[Wo99] *Lars Wolf;* **Multimedia applications in heterogeneous Internet/ATM environments;** INTERNET Research, Volume 9, Nr. 1, 1999, S. 49-57

# Multimedia applications in heterogeneous Internet/ATM environments

*Lars C. Wolf*

**The author**

Lars C. Wolf is Head of the Multimedia Networking Group at KOM, Darmstadt University of Technology, Darmstadt, Germany.

E-mail: Lars.Wolf@kom.tu.darmstadt.de

**Keywords**

Computer architectures, Internet, ISDN, Multimedia, Service quality

**Abstract**

New application types such as distributed multimedia applications have to provide a certain quality of service (QoS) to the users. Since they handle time-critical information such as audio and video data, they need appropriate support from the system components and especially from the network. New protocols and mechanisms have been developed over recent years to offer integrated services by serving both discrete media data (such as text and graphics) and continuous-media data (i.e. audio and video) in digital networks. Internet and asynchronous transfer mode (ATM) are the main players in this area and both possess QoS architectures which allow them to integrate services of data- and tele- communications formerly performed by separate infrastructures. We believe that both will co-exist for a significant amount of time, potentially complemented by other, perhaps simpler, approaches which are currently under investigation, such as differentiated services. Therefore, an interaction between these two architectures is necessary. In this paper, we discuss interaction approaches for the QoS architectures developed for the Internet and for ATM. We base this description on requirements and scenarios of multimedia applications and on the possible communication patterns considering different topological variants for heterogeneous Internet-ATM networks.

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · pp. 49–57
© MCB University Press · ISSN 1066-2243

## Introduction

Multimedia systems have attracted much attention during the last few years in society as a whole and in the information technology field in particular. The use of continuous media such as audio and video in multimedia applications brings new technical challenges due to the time-criticalness of such media which is uncommon for computer systems. Hence, to enable access to time-dependent information such as audio and video data, techniques must be developed and applied which allow for its handling in computer and communication systems and obeying time constraints.

Distributed multimedia applications have several requirements with respect to the service offered to them by the communication system. These requirements depend on the type of the application and on its usage scenario. For instance, a non-conversational application for the retrieval of audiovisual data has different needs than a conversational application for live audiovisual communication (e.g. a conferencing tool). The usage scenario influences the criticality of the demands. For example, a home video-conference, say between parents and children, is not as critical as a video-conference used as part of a remote diagnosis by a physician.

Typically there are various entities in distributed multimedia applications which co-operate in order to provide real-time guarantees such that the information can be presented at the user interface in time. The application requirements are defined in terms of quality of service (QoS) – the set of parameters which defines the properties of media streams. One of several QoS layers in networked multimedia applications (Steinmetz and Nahrstedt, 1995) is "Network QoS" describing requirements on network services (like network load or performance).

There are currently two main lines of work on network QoS; Internet integrated services (IntServ) and ATM (see Wolf *et al.* (1997) for

The work performed by the members of the "multimedia networking" group at KOM, especially Jens Schmitt, contributed significantly to this work. This work is supported in part by Volkswagen-stiftung, Hannover, and by Deutsche Telekom AG, Darmstadt.

an overview about general multimedia communication issues and techniques). It is often doubted (especially by data communications people) that ATM will really ever become a global internetwork. However, ATM plays currently and will play in future as well, a certain role, e.g. as backbone and link layer technology, even if it does not reach many desktops. IntServ has been designed over recent years as the QoS architecture of the Internet.

Recently, the questions of whether a reservation-based QoS is necessary and whether it can be provided in a cost-efficient and scalable manner have been raised again, similar to discussions in the early days of work on QoS in the late 1980s and early 1990s. This leads currently to the development of a separate line of differentiated services which allow us to treat traffic in different ways, but not keep much state information within the network. Nevertheless, we believe that reservation-based QoS methods are needed for several application classes for at least the middle-term future and also that IntServ and ATM will co-exist for that time. Hence, interaction mechanisms among them are needed.

