

Provision of QoS for Legacy IP Applications in an ATM-over-HFC Access Network

Jürgen Jähnert¹, Stefan Wahl², H. C. Leligou³

¹ Supercomputing Center University of Stuttgart
Stuttgart, Germany

² Alcatel Corporate Research Center
Stuttgart, Germany.

³ National Technical University of Athens
Athens, Greece

Jaehnert@rus.uni-stuttgart.de
swahl@rcs.sel.de
Nelly@telecom.ntua.gr

Abstract. With the increasing needs for delivering higher bandwidths to residential areas, new technologies (e.g. xDSL) are emerging that are enabling broadband access. Network technologies, which are based upon upgrades of existing network infrastructures, are, economically, one of the most promising approaches. One of these, competing with several other approaches, is hybrid fibre coax networks (HFC) currently being used for CATV broadcasting. The various access network technologies will be interconnected via broadband core networks resulting in heterogeneous networks which, have to provide end-to-end applications based on the Internet Protocol. Since the CATV network is of a shared nature, the provisioning of IP-based Quality of Service (QoS) requires suitable inter-layer signalling in order to provide QoS-related information as soon as possible to the relevant instances. This paper presents a system architecture using ATM-over-HFC for high-speed Internet access providing IP-based QoS. The architecture's seamless interoperability with other network technologies supports both IP-based and native ATM applications.

1. Introduction

Much progress towards high-speed networks has been made in recent years with respect to the bandwidth capacity of the backbone infrastructure. Also, in face of higher available resources in this section of the network, an effective exploitation of available resources is highly recommended since, from the application point of view, an increasing demand of this scarce resource can be identified. Solutions are required that enable customers to take advantage of the possibilities that future core networks could provide.

One promising approach is the Hybrid Fibre Coax (HFC) network. It consists of a fibre part for large-scale data distribution in city areas and a coaxial part into customer households. HFC networks form the basis of current cable TV infrastructures for the distribution of broadband TV signals. Thus, HFC systems combine two main advantages; firstly, they are widely installed and secondly, they are specifically designed for broadband communications.

HFC networks are characterised by a tree and branch topology. At the root of the tree, a head-end entity controls the traffic. The HFC network is a point-to-multipoint bus access network in the downstream direction, and a multipoint-to-point bus in the upstream direction. However, in order to support bi-directional broadband communication, the network has to be upgraded. For the physical layer, this requires the extension of the unidirectional amplifiers to realise an upstream transmission capability, introduction of medium access control (MAC) mechanisms and additional effort to guarantee privacy in a shared medium. On top of this physical layer, a transport layer is required to provide data traffic with a distinct QoS and a dynamic response to fluctuating demand. Future applications will be based upon IP and will be supported by several QoS mechanisms. Since interoperability between the QoS mechanism and existing applications is highly recommended, the chosen QoS protocol should be modular and transparent to these kind of applications. For this reason, within the ACTS project AROMA, such a system was developed, implemented and verified.

Section 2 describes briefly the overall architecture of the AROMA system, its individual network elements, concepts and issues concerning the physical and transport layer.

Section 3 follows with an introduction of the QoS service-provisioning concept of the AROMA access providing seamless interoperability with both, other access technologies and the core network.

Section 4 provides a summary of initial measurements made in the real AROMA network focussing on IP-based QoS for legacy IP applications.

The paper closes with a conclusion.

2. The Aroma HFC Architecture

Within the European funded ACTS project AROMA 13, an HFC access network has been designed and implemented that provides both a multiple services environment and Quality of Service support 1.

ATM is chosen as the mechanism to provide connection oriented traffic with a distinct QoS. Since the most important protocol used at the application level is the Internet protocol which is deployed in a classical IP-over-ATM environment, a suitable inter-layer signalling is required in order to provide QoS-related information as soon as possible to the relevant instances. This is to effectively exploit the scarce resources of the shared medium and to offer a core-network compatible QoS concept.

Fig. 1 depicts the deployed system architecture. The relevant access network elements are the Access Network Adaptation (ANA), the Access Network

Termination (ANT) and the Access Control Unit (ACU). The ANA is connected to the ATM core network and adapts and routes the ATM data stream to the HFC cell.

