Project	**IEEE 802.16 Broadband Wireless Access Working Group**	
Title	**A MAC proposal addressing 802.16 System requirements**	
Date Submitted	**1999-11-01**	
Source	Willie Lu, Mohan Maghera, Charles Bry Infineon Technologies 1730 North First Street San Jose, CA 95112	Voice: 1-408-501-6591 Fax: 1-408-501-2403 E-mail: willie.lu@infineon.com
Re:	Call for contribution – October 29th, 1999 – MAC Task Group	
Abstract	A MAC proposal addressing system requirements and based on fixed packet size to transport IP, LLC, or ATM directly is presented. Dynamic bandwidth allocation and wireless Quality-of-Service are guaranteed by using this protocol.	
Purpose	Submitted for discussion and evaluation	
Notice		
Release		
IEEE Patent Policy		

A MAC proposal addressing 802.16 System requirements

Willie Lu, Mohan Maghera, Charles Bry
Infineon Technologies

Introduction

The access method and duplexing scheme (FDD or TDD) should target the basic system requirements.
Basic system requirements, as we understand them, address a P-MP system operated at mm-waves.
The orientation of the regulatory allocation of bands presently in place is open to both duplexing schemes: single block allocations (TDD) where RF bands are very articulated (e.g. LMDS block A, ...) and paired allocation (FDD) which is the most adopted scheme outside USA (e.g. 26 GHz, 38 GHz ETSI and 38 GHz USA, ...).

The inherent higher cost of systems operating at mm-waves compared to cost of systems below 6 GHz (such as 3.4/3.6 (3.7) GHz, UNII, MMDS, ...) dictates a basic de facto difference between two classes of systems:
1) mm-wave systems: support of multiservice (IP, CES, FR where CES is a very qualifying portion of the overall traffic). Systems in this class are intended for Carrier operators (mainly CLECs), have a strong requirement on high availability (> 99.99%) and QoS support and final customers are large, medium (and small) business.
2) below 6 GHz systems: economical systems, with application in non line of sight and / or non licensed bands. Principal service: IP (including delay sensitive applications like VoIP and Video). System of this calss are intended for ISPs and final customers will be small business, SOHO and, in perspective, residential users. Availability requirements are much more relaxed.

From the above considerations we believe that the MAC access scheme for systems in the mm-wave bands should address primarily the following requirements:
- open to TDD and FDD
- multiservice support with a portion of CES
- QoS guarantees.

In addition packet-by-packet adaptive modulation support would be helpful.

Taking the above into consideration we recommend that the following proposal is examined.

Access method and duplexing

The proposed access method is TDM/TDMA over an FDD or TDD duplexing scheme.

Protocol stack Reference Model

The proposed Protocol Stack Reference Model is shown in the following picture:

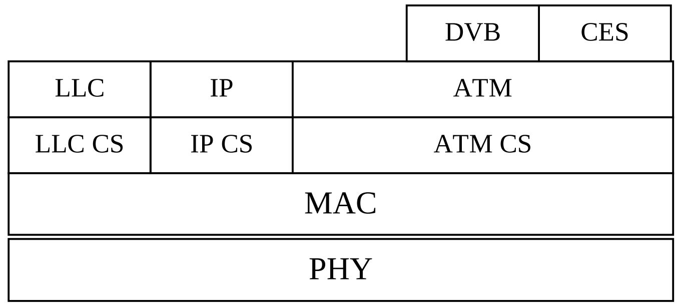

Fig.1

IP, LLC and native ATM are carried directly by the MAC while the CES services and Digital Audio/Video Multicast are transported via ATM through proper ATM Adaptation Layers according to the QoS required by the services themselves, while.

The MAC frame is depicted in the figure 2 along with physical layer encapsulation.
Connection ID and Payload is proposed to be FIXED SIZE. A size which is optimal for ATM transfer is 52 bytes. Different sizes can be chosen to trade off overhead for protocol speed and granularity.
A rationale for fixed size choice is reported ahead in this document.