Users are usually only interested in a good overall performance and the quality of their applications and do not care much about the underlying infrastructure and technologies used. To get satisfying overall application presentation quality, users are interested in a seamless and transparent interworking of their applications also in heterogeneous network environments. Hence, for an interaction between the Internet and ATM, the seamless interworking between the QoS architecture of ATM and the Internet integrated services architecture is important. That means, the provision of QoS end-to-end regardless of what is inside the network and whether both ends are located in the same world or not should be possible, i.e. providing a homogeneous service over a heterogeneous network.

This paper discusses application requirements and interaction approaches for Internet and ATM QoS architectures. In the next section we discuss application scenarios to see what kind and amount of interaction is needed. The next two sections then give a brief overview of existing ATM / Internet interaction models and compare their QoS architectures. This is followed by an examination of the communication patterns in heterogeneous ATM-IP (Internet protocol). Based on that, several interaction models for the QoS architectures of the Internet and ATM are derived. Finally, we conclude the paper.

## From applications to interactions

To deduce possible interaction patterns for the interworking of QoS architectures, we will consider application scenarios in this section. Several applications are envisioned for integrated service or multi-service networks as ATM and the new integrated services Internet claim to be. In Figure 1 some of the most prominent and demanding applications, as well as some traditional applications, are shown with respect to their bandwidth and delay requirements.

While all the applications under the dotted gigabit domain line could principally be served by existing networks (at least if there are not too many users requesting these services), the applications above the boundary line need the enhanced services of new networking technologies due to their high bandwidth and delay requirements. Although the new networks should provide enough bandwidth to serve these gigabit applications, it is important for them not to neglect the traditional applications and their requirements. That means that some kind of service differentiation is essential for these new networks, i.e. they must provide for parameterised services for the new applications.

Most of the complexity from these new applications stems from the media they are processing. These also now comprise, besides discrete media, continuous media like audio and video, which could also be categorised as soft real-time data. However hard real-time data also comes into play for real-time control of complex processes, and totally new media forms like virtual reality data have to be considered as well.

Examples for applications where audio and video streams are processed are videoconferencing and video-on-demand (VoD). Telemedicine, and especially telesurgery, provide examples of applications where, besides audio and video, hard real-time data for control of surgery equipment comes into play. Other examples are distributed interactive

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · 49–57

**Figure 1** Envisioned applications for integrated services networks

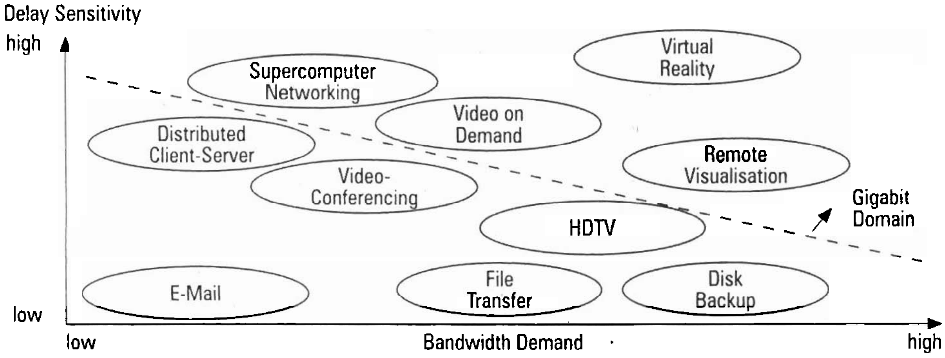


**Source:** Schill *et al.* (1997)

simulations, remote education, teleshopping, multimedia kiosks and distributed interactive games.

These applications can be classified in various ways, e.g. from an application domain or from a communications point of view. In the latter case, interactive applications can be divided further in conversational (e.g. videoconferencing with application sharing, or videotelephony), messaging (multimedia mail) and retrieval applications (video servers, multimedia-enhanced World Wide Web), tele-action applications (e.g. alarm and surveillance with a remote camera control), and tele-operation applications to exert remote operations like telerobotics combined with audio-visual control. Distribution-oriented applications are, for example, pay-per-view television, and near video-on-demand, while true video-on-demand is a hybrid form of a distribution-oriented application with some elements of interaction with the user.