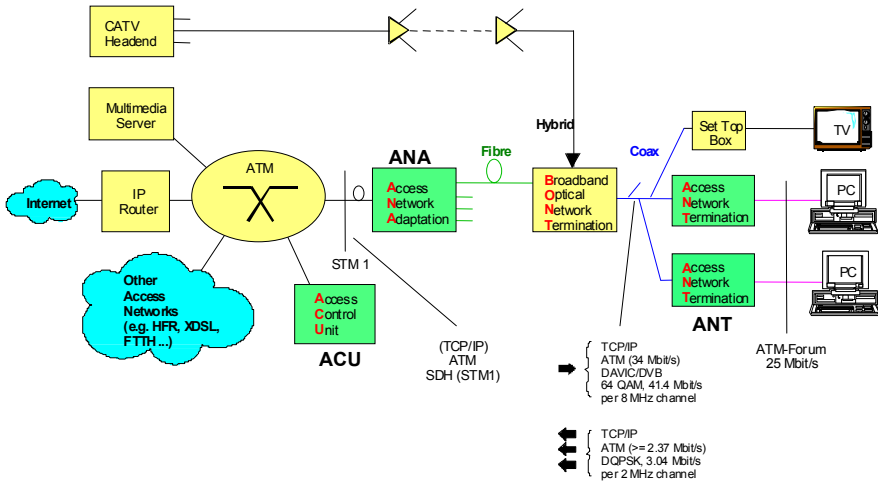


Fig. 1. The AROMA system architecture

A MAC controller in the ANA is defining, on a slot basis, which ANT is allowed to transmit data, controls access to the shared upstream channels. At the subscriber premises, the Access Network Termination (ANT) – also known as cable modem – provides an ATM Forum 25 Mbits/s interface to the AROMA end-system. Within the ANT, multiple upstream queues with different serving priorities are implemented. Furthermore, a hardware based piggyback feedback allows an individual and fast upstream bandwidth allocation on a per service-class basis. The Access Control Unit (ACU) performs the overall HFC access network management comprising the network element initialisation, user authorisation and the HFC resource management.

2.1 AROMA Dynamic MAC

The allocation of upstream slots in a tree-topology access system is based upon a reservation method which allows a dynamic adaptation of the bandwidth distribution to traffic fluctuations. The service policy of the MAC governs the distributed multiplexing from a central point situated at the head-end. Because of the much larger reservation delay and statistical behaviour of the aggregations from many customers, special care must be taken to safeguard QoS to sensitive traffic.

In the AROMA system, a mix of native ATM, IP-based quality services, better than best-effort and plain best-effort services are concurrently supported. The main tool is a QoS aware MAC protocol 7 with a suitable mix of bandwidth allocation mechanisms. This protocol relies on the following three innovative features: a DiffServ-aligned prioritisation architecture, the simultaneous collection of multiple requests from all class queues and polling of requests to guarantee worst case performance. The speed of the upstream is 7.2Mbps. The MAC protocol that was

designed and implemented within the AROMA project supports four different priority connections, which are sufficient to offer QoS as justified in 12. Logically separate queues for each priority connection are necessary on the cable modem side for the proper operation of the prioritisation scheme. Their service is based upon reservations.

Services with a similar set of requirements are grouped into 4 behaviour aggregates (classes) supported by four different priorities connections:

All delay sensitive periodic Constant Bit Rate (CBR) traffic is serviced through the highest priority connection (1st priority) which is supported by unsolicited permits scheduled during the connection establishment stage. Traffic generated from voice and video applications should be injected into the system through this class. In the DS context, this would be the expedited forwarding PHB. This topic is discussed in greater detail in section 3.

Real-time Variable Bit Rate traffic, such as video applications, should be injected through a 2nd priority connection, which employs the reservation scheme and is guaranteed to be serviced up to the contracted peak rate. This class is targeting the higher assured forwarding classes.

The third priority traffic class is suitable for services requiring a better than best-effort treatment. A minimum cell rate service is guaranteed, while traffic exceeding this rate is relegated to the 4th priority. This mechanism is suited to the support of all four or the lower three AF classes.

Plain best-effort services are grouped into the 4th priority, which is serviced on a request basis without any guarantee but with a fair distribution of the bandwidth not consumed, by the three higher priorities amongst all users in a round-robin fashion. This priority can accommodate the lower assured forwarding classes.