UW = Unique Word for physical layer slot alignment
GAP = Silence between two different station transmit
PRE = Physical layer synchronization preamble
MAC = MAC fields (detailed ahead)

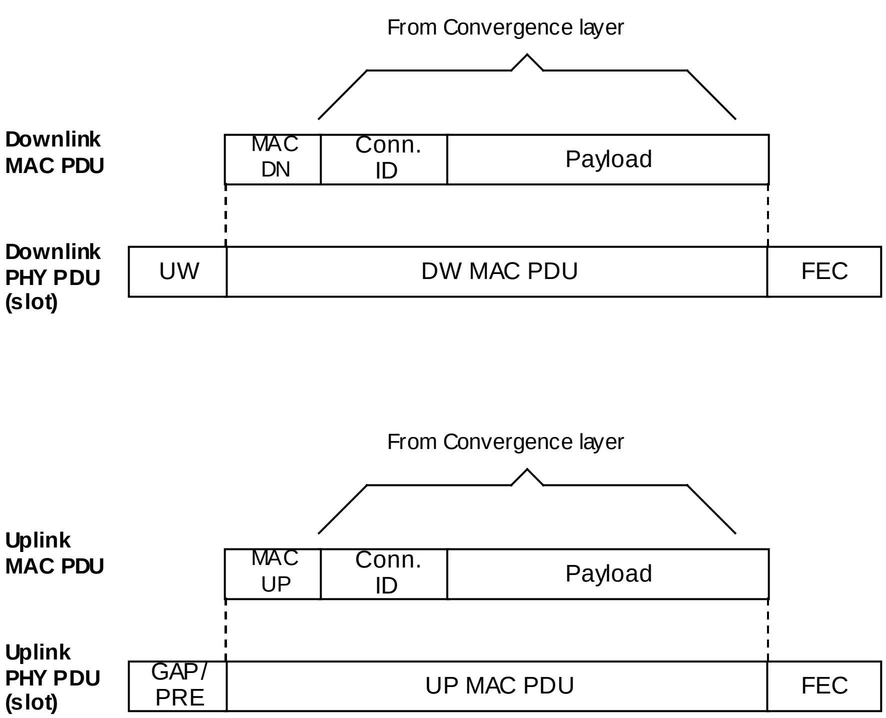

Fig. 2

Conn ID = Connection identifier (station ID + connection for IP and LLC, ATM header for ATM).

MAC DN field

It is composed by two fields: Grant, Type_down (possible total size: 1 byte).

Grant

The Grant contains the **internal identification code** of the selected subscriber station (6 bits for 64 subscriber stations) that is invited to talk in the next upstream slot. Note that this type of grant is "generic", in the sense that it doesn't address any particular connection inside the subscriber station. The subscriber station decides which connection to serve each time a grant is received via a scheduling algorithm. The scheduling is therefore performed in two levels: the base station assigns bandwidth to subscriber stations and each subscriber station redistributes the bandwidth to the connections. This doesn't mean that the MAC in the base station can operate independently of the connections which it must be aware of.

Type_down

It indicates in 2 bits if the downstream PDU is dedicated to an IP or LLC segment an ATM cell or a PLOAM, (Physical Layer O&M cell). PLOAM cells are specifically used to support physical layer functions, as indicated later on.

In the case of type field = IP (or LLC), the **connection ID** of the MAC PDU is the "**subscriber station** + **connection identification** inside the station" information plus the **final segment indication bit** (from 2 to 4 bytes are proposed for these fields according to the number of supported connections), while the payload field contains the segment of the IP packet.
IP packets are segmented and the last segment is marked; the segmentation and reassemble is thus very simple and relies on IP packet length field in the IP packet itself.
In the case that the use of the IP packet length field is not acceptable, it is possible to use the first payload byte of the last segment to carry the number of padding bytes in the last segment itself.

In the case of type field = ATM cell, the connection ID and payload fields of the MAC PDU are respectively the standard ATM header (with HEC byte suppression) and the ATM payload. ATM header compression is also possible.

Even if connection ID size can change between ATM and IP, the connection ID + payload length has to be constant (see rationale ahead). For example if a decision is made that the whole size is 52: in the case of ATM, connection ID is 4 and payload is 48 byte long; in the case of IP, the connection ID could be 2 bytes and the payload would be 50.

MAC UP field

It comprises the request information from the subscriber station and the type up information (possibile total size: 1 byte).

Request

6 bits to hold queue status for up to 6 different priorities. The queue status is coded in one bit as queue empty / not empty.

Note: the per-priority signaling and per-subscriber-station granting does NOT limit the QoS service guarantee because MAC only deals with bandwidth assignment TO STATIONS and CAN actually work on aggregate queue status and aggregate bandwidth requirement of all the connections inside one station. The station, though, will have to hold SEPARATE queues if QoS has to be guaranteed to connections.

Requests are piggy-backed in each traffic slot; special “minislot” are defined and assigned to silent subscriber stations to allow their wake-up. This assignment can be preallocated (via a frame delineation based on UW change) without need for a granting mechanism.

