Interaction Approaches for Internet and ATM Quality of Service Architectures

J. Schmitt[1], L. Wolf[1], C. Siebel[2], Y.-O. Lorcy[2, 3], R. Steinmetz[1,4]

[1] KOM, Darmstadt University of Technology, Merckstr. 25, D-64283 Darmstadt, Germany
Tel.: +49-6151-166150, Fax: +49-6151-166152
{Jens.Schmitt, Lars.Wolf, Ralf.Steinmetz}@kom.tu-darmstadt.de

[2] Deutsche Telekom AG, Technologiezentrum Darmstadt,Germany
siebel@tzd.telekom.de

[3] ENSEEIHT, Toulouse, France
lorcy@campus.fz-telekom.de

[4] GMD IPSI, German National Research Center for Information Technology,
Dolivostr. 15 • D-64293 Darmstadt • Germany

Abstract

Internet and ATM possess Quality of Service (QoS) architectures which shall allow them to integrate services of data- and telecommunications formerly performed by separate infrastructures. We believe that none of them will be able to oust the other. That means both will exist for at least the middle-term future. Therefore, an interaction between both is necessary, especially in the field of distributed multimedia applications where both worlds 'meet' first. These applications require a certain QoS to perform gracefully. Hence, in particular the interaction between Internet and ATM QoS architectures is important. In this paper, interaction approaches for the QoS architectures developed for the Internet and for ATM are discussed. We do not restrict on common approaches but also derive more unconventional models by regarding the possible communication patterns based on different topological variants for heterogeneous IP-ATM networks.

Keywords: Signalling in Heterogeneous Networks, IP/ATM Networks, Integrated Services, QoS.

1. Introduction

ATM and Internet are the two major players in communications at the moment. At first glance they are counterparts, at least in some respects, with different strengths and weaknesses, and need to compete with each other – the Internet defending its position as a global internetwork and ATM trying to become one. That ATM will really ever become a global internetwork is often doubted (by data communications people). However, ATM will play a certain role, even if it does not become the successor of the Internet. We believe that both will exist for at least the middle-term future. Thus, they must be able to interact with each other.

The interaction of these two worlds is particularly desired in the relatively new field of distributed multimedia applications where both worlds meet first. These applications depend on the provisioning of QoS mechanisms by the communication system. This need has been observed in both worlds, and both have developed independently of each other QoS architectures that shall be able to provide integrated services. So one of the most important aspects of the interaction between both worlds is the seamless interworking between the QoS architecture of ATM and the Internet integrated services architecture (IntServ). That means to enable 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, i.e. providing a homogeneous service over a heterogeneous network.

In this paper, we discuss interaction approaches for Internet and ATM QoS architectures. In the next two sections we give a brief overview of existing ATM / Internet interaction models and compare their QoS architectures. Then we examine the communication patterns in heterogeneous IP-ATM networks in Section 4. Based on that, several interaction models for the QoS architectures of the Internet and ATM are derived in Section 5. Finally we conclude the paper.

2. Existing Interaction Approaches

Over the last few years, several approaches have been developed for the interworking of ATM and legacy networks *without* QoS support, e.g., IETF's Classical IP over ATM [15] with its extensions for multicasting, MARS (Multicast Address Resolution Server) [9], and short-cuts, NHRP (Next Hop Resolution Protocol) [16], and ATM

Forum's LAN Emulation (LANE)[3] and Multi-Protocol over ATM (MPOA)[4]. IP switching [17] and similar solutions can be seen as more revolutionary approaches, which try to identify data flows and build up VCs (Virtual Circuits) 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 asynchronous data communication. In this paper we reconsider whether this is sufficient.

A different approach has been taken by AREQUIPA [1] where ATM endsystems try to detect whether ATM connectivity is existing up to the communication peer. If this is the case, an ATM connection is set up for further communication. This approach is only viable if end-to-end ATM connectivity exists.

3. Comparison of ATM and Internet QoS Architectures

In Figure 1 the most important components of both QoS architectures and their approximate semantic mapping onto each other is illustrated. We base our comparison on the latest (at the time of writing) specifications of the ATM Forum ([6], [7]) and the proposed standard RFCs of the IETF ([12], [19], [21]).