This flexible and robust scheme offers different QoS levels, bounded worst case performance, efficient bandwidth utilisation and a flexible pricing strategy. Under light load conditions, all the connections are serviced leaving the prioritisation scheme to provide the highest priority service with a better performance. Under heavy load conditions, the performance of each priority connection cannot be worse than that guaranteed leaving the 4th priority to suffer more system congestion. This behaviour will be verified in Section 4.

2.2 The Role of ATM in AROMA

The AROMA HFC access network is a pure ATM-based system. It provides the user with an ATM Forum UNI 3.1 interface and the capability to establish Switched or Permanent Virtual Connections (SVC, PVC) with accommodated QoS and traffic parameters. To support IP in the ATM environment, the concept of CLIP 16 has been applied. CLIP defines the address resolution for IP-over-ATM networks and the data formats of the data packets.

PVCs with dedicated traffic and source parameters are established between the connected AROMA end-systems and the IP gateway router. If an end-system has to send an IP packet to another IP instance, it has to request the corresponding ATM VC using the ATM Address Resolution Protocol Server. Knowing this ATM VC, it can encapsulate the IP packet according to 11 and send it via this ATM VC. The traffic

shaping functionality on the ATM Network Interface Card (NIC) has to take care of the negotiated/agreed traffic parameters for this VC. Leaving the ATM NIC, the ATM cells are transported via the ATM Forum 25 Mbps interface to the Access Network Termination (ANT). Within the ANT several ATM layer functions are performed:

- ATM cell filtering
- VCI/VPI translation for the HFC segment
- Payload encryption using a connection specific encryption key
- Routing the ATM cell to the connection-defined MAC service class

Reaching the Access Network Adaptation (ANA) at the head-end site, the VCI/VPI values of each ATM cell are extracted and are used to address the ATM connection context table. On the basis of the connection context, it is first verified whether the received ATM cell belongs to a known and active ATM connection. If it this is not the case, the ATM cell is immediately dropped and counted. Otherwise, the connection context is applied to decrypt the ATM cell payload and the VCI/VPI values are translated for the ATM core section. The ATM connection context table defines also whether the ATM cell is routed towards the ATM core network or whether the ATM cells are looped back into the HFC network. In the downstream direction, the ATM cells experience the same processing with the exception that there exists no multiple downstream service priorities.

PVCs can be installed and released via the management interface of the AROMA system involving the HFC Resource Manager (HRM). The HRM is responsible for the management of all HFC resources. These include, for example, the traffic classes, service priorities, upstream/downstream bandwidth, VPI and VCI ranges. If HRM accepts the set-up of a new PVC, it ensures that required ATM connection context values are installed within the related network elements (ANA, ANT). Furthermore, the HRM calculates the MAC control values from the service class and traffic parameters and transmits these values to the MAC controller in the ANA. In the case of a PVC, it frees the allocated HFC bandwidth and the ATM connection is set to the unknown state in the ANA and the ANT.

Ranging and keep-alive checks running on the ANA identify quickly whether an ANT is switched on or off. Each change of an ANT activity state is forwarded to the HRM. Beside the above-mentioned HFC resources, the HRM also manages all the network elements in the HFC network and performs the user authentication and authorisation. These comprise the user-related functionality, which is required for supporting Differentiated Services with user individual Service Level Agreements. Thus, the HRM is capable of installing all the required PVCs for a user as soon as he switches on his ANT. This includes multiple PVCs between two identical endpoints, each with different service and traffic parameters. The ATM ARP server has been modified as a multi-service ATM ARP server to support the service differentiation, as explained in the following section 2.

For the introduction of Switched Virtual Connections (SVC), all the above mentioned functionality is also required. One difference in comparison with the installation of PVCs, is that the SVCs are set-up and released by the ATM connection end-points according to the procedures and messaging defined in the ATM Forum UNI specifications. An HFC Signalling Manager (HSM) processes all the ATM signalling messages and has also an interface to the HRM for requesting the set-up of a new connection within the HFC section and for informing the HRM to release an

ATM connection. Adding appropriate rules to the multi-service ATM ARP set-up of SVC with application dependent traffic parameters can be initiated.

3. The Service Provisioning Concept of AROMA

The objective of the AROMA project was the development of an intelligent resource management system to be applied to an existing multilayered ATM-based HFC access network connected to ATM core networks which, in turn, are connected to the global Internet.