Type_up

It indicates if the upstream slot contains an IP or LLC segment, an ATM cell, or a PLOAM. Same comments as for Type_down.

Support to the Physical Layer and other functions

The MAC Layer will support the procedures and communications channels needed for the following operations (details will be discussed in the future if our proposal will be taken into consideration):

- adaptive modulation support
- admission of new subscriber Stations
- automatic ranging of subscriber stations.
- control of the transmitted power (ATPC) and time delay in the upstream direction
- alarms gathering
- encryption key exchange
- Request transfer not dependent on traffic slot availability (special minislots assigned periodically to subscriber stations serve the purpose).
- Grant transfer independent of downstream slot for use with TDD asymmetrical allocation (when upstream bandwidth is greater than downstream so that it is not possible to use downstream slots to carry upstream grants: this procedure would be slower but allow for full flexibility; the two procedures can coexist).

These and other operations are supported by PLOAM slots and special minislots.
It will be a task of the MAC Layer to schedule the suitable capacity for these operations.
Detailed discussion of these items will be presented at the next meeting if applicable.

Support for adaptive modulation

Details will be disclosed in the full proposal, if applicable, at the next meeting. An example is reported anyway to show the feasibility of adaptive modulation support with little overhead and additional complexity.

In principle our concept is based on a fixed time slot TIME (not number of bytes). If this allows carrying 52 bytes of Connection ID + Payload with 4-level modulation, it will allow carrying 104 bytes with 16-level and 156 with 64-level. System synchronization will remain simple, granting mechanism will not change. In the following the slot for a 16-level modulation in the adaptive modulation option is depicted.

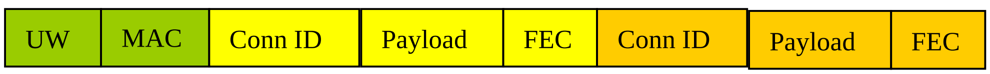

Fig.3

GAP, PRE, UW and MAC fields are always based on the 4-level modulation.
Connection ID, Payload and FEC can use different modulation levels.
Information in the MAC field should be changed as following: the Type (up or down) information will also contain the modulation scheme code.

Rationale for FIXED packet size

There are substantially 3 options for carrying IP:
1. Variable slot size unbounded: all the IP packets are transferred without segmentation no matter how big they are.
2. Variable slot size bounded: variable packet size avoid padding but segmentation function is kept to avoid too big packets.
3. Fixed slot size: the packet is segmented into fixed size slots with padding.

Option 1 can be hardly supported because packet size distribution in the figure 4 shows a substantial presence of packets that can be really huge. Their transmission could reserve the channel for too long (which is not a problem for a system carrying best effort traffic only but will certainly pose some limitations to systems carrying also QoS IP, CES and/or ATM).
Just to give an example: a 10 Mbit/s system in which a CES requires to send a burst every 1 ms can transmit 1250 bytes between two CES bursts. An IP packet of 1500 byte cannot fit in there, so it could not be transmitted. With ATM QoS the problem is even worse (and less predictable).

We will therefore compare option 2 and option 3.

Option 3 has a major disadvantage in padding overhead.
This can be calculated with current packet size distribution as seen in the figure 4. For a fixed slot size of 48 bytes the mean **padding overhead** turns out to be of **5.51%** (this is calculated as (total overhead bytes / total payload bytes) * 100)**.** The **MAC protocol overhead** in its basic form presented in this paper is around **2%** (also under the hypothesis of packets 48 bytes long).

Packet size distribution measured inside the Sprintlink network, on a transatlantic link and on a router at MAE-East (Peter Luthberg).Around 50% of packets are 40 bytes long, 15 % are 552 or 576 bytes long and 7 % are 1500 bytes long and the rest is equally distributed between 40 and 1500.
552 and 576 is an effect of the old Arpanet; 1500 is the Ethernet MTU and 40 is the acknowledge packet size.

