assured connection leads to retransmissions. Even if selective retransmission is used (i.e. the retransmission is requested of the missing data only), a certain number of packets will have a longer end-to-en delay; moreover it is quite difficult to guarantee a good quality of service for packets with P=1.

11.2 Feedback

The traffic sources which cause the congestion must be informed about the congestion event so that the source can reduce or suspend the data flow till when the congestion is finished.

The problem is to detect when the station is congested and how to inform the source(s) causing the congestion. When congestion is detected, the PTI field of each SATM cell of particular virtual paths or virtual channels contributing to the congestion is set to the "*congestion in course*" value. When received at the destination, this indication is sent to the higher protocol lays which must communicate to the end-user application(s) the necessity to reduce the traffic flow. Unfortunately, due to the high Round Trip Delay typical of the satellite networks, the end-user application cannot be really sure about the current congestion situation.

11.3 How to detect congestion

This mechanism is implementation-dependent.

REFERENCES

[1] N.Celandroni, E.Ferro
"The FODA-TDMA Satellite Access Scheme: Presentation, Study of the System, and Results", IEEE Transactions on Communications, Vol.39, No.12, Dec. 1991.

[2] N.Celandroni, E.Ferro, N.James, F.Potorti
"Design and Implementation of a flexible, software based TDMA System", IEEE Global Telecommunications Conference, December 6-9, 1992.

[3] N.Celandroni, E.Ferro, V.Mihal, F.Potorti
"FODA/IBEA-TDMA, Satellite Access Scheme for mixed traffic at variable bit and coding rates. System Description", Technical CNUCE Report C92-05, April 1992.

[4] P.Resele
"Design and Implementation of a Transputer Satellite Access Controller", Diplomarbeit, 1991.

[5] P.Resele, O.Koudelka, H.Klausen, B.Nocker, O.Lendl, U.Hofmann, C.J.Adams
"Satellite Transmission and Networking, Work Package 2.1, Analysis Report", ESTEC Work Order Frame Contract No. 10193/92/NL/LC, ASTP 4 ITEM BF 2.8, August 1994.

[6] P.Resele, O.Koudelka, H.Klausen, B.Nocker, O.Lendl, P.Burgstaller,
"Satellite Transmission and Networking, Work Package 2.2, System Design", ESTEC Work Order Frame Contract No. 10193/92/NL/LC, ASTP 4 ITEM BF 2.8, August 1994.

[7] M.May, P.Thompson, P.Welch
"Networks, Routers & Transputers: Function Performance and Application", INMOS, IOS Press, 1993.