Providing new innovative HFC-resource management components targeted at the support of ATM and harmonising them with resource management which, in turn, requires inter-layer communications between relevant network layers providing guaranteed QoS both to native ATM- and legacy IP-based applications.

The nature of the MAC layer as already presented and its interaction with the ATM layer managed by the resource manager on one side and the customers with their demand to run applications over the access network on the other side, requires a mapping or signalling of customer requirements as described in Service Level Agreements (SLAs) for the ATM network layer. Applications could be based on several technologies. On the one hand, there are QoS-aware applications based on native ATM or supporting IP based signalling such as RSVP, and on the other hand, there are various applications currently used by the whole Internet community which could be regarded as QoS unaware.

The Service-provisioning concept of AROMA addresses both groups. Since the Internet is based upon a quite simple technology which is widely accepted, from the customer point of view it is not relevant what kind of technology is provided by the network, but it is required that existing applications which are partly long-term investments will work over it. Various, mostly asynchronous services such as e-mail completely comply to the user requirements. Other, mostly synchronous applications do not yet fulfil user needs because of the missing QoS support from the network in general.

Targeting these kind of synchronous applications, e.g. videoconferencing, appropriate mechanisms interconnecting the application layer with the AROMA-specific ATM layer must be provided.

The MAC developed within AROMA supports a differentiation into 4 traffic classes with each class different serving strategies. This is priority based and can guarantee QoS only for aggregated data streams corresponding to the 4 MAC classes. Reaching this point, it became obvious not to introduce RSVP as proposed by relevant standard organisations like DOCSIS 3 for QoS provisioning in this specific section of the network, but to use the more appropriate DiffServ approach because of its queuing mechanisms were similar to the AROMA specific MAC protocol. A further reason is the goal of AROMA to provide QoS also for legacy IP applications (QoS-unaware).

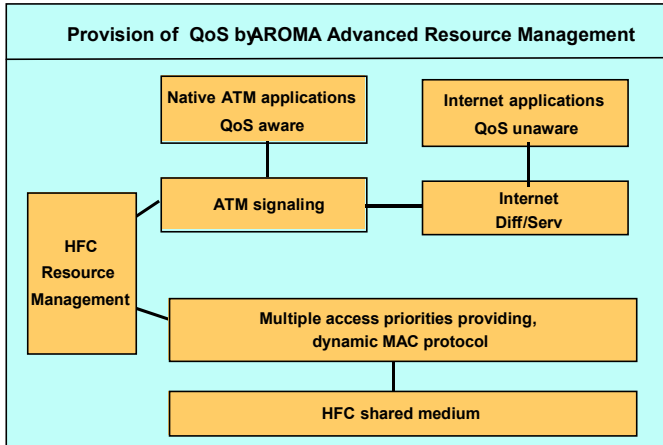


Fig. 2. AROMA Service provisioning concept

Fig. 2 shows the service-provisioning concept of the AROMA system. Since this paper is focusing on the provision of QoS for QoS-unaware applications, such as legacy IP applications, only the right vertical branch is dealt with within the following sections of the paper. Here QoS unaware applications are enriched with QoS in the AROMA system due to the deployment of the differentiated services (DiffServ) strategy. A DiffServ module groups the different services into behaviour aggregates (QoS classes) with a similar set of requirements. These service classes are attuned to the MAC control functions.

3.1 IP Based QoS Provisioning in AROMA

The differentiated services (DS) 9, 6 strategy recently adopted by IETF as a scalable and relatively simple methodology towards enriching IP services with a QoS, is applicable and quite appropriate for the case of tree-shaped access systems where IP services are dominant. To align such an access system to the differentiated services concept on the AROMA specific network infrastructure, requires the incorporation of provisions in the MAC function for the appropriate handling of each flow aggregation in respect to its requirements. Unlike the router case, it is not possible to offer such functionality as a software upgrade since the MAC requires a fast H/W implementation.

Each QoS class must encounter the specified Per-Hop-Behaviour (PHB) across the multiplexing points against the competing flows or at least not be delayed in a way that can not be retrieved in the next fully DS compliant node. This can not be realised in an HFC system without embedding suitable differentiation mechanisms into the MAC control function.

The basis of the approach is the use of access priorities in the reservation system, which can be programmed to fit in with required PHBs by means of the mapping of flows to priorities.