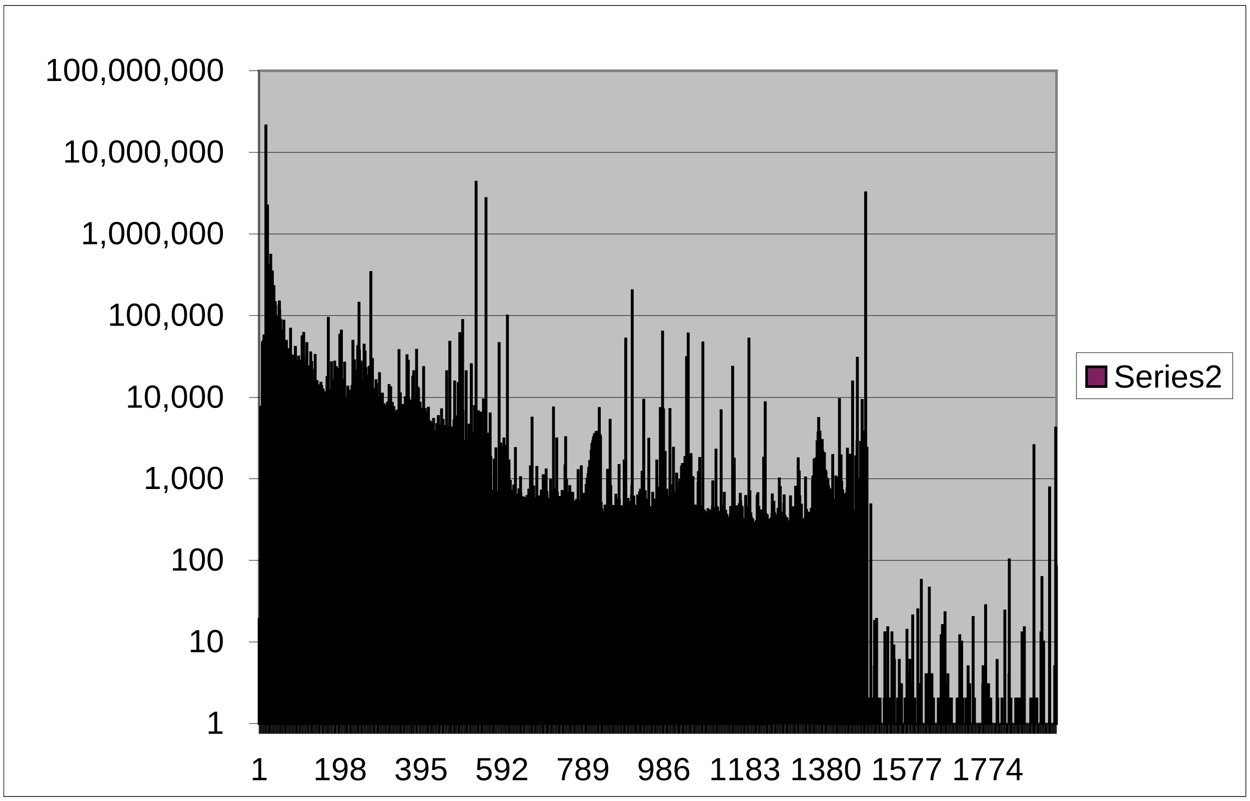

Fig.4

Option 2 on the contrary has several disadvantages:

- It still has a **protocol overhead** due to pointers and variable packet size management. The MAC protocol overhead of option 3 is around 2%. The protocol overhead of Option 2 will be in general higher and dependent on mean packet size. The final effect is to **partially reduce the Option 2 advantage related to the padding overhead saving, especially for small packets**.

Assuming a frame size of 1 ms (maximum frame size to support CES efficiently) and a minimum protocol overhead of 8 bytes per subscriber station and 2 bytes per connection that has to transmit in each 1 ms frame[1] it is easy calculated that the percentile overhead is given by:

$$\text{OH\%}= 100\left(\frac{2}{packetsize}+\frac{64000t}{bitrate}\right)$$

t (number of terminals) = $\min\left(\frac{bitrate}{8000\,packetsize},64\right)$ which means all the terminals that can fit in 1 ms up to 64.

Putting in in a graph of OH% versus packet size (different curves for different bit rates: blue=10,green=25 and red=50 Mbit/s).