Here, we will concentrate on video-conferencing and VoD applications. We will use them for the analyses of the possible interaction patterns between Internet and ATM, since they are the most prominent representatives of interactive respectively distribution-oriented application classes.

### Video-conferencing

In recent years, many solutions for video-conferencing have been implemented based on different compression schemes and for different communication systems. In general a video-conference may consist of three types of data streams:

(1) audio;
(2) video; and
(3) control data for application sharing.

The bandwidth requirements of a video-conference might vary depending on the participants and over time. Several categories of video-conferences may be distinguished, e.g.

- small;
- medium; and
- large

depending on the number of participants and on the required control mechanisms.

With respect to the interaction of ATM and the Internet, for all kinds of video-conference scenarios, there might be a need for all types of interaction between the ATM and Internet worlds:

- *Internet over ATM*: Internet video-conferences will use ATM network resources for connecting each other especially if the QoS features of ATM can be utilised to improve the quality of the video-conference.
- *Mixture of Internet and ATM*: If we assume established systems for Internet and native ATM video-conferencing, there will be a need for having video-conferences not only with people using the same system, but with people using other systems as well. Thus, at least in the long-term, this would imply a need for a kind of peer-to-peer interaction between ATM and Internet.
- *ATM over Internet*: Even the case where an ATM video-conference participant might use an Internet path to send their data to another ATM participant who is not reachable via an ATM connection is imaginable. However,

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · 49–57

this will be a rather rare event, so that this kind of interaction, at least from the perspective of a video-conference, seems to play a minor role.

### Video on demand

For this application we consider a scenario of a network of possibly widely distributed video servers which offer videos to customers owning either a special set-top-box dedicated to this purpose (or an according component integrated into the TV) or a usual multimedia-enhanced computer.

ATM has been proposed and used in many field trials as the backbone network, while usually the cable network or the asymmetric digital subscriber line (ADSL) technology were used for the last-mile distribution of the videos to the consumer. For ATM networks, originally ATM adaptation layer 1 (AAL1) or AAL2 was provided for such data traffic, but since workstations usually favour AAL5, the ATM forum standardised "MPEG-2 over ATM" to use AAL5 as adaptation layer and available bit rate (ABR) as service category.

For some time the Internet has not played a very important role in the VoD arena. But recently it gained more interest, and as more interactivity comes into play, and hence the more computer-based approaches are considered, the Internet becomes more important. An Internet-based VoD system seems, however, only possible in an IntServ environment, in conjunction with a re-engineering of the major elements of the Internet, giving them much more bandwidth capacity.

With respect to VoD services, the following interaction scenario for Internet and ATM seems to be probable: besides the use of ATM as link layer in the Internet, there may be further interactions. In particular, in addition to pure Internet VoD servers there may also be an ATM VoD server network and Internet clients that would like to access the VoD services. While this is a service provider/service user relation on the application level, it would mean a peer-to-peer relation on the networking level. While, as we will see later, this is difficult to achieve, it might pay off for the ATM VoD service provider to allow Internet clients access to its videos and therefore to extend its customer basis by a potentially vast number.

### Other application areas

The application types discussed above, such as video-conferencing or video retrieval, are also useful in scenarios other than with "desktop computers" in homes or offices. For example, in manufacturing environments, video surveillance of certain factory parts can be very desirable. A camera is positioned in, e.g. a prototype manufacturing area, for the responsible engineer to watch the video from his office. If any problems occur, the engineer would like to communicate with someone supervising the plant area by video-conference. In the offices, typically Internet mechanisms are used, for the factory it can be desirable to use (wireless) ATM. Therefore, an interaction between Internet and ATM seems to be very useful in such a scenario.