Within the ATM Forum, several proposals to implement DiffServ on top of ATM in general were discussed. Within AROMA, the so-called VC-bundle approach proposed by Cisco 8 was deployed.

A VC-bundle is a set of independent VCs terminating one IP address at each leaf end. Packets classified by the end-system are sent out according to their classification respectively and according to the Service Level Agreement (SLA) on different PVCs supported by the AROMA-specific MAC protocol over the HFC network to a DiffServ capable router terminating the VC bundle. Here, packets are leaving the AROMA specific Differentiated Services Domain (DS Domain) and are handed over to the Internet – to the next DS Domain. Exploiting features of the ATM NIC such as shaping the ATM layer could be regarded as access point for the DiffServ described in 8. On the other hand, the mapping from ATM connections to the different MAC queues combined with a piggy backing mechanism which in turn controls via MAC controller the resource allocation in the system up to the IP layer, a kind of interworking between DiffServ and the physical resources available in the system is established using ATM. Between the IP/DiffServ and ATM layer, the Cisco proposed VC-bundle concept comprising the second phase of the IP to ATM Class of Service (IP to ATM CoS) seems to be a promising solution for that specific network topology. ATM VC-bundle management allows configuring multiple virtual circuits that have different quality of service (QoS) characteristics between any pair of ATM-connected routers terminated each with one IP address. These VCs are grouped in a bundle and are referred to as bundle members. Using VC bundles, differentiated services can be created by flexibly distributing IP precedence levels over the different VC bundle members. It is possible to map a single precedence level or a range of levels to each discrete VC in the bundle, thereby enabling individual VCs in the bundle to carry packets marked with different precedence levels. Fig. 3 shows the AROMA adopted VC bundle concept.

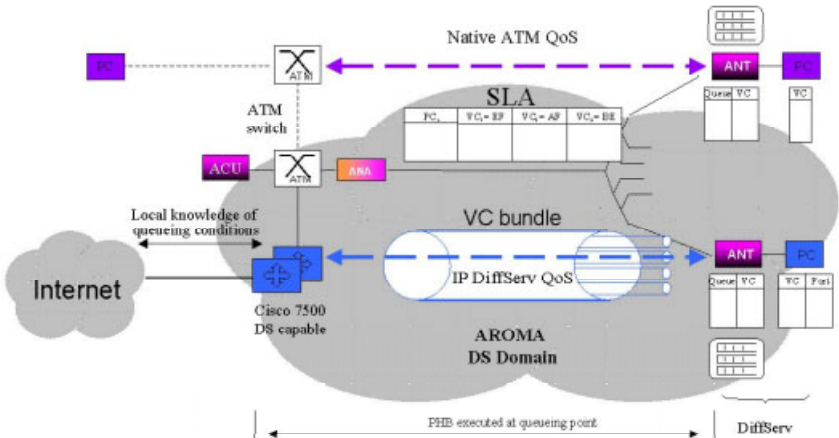


Fig. 3. AROMA DiffServ concept

To determine which VC in the bundle to use to forward a packet to its destination, an ATM VC bundle management software matches precedence levels between

packets and VCs. IP traffic is sent to the next hop address for the bundle because all VCs in a bundle share the same destination. However, the VC used to carry a packet depends on the value set for that packet header in the IP precedence bits of the ToS byte.

3.2 MAC Parameter Mapping

The ATM cell sized slots are the quantum of MAC assigned bandwidth allocation in AROMA. Several slots, enough to accommodate an IP frame, can be successively assigned, and this can be used to advantage in the context of the DS architecture over a shared link such as HFC. Namely, it allows suspending the transmission of a lower priority packet on the boundary of a slot (cell), and transmit delay-sensitive packets before resuming the lower priority transmission. In addition, fixed slots are easier to handle in a H/W based MAC. Such ATM slotting is, anyway, the approach followed by main standards bodies, such as DAVIC 17, DVB/ETSI 4, and IEEE 802.14 5. As in the centralised multiplexer case, flows with demanding QoS (e.g. Expedited Forwarding or Assured Forwarding) must be identified and receive properly differentiated treatment from the plain best effort traffic. This is accomplished in the ANT by placing the corresponding cells in the high priority queues which will subsequently place higher priority requests and activate the higher priority permit allocation MAC algorithms.