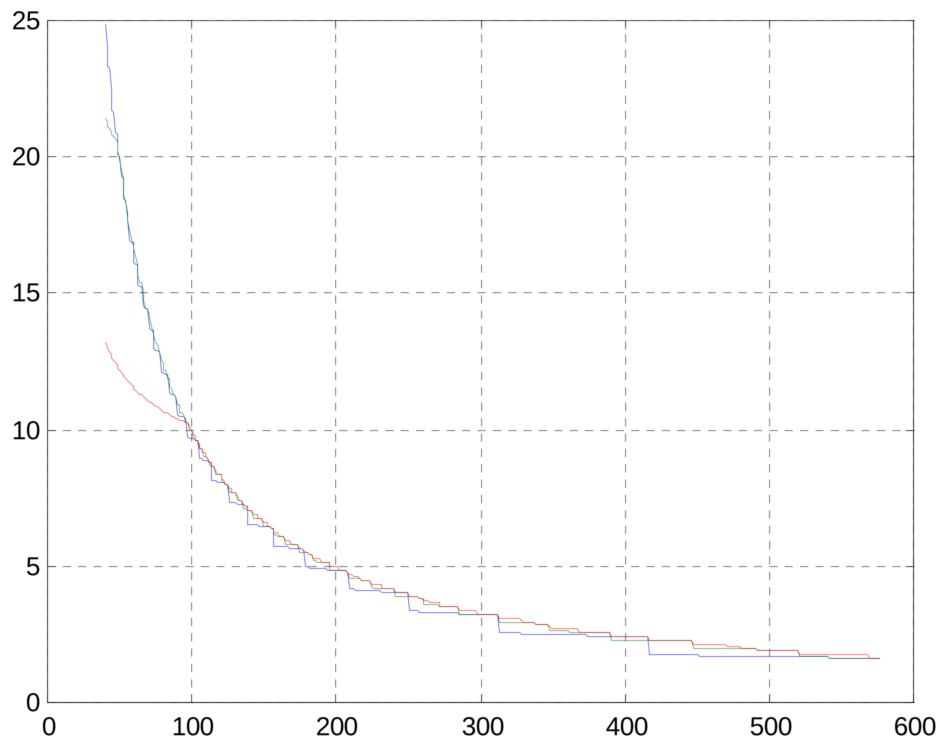

Fig.5

[1] This is in a way arbitrarily but reasonable; it is based on a known variable size protocol based on pointers to symbols to transfer indications of subscriber terminal start of segment and connection start of segment inside the terminal assigned segment. Grants follow the same mechanism.

This means a protocol overhead ranging from **13-25%** (when many **small packets** 40 bytes long are transmitted from different stations) **downto 1.7 %** when all the packets are **576 bytes long**). We did not consider packets longer than this because same considerations applicable to Option 1 hold in that case.

- The overhead mentioned in the previous point has to be paid also for **native ATM (and CES)** as an ATM cell is only a small packet. No padding overhead is paid of course for ATM in Option 3; therefore efficiency in the transport of ATM is **8% worse** than option 3.
- With variable packet size, **less flexibility is available for FEC choice** (impact on physical layer design); a convolutional code must be used (no block codes available) and convolutional codes happen to have an higher overhead at very low target BER; moreover they cannot be used for detection of errors (while block codes can) and if detection is required to discard the packet a second code shall be added.
- **Less robustness** due to critical information: pointers are very critical because the loss of a pointer causes the loss of an entire packet, error in the reception of a grant can cause the collision of two transmissions on the upstream channel and the loss of both; therefore pointers has to be protected more heavily than other information bits with both error correction and detection codes.

Based on these considerations and taken requirements into account we propose option 3 (FIXED PACKET SIZE).
We also point out that option 3 is seen as the most appropriate candidate to support adaptive modulation (slot-by-slot), which we will explain in the detailed description, if this proposal will be forwarded to the next meeting.

Evaluation on the grid

1. Meets system requirements: meets requirements in the system requirement document.

2. MAC delays: the lowest possible thanks to fixed (and small) packet size and the aggregated information transfer (that can be updated frequently)

3. Payload and bandwidth efficiency: see “Rationale for FIXED packet size” for fixed packet size overhead comments; MAC overhead is minimum thanks to per-subscriber station signalling and per connection scheduling.

4. Simplicity: never seen anything simpler.

5. Scalability: No known limits or critic points in scalability.

6. Service support flexibility: Apart from efficient support of ATM, IP and LLC all the other services are supportable via ATM encapsulation

7. Robustness: the loss of a MAC field has a very limited impact; the MAC recovers immediately from loss of link because it has no complex protocol and it is very fast.

8. Security: payload can be encrypted, key exchange is supported by PLOAM slots.

9. Physical channel configurability: 4, 16 and 64 QAM with fixed slot size (1, 2 or 3 segments transmitted per slot) can be supported. Details in the next paper.

10. Maturity: a very similar MAC is applied in PON (ITU-T G983.1)

11. Convergence with existing technologies: see point 10.

12. Independence of physical layer: supports FDD and TDD, no impact from constellation.

MAC functionality

In the following a possible MAC scheduler is shown. Main purpose of this section is not to get into details about MAC implementation but to show that a connection oriented MAC with a per-subscriber -station signaling allows very good efficiency (low overhead, fast reaction to changed traffic conditions) and easy QoS guarantees both on IP and ATM. This MAC is based on fixed slot size.