## Existing interaction approaches

Over recent years, several approaches have been developed for the interworking of ATM and legacy networks without QoS support, e.g. Internet engineering task force (IETF's) classical IP over ATM (Laubach, 1994) with its extensions for multicasting, multicast address resolution server (MARS) (Armitage, 1996), and short-cuts, next hop resolution protocol (NHRP) (Luciani *et al.*, 1998), and ATM Forum's LAN (local area network) emulation (LANE) (ATM_Forum, 1995) and multiprotocol over ATM (MPOA) (ATM_Forum, 1997a). IP switching (Newman *et al.*, 1996) and similar solutions can be seen as more revolutionary approaches, which try to identify data flows and build up virtual circuits (VCs) for them if they seem to be long-lived. The signalling protocols that build up the VCs are especially tuned for this kind of purpose and are no longer the original ATM signalling protocols. So IP switching might be viewed not as an interaction approach with ATM, but as a competing approach to ATM since essentially only the switching hardware of ATM is being used.

All these approaches do not support data flows requiring a predictable QoS since they were designed for asynchronous data communication only. Further approaches try to integrate IntServ and ATM's QoS architecture by extending these existing approaches for

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · 49–57

asynchronous data communication. In this paper we reconsider whether this is sufficient.

## Comparison of the ATM and IntServ QoS architectures

The QoS architectures of ATM and the Internet have very different capabilities and characteristics with regard to the signalling and the QoS models. These discrepancies must be overcome for the interworking between ATM's QoS architecture and the Internet's IntServ QoS architecture. The most salient differences between the QoS models, i.e. the ATM Traffic Management 4.0 (ATM_Forum, 1996b) and the integrated services (IntServ) specifications (Shenker *et al.*, 1997; Wroclawski, 1997) are:

- cell-based versus packet-based traffic parameters and performance specifications;
- the handling of excess traffic: tagging or dropping versus degradation to best-effort;
- different service classes and corresponding traffic and service parameters.

While the traffic characterisation of both QoS models is quite similar, the service definitions differ substantially such that a simple one-to-one mapping is too "semantic-lossy". Thus we think a mapping might have a dynamic or even adaptive n:m character, i.e. the mapping is not fixed, it might adapt itself and one service class of IntServ might, depending on the actual values of the specified parameters, be mapped on different service classes in ATM and vice versa. The QoS parameters must be mapped as well. While the two parameter sets have an intersection, they are neither a subset nor a superset of each other, thus making an easy mapping impossible. As an example for the mismatch between the service classes, IntServ's controlled load (CL) service can be considered, which seems to have no equivalent in ATM. CL is attractive for adaptive applications which have not been sufficiently recognised in ATM's service model so far.

Another problem is the treatment of non-conforming traffic, which in IntServ becomes best-effort traffic while it is at best being tagged with the cell-loss priority (CLP) bit in ATM (but could also be directly discarded depending on policies). Thus, it is treated worse than ATM's best-effort traffic. This means that traffic (which is flowing from an Internet part to an ATM part) that is non-conforming in front of the ATM cloud, i.e. in the Internet, would be treated better than traffic which is not conforming inside the ATM cloud – an obvious mismatch.

There are many differences between IntServ's signalling protocol resource reservation protocol (RSVP) (Braden *et al.*, 1997) and ITU-T's Q2931 which can be viewed as the basis for all ATM signalling protocols. Prominent examples for the differences are:

*Heterogeneous versus homogeneous QoS*: While ATM only allows for homogeneous reservations, RSVP allows heterogeneity, first for different QoS levels for receivers, and second, for simultaneous support of QoS and best-effort receivers. This mismatch in the semantics of RSVP and Q2931 is a major obstacle to simple solutions for the mapping of the two.

*Dynamic versus static QoS*: RSVP supports a dynamic QoS, i.e. the possibility to change a reservation during its lifetime. ATM's signalling protocols, however, have so far been providing only static QoS (QoS renegotiations are currently under discussion as possible future extensions of ATM signalling protocols).