To reduce complexity in the cost-sensitive residential access system, services are grouped into behaviour aggregates (classes) with a similar set of requirements. This is in line with the DS philosophy of flow aggregation for better scalability and flexibility. The characteristics of the four aggregation levels/priorities are as follows:

The high priority is devoted to delay sensitive periodic CBR traffic, which is supported by pre-allocated pre-arbitrated unsolicited permits issued on a periodic basis by the MAC controller. This class is suitable for services with very strict delay requirements, which undergo strict traffic profile control (traffic conditioning) such as the EF (Expedited Forwarding) service 14.

The second priority level is devoted to real-time variable rate flows, such as video services or VoIP and it is provided with peak rate policing for guaranteed QoS. MAC exercises a policing function by rate checking before issuing the permits. This check is based on credit allocation at the time of subscription or connection set-up. In the DS context it could be used for the top AF (Assured Forwarding) class 15.

The third priority is devoted to data services with higher requirements than best-effort. The traffic profile control assumed for this class aims at minimising the loss of packets and the disturbance to other traffic. The credit scheme is used to guarantee a minimum rate (while credits last) while traffic exceeding this limit is relegated to the 4th priority permit generation. The 3rd priority mechanism is suited to support all four or the lower three AF classes 15.

The fourth priority is reserved for plain best-effort services which employ loss based flow control at the TCP level and can be very disruptive to the other classes when sharing the same queue.

It can be stated that within the AROMA system an inter-layer signalling is realised. This inter-layer signalling spans from the IP transport layer down to the MAC layer.

IP source and destination ports and addresses are used by the multi-service ATM ARP or VC bundle management software to select the appropriate ATM VC, which has been set-up before with the appropriate parameters. This ATM VC is assigned to an appropriate MAC service priority whereby the according MAC parameters are derived from the ATM VC traffic parameters.

4. Concept Verification

To verify the concept, initial tests in a heterogeneous ATM environment were undertaken. Here an ATM25 – ATM155 infrastructure as part of an underlying HFC infrastructure testbed was set up. Fig. 4 shows a conceptual testbed set-up to demonstrate the interworking of all three levels involved in QoS provision – IP – ATM – HFC/MAC.

Hidden from the terminal’s CLIP driver, a VC-bundle was established between the AROMA end-system with the DiffServ module and the commercial edge router. Served by different MAC queues according to the MAC’s QoS framework, the VCs of the bundle inherit the MAC QoS. Knowing the mapping VC-to-MAC which reflect the Service Level Agreement (SLA), IP packets eventually are fed into VCs according to the targeted QoS, the latter possibly covered by a SLA. According to the AROMA constraints, the VC-bundle consisted of 3 VCs and according to the IP application identified by the port number/IP destination, different VC’s are selected to transport the data across the ATM network.

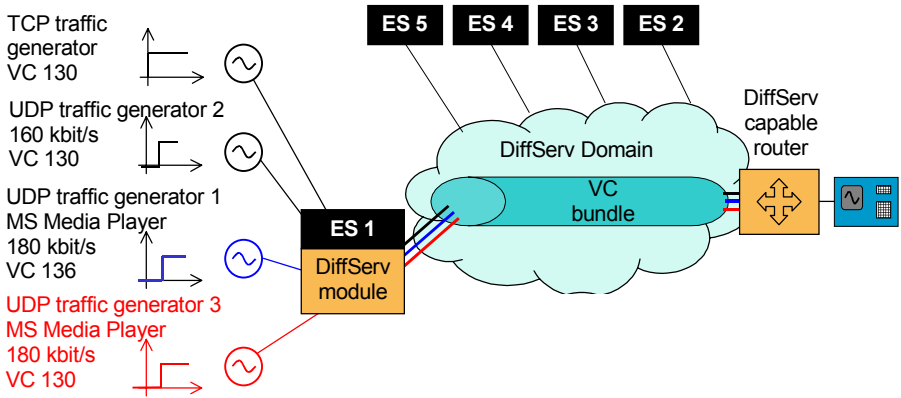


Fig. 4. :Measurement configuration

Next to this VC-bundle based leaf-end node, four further end-systems were connected to the shared access network receiving representative data streams. The tests were carried out with an ATM tester connected directly to the DiffServ capable router at the head-end side. Only non-corrupted cells were measured. As mentioned, QoS for Legacy IP applications are targeted, whilst off-the-shelf applications were used as follows:

The TCP traffic used in the test was a standard WIN2000 shell application.