Each time slot is reserved to transmission from a single subscriber station that is activated in that particular slot by a message, called Grant, sent by the base station on the downstream channel.
MAC functionality, located in the base station, is in charge of generating these 'grant' messages in order to satisfy bandwidth requirements of subscriber stations.
MAC functionality has to be able to guarantee that bandwidth assigned to each subscriber station allows the fulfilment of QoS parameters of each connection belonging to one of the different classes of traffic defined in the international specifications.
Present MAC protocol is capable of not only serving connections related to narrowband and broadband services with required efficiency, but also to satisfy traffic parameters of broadband services, differentiating thus, within broadband services, treatment of connections belonging to different traffic classes. In the meantime it guarantees efficient use of all available transmission capacity of the upstream channel.
According to present MAC technique, upstream transmission is assigned by the base station to subscriber stations, following two different modalities:

- static allocation
- dynamic allocation

Static allocation is the periodic assignment of transmission slots to subscriber stations with a fixed rate.

This kind of assignment is used for serving constant bit rate services (ATM CBR, CES, IP GS) complying with strict delay and delay variation requirements.
Dynamic allocation is made according to two different mechanisms:

- one mechanism called 'Dynamic guaranteed bandwidth' allocation guarantees that a subscriber station has always the possibility to be assigned, for certain connections, a certain amount of "minimum guaranteed bandwidth" defined during connection setup phase. The subscriber station can momentarily free a part of or all this "minimum guaranteed bandwidth", according to its instantaneous bandwidth requirements;
- a dynamic mechanism to partition bandwidth that remains available distributes "available bandwidth" in equal parts among all subscriber stations that have traffic to be transmitted, thus implementing a fairness model perfectly attending criteria specified by standardisation bodies.

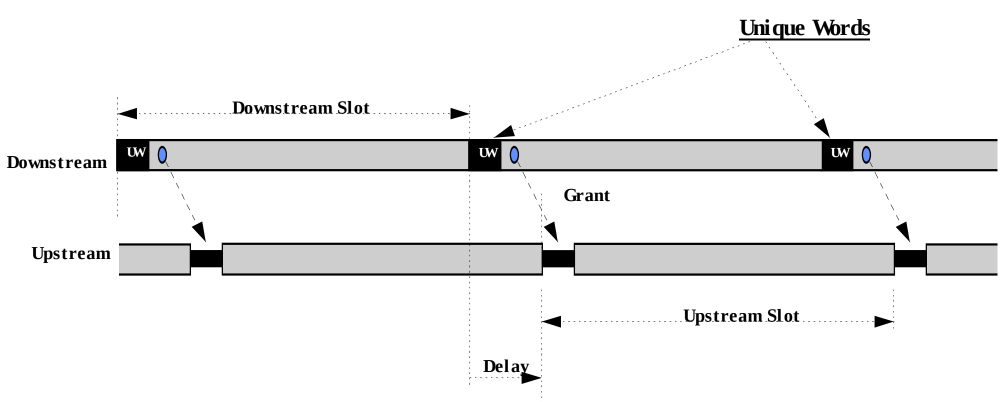

Fig. 6

The combination of these two dynamic bandwidth assignment mechanisms are powerful for serving IP with guaranteed minimum and ATM VBR, ABR, GFR connections, that is all the service classes that have minimum bandwidth requirements but can be statistically multiplexed. ATM UBR and IP best effort only rely on the available bandwidth.

- It has been previously mentioned that a subscriber station can momentarily free guaranteed dynamical bandwidth: it can modulate assigned bandwidth according to on/off mechanism. The subscriber station sends an aggregated information about queue status related to those queues which rely on the dynamical bandwidth allocation; this information, indicating whether traffic to be transmitted is present or not can be used by the centralised unit to momentarily inhibit bandwidth assignment to the subscriber stations that don't have traffic waiting to be transmitted in connections with dynamical bandwidth allocation.

The described protocol is able to treat in an optimal fashion bandwidth allocation for all current versions of QoS under discussion in IETF and all ATM classes standardised by ATM Forum.

The proposed MAC protocol offers instruments that allow base station to assign to each subscriber station a bandwidth which is evaluated on a instant by instant basis according to state information received by subscriber stations, always guaranteeing respect of traffic parameters and, in particular, of a minimum guaranteed bandwidth.