*Receiver- versus sender-orientation*: The different designs with regard to the initiation of a QoS reservation reflect the different attitudes regarding centralised versus distributed management, and also that the IntServ architecture had large group communication in mind, while the ATM model rather catered for individual and smaller group communications.

*Hard state versus soft-state*: The soft-state of RSVP leads to a robust behaviour of the protocol in case of link failures, whereas ATM's hard state needs explicit handling of such situations. However, hard state allows for a more accurate and reliable QoS provision than soft state, especially if no link failure or similar situations occur.

*Resource reservation independent or integrated with setup/routing*: The separation of RSVP from routing leads to asynchronicity of reservation and flow set-up; it also simplifies the support of dynamic QoS. However, QoS routing is much more difficult to achieve in such an environment than with an integrated approach such as followed by ATM; this might become an important issue in future (private network-node

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · 49–57

interface ((P-NNI) (ATM_Forum, 1996a) already supports a form of QoS routing).

*Multicast model*: while IP multicast allows for multipoint-to-multipoint communication, ATM only has point-to-multipoint VCs to emulate IP multicast by either meshed VCs or a multicast server. These are both workarounds which can be shown to be sub-optimal for certain scenarios (Talpade, 1996). The proposed solution at the moment is MARS which however does not seem to be scalable enough for some applications envisioned in the Internet like distributed interactive simulations (DIS) with approximately10,000 group members joining and leaving rapidly.

*Transmission of control messages*: while in ATM separate control channels are used for the transmission of control messages of the signalling protocols, RSVP uses best-effort IP to send its messages.

Many of the differences in signalling can be traced back to the roots of the two signalling mechanisms:

(1) RSVP is based on the observations made during the experimental MBone multicasts of the IETF meetings and therefore multicast is seen as very closely related to QoS in the IETF (Braden *et al.*, 1994).
(2) Conversely, Q.2931 is based on the traditional plain old telephone system (POTS) signalling and its successor N-ISDN with its signalling protocol Q.931.

## Interaction patterns

In order to identify the important interaction patterns which should be supported by interaction models, the potentially possible patterns should be studied. We can derive three such interaction approaches for heterogeneous ATM/Internet networks based on topological observations (see Figure 2):

(1) *IP × ATM × IP*: Communication between IP end-systems with an ATM subnet on the transmission path, e.g. IP-E1 wants to send data to IP-E2 (IP-E1ØIP-N2Ø ED1ØATM-N1Ø ED2Ø IP-N2ØIP-E2), symbolized by communication pattern 1 (CP1) in Figure 2. An example for this is an Internet video-conference, which utilizes ATM's QoS abilities when transmitting over an ATM subnet.
(2) *ATM × IP*: Communication between an ATM end-system and an IP end-system, e.g. IP-E1 would like to send data to ATM-E1 (IP-E1ØIP-N2ØED1ØATM-N1ØATM-E1), symbolized by CP2 in Figure 2. An example application for this communication pattern could again be a video-conference, which, however, this time consists of some participants with IP-connectivity and some with ATM-connectivity.
(3) *ATM × IP × ATM*: communication between ATM end-systems with an IP subnet on the transmission path, e.g. if ATM-E1 wants to send data to ATM-E3 (ATM-E1ØATM-N1ØED2ØIP-N2ØED3ØATM-N2ØATM-E3), symbolised by CP3 in Figure 2. An example application for this could be the connection of two ATM LANs via the Internet for the purpose of building up a virtual private network. In order not to lose too many of the guarantees given by ATM, it would be favourable to be able to utilise IntServ/RSVP flows for the linking of the two ATM LANs.