The two real-time applications used were standard MediaPlayer data-streams directed to different clients differentiated by the IP destination address parameters. Furthermore, a UDP generator was used to disturb the low priority data traffic.

In this test scenario, the total upstream capacity is about 6.9 Mbit distributed to three upstream channels. Of the five end-systems (ES) plus modems 1 – 5, configured in the test-bed, systems 2 – 5 were silent but given an “SLA” of 1.6 Mbit/s for Queue 1, the pre-arbitrated queue. This artificially confines ES/Modem 1 to the total upstream capacity of about 550 kbit/s, which is allocated according to Table 1. Because the ATM NIC drivers do not support VBR, all VCs are using ATM CBR exploiting only the Peak Cell Rate parameter of the latter. In the given set-up, these pre-configured VCs are mapped to MAC Queues 1 and 4 respectively. With respect to the bandwidth figures given in Table 1, we can identify the following:

Firstly, the TCP rate and the sum of all queue values (in kbit/s) could have been increased to the theoretical maximum remaining upstream capacity for System 1 which is 517 kbit/s, but the ATM NIC does restrict it to 500 kbit/s per each VC.

Table 1. : SLA-MAC parametrisation

ES/ Modem	Destination/ Modem	IP	ATM				HFC MAC		
			VC	TC	Max Kbit/s	Mapped to HFC queue	Queue 1 Kbit/s	Queue 2 Kbit/s fix credit	Queue 4 Kbit/s
1	2	TCP	133	CBR	500	4			317
1	2	UDP	133			4			
1	3	UDP	136	CBR	500	4			
1	4	UDP	130	CBR	500	1	200		
2	-	-	-	-	-	1	1600		
3	-	-	-	-	-	1	1600		
4	-	-	-	-	-	1	1600		
5	-	-	-	-	-	1	1600		

For this setting the 200 kbit/s in Queue 1, allocated to the 1st priority queue of System 1, was definitely more than adequate. As can be seen from the graphs below (Fig. 6), the constant traffic shape reflects this fact. Initially, a TCP traffic stream using the FTP protocol was generated. This data flow was allocated the freely available capacity of the system.

The maximum transfer rate is limited by the ATM NIC parameter. After five seconds, three different UDP data streams (two real-time streaming applications – Windows MediaPlayer and one UDP generator) with different SLAs were generated (see Fig. 4).

According to Table 1, one real-time streaming application had a different SLA than the other three applications. The three data sources competing with each other mapped to the 4th quality group, having the lowest priority, transmitted at this point of time with a total bandwidth of about 800 kbit/s which was more than the system can transport since only about 350 kbit/s were available to this traffic class. Streaming data with the highest priority had an allocated bandwidth of 200 kbit/s, which results in an undisturbed traffic profile. It can be seen, that the running TCP application had immediately suffered from the UDP data transfers and rapidly decreased its sending rate. Cell loss occurs in the system but is restricted to the traffic mapped to the 4th queue – the traffic with the lower SLA. The measurements further show, that the TCP