A second level scheduler will manage the assigned bandwidth between connections and classes of service inside the subscriber station.

MAC messages

In the following, MAC messages, an essential part of the protocol are defined.
In downstream direction the MAC master processor must send messages to assign the use of each upstream slot. The message that carries assignment information is called "Grant"**.**
Grants carry the following information fields:

1. Type of slot to be transmitted (user traffic, operation and maintenance, etc.);
2. Identification number of the subscriber station to which the slot is assigned (Connection ID or type of traffic is NOT carried).

In the upstream direction, each subscriber station sends to the base station indications about (instantaneous) queue status and instantaneous bandwidth needed for dynamic bandwidth allocation. This message is called "Request".

Information fields are:

1. Aggregated queue status: the field carries information about queue status in the subscriber station without distinguishing among connections but cumulating information of all connections that use dynamic bandwidth allocation. Static allocation connections are excluded.

This information reaches MAC master processor with a certain minimum periodicity, determined by a proper programming, as to guarantee fulfilment of all traffic parameters of active connections.
Request transport can take place with two different and coexistent modalities:

1. piggy-back in traffic slot: in the upstream slot transmission, the subscriber station to which the slot is assigned, transmits a Request as well as user traffic;
2. mini-slot: in special upstream slots of shorter duration, called mini-slots and assigned through particular Grants, a subscriber station transmits only a Request (plus physical layer preamble).

MAC processor

As already stated the base station hosts the MAC master processor.
The MAC master processor, through proper Grant generation, performs bandwidth allocation in order to guarantee:

- static bandwidth assignment (meaning by static that it can be modified by commands sent to MAC master processor, but non according to subscriber station queue status) of a certain amount of bandwidth to each subscriber station.
- dynamic assignment of bandwidth to priority connections. This allocation scheme can be influenced by subscriber station queue status.
- ex-equo assignment of bandwidth not assigned with previous mechanisms to all the subscriber stations that have traffic in the queue.

Static allocation corresponds to the assignment of a fixed capacity, equal to a constant grant rate, to a certain station, that is to a certain group of connections with constant traffic profile.

Configuration information for static bandwidth allocation transits through proper maintenance and control interfaces and reaches MAC master processor control interface in the Base station.
MAC master processor is not influenced by the status of the queue related to static allocation connections: for these connections no Request information is transmitted.
Static modality must guarantee a continuous Grant generation stream, correspondent to a certain predetermined capacity to each requesting station.

Grants needed to statically assign a certain constant capacity to a certain station are inserted in a table in proper positions; the table is shown as a circular stripe in the figure 7 to indicate that the table is cyclically scanned.
Grants extracted from the table during cyclical scanning are transmitted to the subscriber stations. Table dimension must be such as to obtain granularity and minimum allocation as requested by system specification.
Each position in the table corresponds to a slot and can contain:

- a static allocation user traffic Grant for a particular station
- an Operation and Maintenance Grant for a particular station
- a Grant for a group of minislots thus addressed to a group of stations
- an available grant for dynamic allocation (free position)

Dynamic allocation is mainly used for traffic with variable traffic profile.
In the figure 7 MAC master processor functional scheme that implements static/dynamic Grant generation is shown.

Functional blocks in Figure 7 are:

A *Priority Scheduler*, that selects the right FIFO from which Grant must be extracted in each slot choosing on the basis of FIFO status and priority. Static traffic is the highest priority.

As many *Grant Managers* as the number of handled priorities.
Slots that aren’t assigned statically are dynamically assigned to other traffic types. In the figure 7 three priorities are shown.
Grant Manager function is repeated for each priority; only one priority will be therefore described in the following.

Request Processor & Queue Status Mirror elaborates Requests sent by subscriber stations for that traffic

priority, rebuilding in dedicated registers the mirrored queue status related to each station.
When the mirrored status indicates the queue of a station is not empty, a ‘Pending request’ information related to that station is active inside the MAC master processor.
‘Pending request’ can’t be served until a down-counter associated to that priority of that station has expired.