Similar topological observations are made in Borden *et al.* (1995), however, it considers only CP1 in more detail. In the same vein are all IETF models for the interaction between IntServ and ATM's QoS architecture based on the support of that pattern. While it is clear that CP1 is one (and probably the most) important interaction pattern, it might be too limiting to constrain on one of the possible communication patterns. However, which of them are really worth being investigated depends on the topology of the future networking infrastructure. We perceive two possible topological scenarios with two variants each for a future IP/ATM network:

(1) ATM in the core of the network surrounded by other network technologies to which users might be connected. This might not be realistic in the Internet, yet it is at least possible for corporate networks. With this scenario, we can restrict the view on patterns CP1 and CP2. The two variants depend on whether ATM will be able to come to the desktop or to other kinds of end systems in general. If ATM will really play a significant role in end systems, e.g. the LAN environment or for residential users, then both patterns, CP1 and CP2, have to be

Internet Research: Electronic Networking Applications and Policy
Volume 9 · Number 1 · 1999 · 49–57

**Figure 2** Possible interactions between Internet and ATM

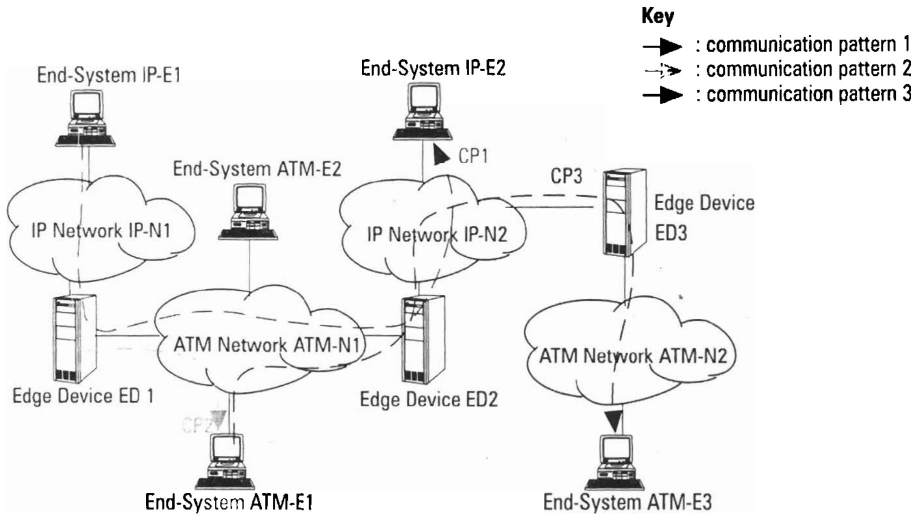


taken into account. If ATM will only be a wide-area network (WAN) solution then a possible interaction model only needs to care about CP1.

(2) The alternative scenario is that ATM is just one of many link layer technologies. In this case all communication patterns might have some importance, even CP3. Again, the variants of this scenario depend on the question of whether ATM will make it to the desktop or not. If ATM will not be solely a WAN solution then all three communication patterns will have to be taken into consideration. Otherwise, i.e. if ATM will be one of many WAN technologies, mainly CP1 (and sometimes also CP3) will have to be supported by the interaction model.

Whether ATM will play an important role on the desktop depends on various factors, e.g. costs, technical suitability, manageability, and the availability of native ATM applications. Another point is that the interaction approaches should also be made dependent on what the purpose of the internetwork is:

- global internetwork;
- private internetwork with centralised administration and control (of network engineering and protocol usage); or
- private internetwork with distributed management by independent organisations but on a scale that is still moderate.

The global case will probably be dominated by Internet technology, owing to its support of heterogeneity. The last two cases might be a niche for ATM – homogeneity may be achieved at least in the backbone, especially in the centralised case. In other than the "desktop computer" scenario different rules and considerations may apply, e.g. for manufacturing purposes using one technology particularly tailored for the appliances in the industrial environment might be important which has then to be connected to general purpose equipment in the supervising offices.

## Interaction models

### Traditional model (straightforward solution)

The traditional model serves situations corresponding to pattern CP1. However, since it is straightforwardly implemented, it does not make clever use of the features provided by ATM. ATM switched virtual circuits (SVCs) or permanent virtual circuits (PVCs) are used as fast bit pipes and the QoS provisioning is done solely by the IntServ architecture – this is essentially classical IP over ATM "abused" for IntServ. While it is inexpensive in terms of invested effort, it ignores all the features provided by ATM and is very expensive in terms of usage of resources. It operates on ATM as if it were a "dumb" point-to-point network or a leased line and does not make any use of the features provided by ATM. Instead it duplicates these functions in the IntServ architecture (in software, which is usually less efficient).