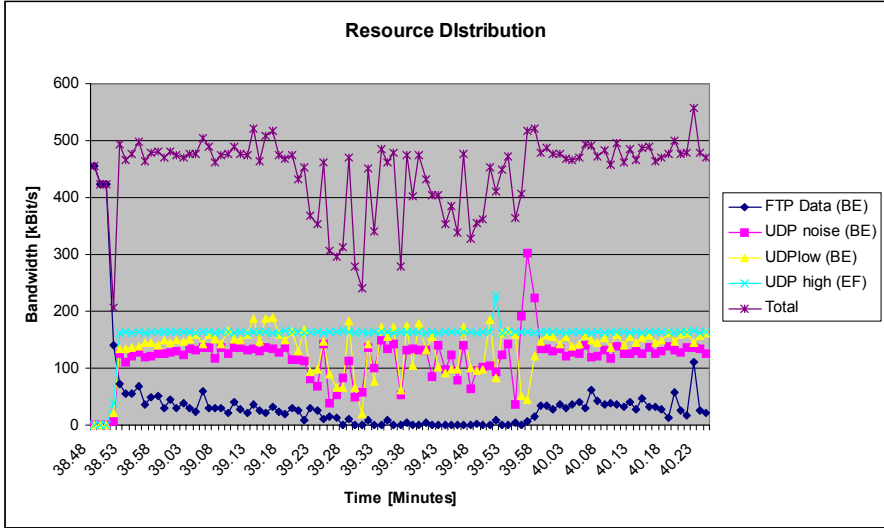


Fig. 5. Measurement result

flow does not really disturb the UDP streams in case of limited capacity because of its flow control mechanism. Further, the measurements verified that the MAC, which executed the QoS is highly dynamic and free capacity not allocated to a high priority traffic, but not used, can dynamically distributed to low priority traffic. This initial measurements verified the resource distribution and interworking between relevant layers in the simpler case of bandwidth distribution allocated to one end-system. Further measurements are required to verify the bandwidth distribution between different end-systems competing with available resources. The lack of early packet discard mechanisms in the system lead to the observed variance of the total traffic since only non-corrupted IP packets are measured.

Conclusion

The introduction of multimedia applications raises the need for adequate access networks capable of satisfying residential users with their requirements concerning bandwidth, delay or latency. This work has described an overall concept for a new network and service architecture, as implemented in AROMA, enabling high speed Internet access to residential users and providing a QoS concept to effectively distribute resources, especially on the scarce upstream branch of this specific network architecture.

The article focussed on the provisioning of QoS for legacy IP applications, but the architecture is flexible enough to support other technologies.

The architecture is partly based upon existing standards and extracts benefits from all of the above mentioned protocol layers. The required inter-layer signalling challenge solved within this architecture is not only restricted to HFC networks, but could become important also in further shared access mediums in general. This paper

described briefly the system architecture, the inter-layer signalling concept and finally presented some initial measured results. The open system architecture enables an interworking with other system architectures on the same network segment as well as an interworking with other access technologies and core network technologies.

This paper has shown that the architecture dynamically distributes available resources in a flexible manner to different customers according to agreed SLAs. A major goal of the architecture is the support of QoS for legacy IP applications which means that existing applications are supported with QoS without modifications.

References

1. J. Jaehnert, H. Fahner, et.al, " Interworking Technology and Applications - ATM-over-HFC-based Access to the Internet for Residential Users ", Proceedings of Intertworking98, Ottawa, 1998.
2. S. Wahl et. al, "Architecture and Experiences of a Multi-Service HFC Network", Proceedings of Conference on High Performance Switching and Routing, Heidelberg, 2000.
3. ITU-T Rec. J.112 Annex B DOCSIS: „Data over Cable Service Interface Specification“ (Same as ANSI SCTE DSS97-2)
4. ETSI ETS 300 800 Digital Video Broadcasting (DVB); DVB interaction channel for Cable TV distribution systems (CATV); 1st edition, March 1998.
5. IEEE Project 802.14/a Draft 3/R2, „Cable-TV access method and physical layer specification“, January 1999
6. IETF, Differentiated Services Working Group, RFC 2475 „Architecture for Differentiated Services“, December 1998
7. J. D. Angelopoulos, Th. Orphanoudakis, "An ATM-friendly MAC for traffic concentration in HFC systems", *Computer Communications Journal*, Elsevier, Vol. 21, No. 6, 25 May 1998, pp. 516-529.
8. ATM Forum TM 4.1, Enhancements to Support IP Differentiated Services and IEEE 802.1D over ATM, BTM-TM-DIFF-01.02 Draft, Dec '99.
9. IETF, Differentiated Services Working Group, RFC 2474 „Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers“, December 1998.
10. Dinesh Verma, "Supporting Service Level Agreements on IP Networks", Macmillan Technology Series, 1999.
11. IETF, Network working Group, RFC1483 " Multiprotocol Encapsulation over ATM Adaptation Layer 5", July 1993
12. J. D. Angelopoulos, N. Leligou, Th. Orphanoudakis, G. Pikrammenos, "The role of the MAC protocol in offering QoS to IP services over shared access systems", *Globecom'99Conference, Business applications session*, Rio de Janeiro 5-9, December 1999
13. <http://www.athoc.de/AROMA/>
14. IETF, Differentiated Services Working Group, RFC 2598 „Expedited Forwarding Behavior, June 1999.
15. IETF, Differentiated Services Working Group, RFC 2597 „Assured Forwarding Behavior, June 1999.
16. IETF, Network working Group, RFC 1577 "Classical IP and ARP over ATM", January 1994.
17. <http://www.davic.org/speci.htm>