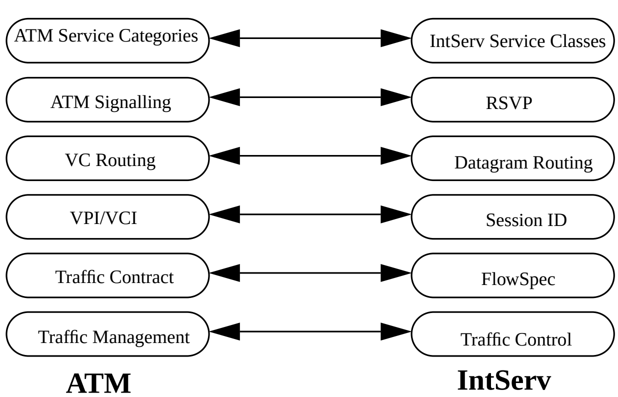

Figure 1: Mapping between ATM and IntServ components (based on [2]).

Both architectures have very different capabilities and characteristics with regard to the signalling (the QoS procedures) and the QoS models (the QoS declaration or interface), but these discrepancies have to be overcome when interworking between ATM's QoS architecture and the IntServ architecture. We look at the QoS models and procedures separately.

3.1 QoS Models

The most salient differences between the QoS models, i.e. the ATM TM 4.0 ([6]) and the Integrated Services (IntServ) specifications ([19], [21]), are:

- packet-based vs. cell-based traffic parameters and performance specifications,
- the handling of excess traffic (policing): tagging or dropping vs. degradation to best-effort,
- and of course different service classes and corresponding traffic and service parameters.

While the traffic characterization of both QoS models is quite similar (token bucket rate+token bucket size/depth vs. PCR/SCR+MBS/MCR*), the service definitions differ substantially, such that a one-to-one mapping seems to be 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 the other way around.

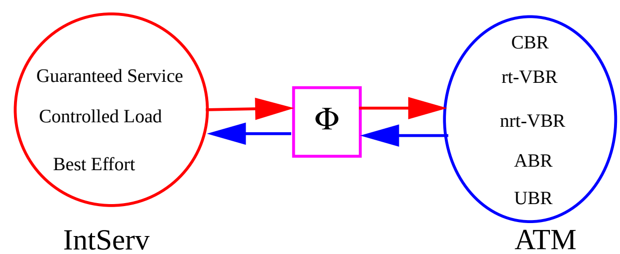

Figure 2: Mapping of Service Classes/Categories between ATM and IntServ.

That the service classes of ATM do not fit together very nicely, can be seen with IntServ's Controlled Load (CL) service which seems to have no equivalent in ATM. That is due to the fact that the applications for which CL is attractive are adaptive applications (also supported by the dynamicity of IntServ's QoS provision, something considered but not yet implemented in ATM signalling), while in ATM's service model adaptive applications seem to be not sufficiently recognized.

Although IntServ's Guaranteed Service (GS) maps easier onto ATM's QoS model there is still no one-to-one mapping possible. While for token buckets with a small depth CBR (Constant Bit Rate) seems to be the right choice as a service category, for larger values of the token bucket depth this would lead to wastage of bandwidth. Therefore to allow for a variable source not to waste too much band-

*. PCR = Peak Cell Rate, SCR = Sustainable Cell Rate, MBS = Maximum Burst Size, MCR = Minimum Cell Rate.

width, GS should rather be mapped onto rt-VBR (real-time Variable Bit Rate) if the ratio of token bucket depth and token bucket rate exceeds a certain threshold value.

Besides the mapping of the service classes also the QoS parameters have to be mapped. While the two parameter sets certainly have an intersection, they are neither a subset nor a superset of each other, thus making an easy mapping impossible. A practical problem in this area is that the parameters are specified in different units: bytes vs. cells, and thus must be translated into each other taking into account the encapsulation and padding overhead.