There is much less implementation complexity in the traditional model compared to other approaches. Despite its obvious deficiencies it must be seriously analysed with regard to the performance loss and resource wastage it incurs in comparison to the more sophisticated models. In the traditional model, ATM is viewed as a black box while the other approaches show a tendency to more and more regard the internals of ATM.

**ATM subordination model**

The ATM subordination model is an extension of the traditional model. It tries to make use of the ATM features as far as possible. ATM is still viewed as a subnet providing services for the IntServ architecture. There are two different forms whereby the interaction can be designed. ATM could be aware of the interaction (and be adapted) or ATM remains unaltered and is passively used by IntServ with all its constraints. In the latter case, the IP over ATM signalling would have to be adapted, since the ATM QoS architecture would be regarded as fixed. For the former, there are radical approaches that want to do away with the ATM signalling and install a completely new signalling for ATM especially in support of overlaying an IP network over ATM.

The IETF favours the passive ATM subordination model (Berger *et al.*, 1998) since they view ATM as an important link layer technology, whose QoS capabilities should be utilised by an integrated services Internet. However, the IETF does not consider a more integrated solution of the QoS architectures such as the one we will present in the next subsection. The reason is that most people active in the IETF expect ATM to be solely a WAN solution, and maybe the WAN solution presenting the backbone of a future Internet, but ATM will never make it to the desktop in their view. So a good solution could be to regard RSVP/IP and ATM as complementary techniques, where ATM is at the core, a place where its QoS routing feature is very desirable, and RSVP/IP is at the edges of the network, where its ease of use is much needed. So they should not be considered as opponents, but rather as partners, though still on a provider-user basis.

**Partnership/competitor model**

The partnership/competitor model serves situations where the pattern CP2 applies, i.e. communication between ATM- and IP-connected end-systems on a peer-to-peer basis. It requires an integrated fashion of interworking between ATM and Internet and accepts ATM as a full-blown protocol stack that is able to operate end-to-end, and not solely as a data link technology as in the ATM subordination models. Currently, ATM on the desktop is practically non-existing, thus, if ever, the peer model will play a role in middle-term future or specific (e.g. manufacturing) environments only. If more ATM end systems would come into existence, then this interaction model is probably necessary.

An example on the application level for this model taken by the ATM Forum is the voice and telephony over ATM (VTOA) Phase 2 work (ATM_Forum, 1997b), which tries to approach the interworking between ATM and Internet voice transportation. However, an interworking at the system layer with "asymmetric" end-systems seems a more fundamental answer to the problem, which however depends on the number of applications the two worlds are really sharing (in a peer-to-peer and not in a provider/user relation).

**Internet subordination**

The Internet subordination model serves situations where pattern CP3 is required. This is the case where an IP network acts as a transit network for communicating ATM-connected end-systems without direct connectivity. While this might look absurd today, it could have some relevance in case that there will be a scattered set of small islands of ATM networks. For example, organisations that have ATM LANs might connect them via Internet and might be interested in preserving the ATM QoS as good as possible by using IntServ's QoS. Nevertheless, the Internet subordination model should have exceptional character since it is not really possible to keep the QoS guarantees given by the ATM network over the Internet section of the transmission path, thereby causing an unpredictable QoS provision. From a technical point of view something similar has been developed as "Cells in frames" (Cells_in_Frames_Alliance, 1996), i.e. ATM cells in Ethernet frames, in order to emulate ATM end-to-end. However, what would be needed is "Cells in packets", i.e. ATM cells in IP packets to be able to cross routers.

## Summary and conclusion

We believe that interaction approaches for the QoS architectures of the Internet and ATM are necessary because both worlds will co-exist for a couple of years. Since they tend to increasingly serve the same applications due to the pertaining convergence process of data and telecommunications, they have to interwork with each other to fulfil application demands.