This constraint realises an automatic flux control on the upstream traffic in the system.
The starting value of the down-counter is such as to generate a grant rate equal to the dynamic guaranteed bandwidth of the subscriber station (within that priority).
The starting value of the counters is chosen at connection setup and can be modified during normal operation by the information carried by the Request.
All the counters in the base station are decremented at each upstream slot time. When a counter after downcounting from starting point to zero, expires, it is reset to starting point and a binary information called ‘Expired Counter’ associated to related subscriber station is set to 1.
Activation of both ‘Expired Counter’ and ‘Pending Request’ at the same time for a particular station corresponds to the activation of a ‘Pending Grant’ information associated to that station.
Another process scans all the Pending Grants at each slot time and inserts in the Grant FIFO a Grant of user traffic type with the identification number of the related subscriber station for each station that has an active Pending Grant. At the same time the process sets to 0 the ‘Expired Counter’ information.
One Grant is extracted from the FIFO at each slot that is not reserved to higher priority traffics.
If the counter expires when ‘Pending Request’ is not active, the counter is re-set to starting value and then it resumes counting, while ‘Expired Counter’ is set and stays at 1, so that eventual later activation of Pending Request will immediately cause Pending Grant activation.

Bandwidth not assigned according to previously described mechanisms (that is when, on Grant generation time, no Pending Grant is active and no static assignment is preallocated) is distributed
in equal shares to all the subscriber stations that have at least one cell in queue in any connection which is handled by the dynamic modality (statically handled connections are therefore excluded) and related Grants are assigned cyclically to all the subscriber stations that have an active ‘Pending Request’.

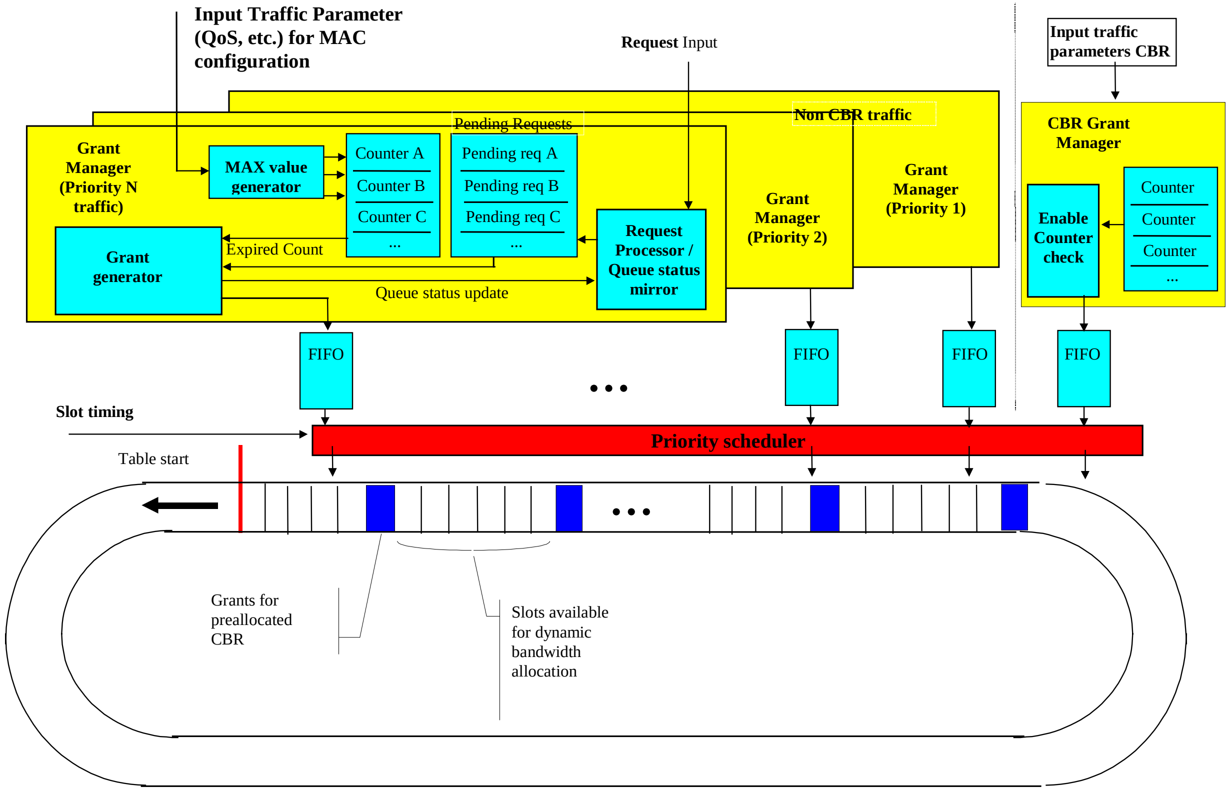

Fig.7

Acknowledgment

This contribution is fully based on the research and development activities undertaken by several project partners within the SIEMENS family. The authors present this contribution on behalf of these partners who are the major contributors to this proposal.