Another problem is the treatment of non-conforming traffic, which in IntServ becomes best-effort traffic while it is at best being tagged (CLP (Cell Loss Priority) bit = 1) in ATM (but could also be directly discarded depending on policies). Therefore, this traffic is treated worse than ATM's best-effort traffic (UBR (Unspecified Bit Rate) or ABR (Available Bit Rate)). This means that traffic that is non-conforming in front of the ATM cloud would be treated better than traffic which does not conform inside the ATM cloud – an obvious mismatch.

A further QoS model mapping problem is caused by the fact that the traffic specification given by the end-systems might not represent the actually generated traffic inside the network, although applications adhered to the traffic contract. This is due to the fact that schedulers can only achieve an approximated fluid model. Therefore reservations based on the traffic description given by the application might lead to situations where the policing functions of the ATM network might throw away data, which was conformant when entering the IP network but non-conformant when entering the ATM network. This is however not the application's fault and hence it should not be punished for it.

Another problem arises from the IntServ concept OPWA (One-Pass with Advertising [18]). This uses a so-called AdSpec to give the receivers an idea of which QoS they could expect from the network. So the question is how to advertise the ATM cloud, which might consist of a very complex ATM network that from IntServ's perspective is however seen as one single hop.

3.2 The QoS Procedures

While it is not easy to map the QoS models of the Internet and ATM, it is even more difficult to map their QoS procedures onto each other. This is due to the fact that they are built upon very different paradigms. While the signalling protocols of ATM are still based on the call paradigm used for telephony, the IETF viewed the support of a flexible and possibly large-scale multicast facility as a fundamental requirement. The most prominent differences between RSVP (Resource reSerVation Protocol [12]), which can be viewed as the Internet's signalling protocol, and ITU-T's Q.2931, on which all ATM signalling protocols are based, are discussed in the following.

Heterogeneous vs. Homogeneous QoS

While ATM only allows for homogeneous reservations, RSVP allows heterogeneity firstly for different QoS levels for receivers and secondly for simultaneous support of QoS and best-effort receivers. This mismatch in the semantics of RSVP and Q.2931 is a major obstacle to simple solutions for the mapping of the two.

Dynamic vs. 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- vs. Sender-Orientation

The different designs with regard to the initiation of a QoS reservation reflect the different attitudes regarding centralized vs. 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 vs. Soft-State

The discrepancies between the ATM QoS architecture and the IntServ architecture in how the state in intermediate systems is realized is another major obstacle to the interworking of both worlds since it leads to very different characteristics of the two QoS architectures. The soft-state of RSVP leads to a robust behavior of the protocol in case of link failures, whereas ATM's hard state is rather fragile to such situations. Yet, on the other hand hard state allows for a much more accurate and reliable QoS provision since RSVP can principally never guarantee that the QoS that was 'promised' by the network to the application can be hold up for the whole duration of the session even 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 setup, and furtherly enables an independent evolution of routing and resource reservation mechanisms. Another advantage is the easy support of dynamic QoS. However a possibly major disadvantage in future may be that QoS routing is much more difficult to achieve than with ATM's integrated connection setup/ resource reservation mechanism (P-NNI [5] already supports a form of QoS routing).

Multicast Model

A further issue is the mapping of the IP multicast model on the signalling facilities in ATM for multi-party calls. 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 suboptimal for certain scenarios [20]. 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 DIS (Distributed Interactive Simulations), with around 10,000 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.

Problematic is that both architectures are still changing quite rapidly, parameters are added and abandoned, new service categories are introduced and earlier ones are abandoned, etc. However, on the other hand these changes could also alleviate the mapping .

It shall be emphasized once more that many of the differences in signalling can be traced back to the roots of the two signalling mechanisms: 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 [11]. Q.2931 on the other hand is based on the traditional POTS (Plain Old Telephone System) signalling and its successor N-ISDN with its signalling protocol Q.931.

Among the capabilities of RSVP which are not supported by ATM are the most important: dynamic and heterogeneous QoS, and sharing and aggregation mechanisms for scalability within a session. These are characteristics which are especially useful in the multicast case.

Capabilities of ATM which are not being (well) supported by IntServ: the accurateness of QoS over the whole lifetime of a connection, the bidirectionality of unicast connections, and the scalability with regard to the number of sessions.