New and, also from an economic perspective, interesting applications, like video-conferencing and video-on-demand services are run, or will be run, in both worlds. It is only natural that a seamless interworking between both worlds is demanded. For example, a video-conference should neither be constrained on Internet-connected participants nor on ATM-connected participants but should allow for mixed video-conferences with participants of both worlds. The ATM subordination model is not able to support such communication scenarios. Solutions for the partnership model do not exist yet, owing to its high complexity. Nevertheless, for real-time data this model has some right to exist, though we are not aware of any approaches in that direction, not even to identify the problems associated with it.

Based on topological considerations and application scenarios we derived the required communication patterns and interaction models for an interworking between ATM and Internet. Which of these will be the prevailing ones, depends on many factors. However, there are technical issues, like the fast introduction of a programming interface to native ATM services and the existence of pure ATM end-systems such as videophones, video-servers, set-top boxes or cameras based on ATM. Conversely, economical and political factors, for example the protection of investments, have to be taken into account.

## References

Armitage, G. (1996), *Support for Multicast over UNI 3.1 based ATM Networks*, RFC 2022.

ATM_Forum (1995), *LAN Emulation over ATM Version 1.0 Specification*, ATM Forum.

ATM_Forum (1996a), *Private Network-Node Interface (PNNI) Signalling Specification*, ATM Forum.

ATM_Forum (1996b), *Traffic Management (TM) Specification 4.0. ATM Forum.*

ATM_Forum (1997a), *MPOA Version 1.0. ATM Forum.*

ATM_Forum (1997b), *Voice and Telephony Over ATM (VTOA) to the Desktop*, ATM Forum.

Berger, L., Crawley, E., Berson, S., Baker, F., Borden, M. and Krawczyk, J. (1998), "A framework for integrated services with RSVP over ATM", Internet draft (forthcoming).

Borden, R., Crawley, E., Davie, B. and Batsell, S. (1995), "Integration of real-time services in an IP-ATM network Architecture", RFC 1821.

Braden, R., Clark, D. and Shenker, S. (1994), "Integrated services in the Internet architecture: an overview", RFC 1633.

Braden, R., Zhang, L., Berson, S., Herzog, S. and Jamin, S. (1997), "Resource ReSerVation protocol (RSVP) – Version 1 functional specification", RFC 2205 – proposed standard.

Cells_in_Frames_Alliance (1996), *ATM over Anything – Cells in Frames. Cells in Frames Alliance*, White Paper.

Laubach, M. (1994), "Classical IP and ARP over ATM", RFC 1577 – proposed standard.

Luciani, J., Katz, D., Piscitello, D. and Cole, B. (1998), "NBMA Next Hop Resolution Protocol (NHRP)", RFC 2332.

Newman, P., Minshall, G. and Lyon, T. (1996), "IP switching: ATM under IP", *Proceedings of IEEE Infocom.*

Schill, A., Hess, R., Kümmel, S., Hege, D. and Lieb, H. (1997), *ATM-Netze in der Praxis*, Addison-Wesley, Reading, MA and Wokingham.

Shenker, S., Partridge, C. and Guerin, R. (1997), "Specification of guaranteed service", RFC 2212 – proposed standard.

Steinmetz, R. and Nahrstedt, K. (1995), *Multimedia: Computing, Communication and Applications*, Prentice-Hall, Englewood Cliffs, NJ.

Talpade, R. and Ammar, M. (1996), "VC consumption in the MARS architecture for IP multicast over ATM", position paper at the IP/ATM Workshop at the University of Washington, DC.

Wolf, L., Griwodz, C. and Steinmetz, R. (1997), "Multimedia communication", *Proceedings of the IEEE*, invited paper in Special Issue on Global Information Infrastructure, Vol. 85 No. 12, pp. 1915-33.

Wroclawski, J. (1997), *Specification of the Controlled-Load Network Element Service*, RFC 2211, proposed standard.