Besides the integration of the QoS models and signalling procedures a practical, realizable solution needs to integrate further components as the security frameworks and the pricing/billing/accounting or policy control framework of both worlds. However these framework components have neither in ATM nor in the Internet reached a consensus, and have so far been postponed so that the interaction between those not yet existing components is difficult to anticipate.

In the next section we will enumerate all possible interaction patterns between ATM and Internet in order to identify the most important ones, which should be supported by an interaction model.

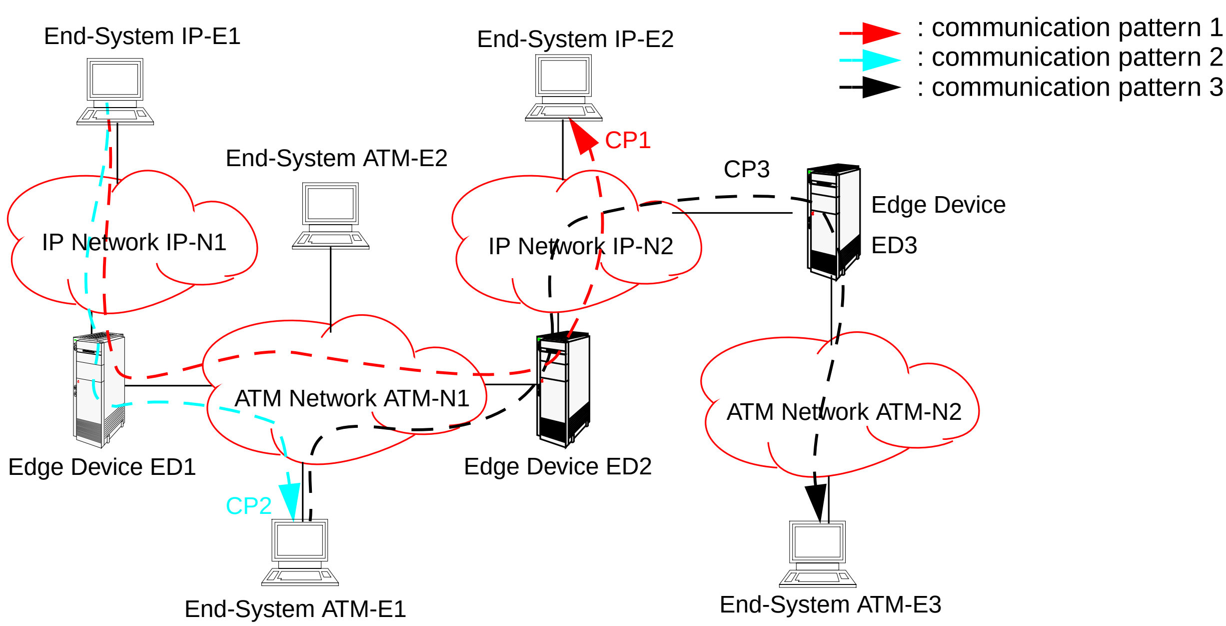

Figure 3: Possible Interactions between Internet and ATM.

4. Enumeration of Interaction Patterns

Based on topological observations we can derive three potentially possible interaction approaches (see Figure 3):

1. 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-E1→IP-N2→ED1→ATM-N1→ED2→ IP-N2→IP-E2), symbolized by communication pattern 1 (CP1) in Figure 3. An example for this is an Internet videoconference, which would certainly be happy to utilize ATM’s QoS abilities when transmitting over an ATM subnet.

2. Communication between an ATM end-system and an IP end-system, e.g., IP-E1 would like to send data to ATM-E1 (IP-E1→IP-N2→ED1→ATM-N1→ATM-E1), symbolized by CP2 in Figure 3. An example application for this communication pattern could again be a videoconference, which however this time consists of some participants with IP-connectivity and some with ATM-connectivity.
3. 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-E1→ATM-N1→ED2→IP-N2→ED3→ATM-N2→ATM-E3), symbolized by CP3 in Figure 3. 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 much of the guarantees given by ATM, it would be favorable to be able to utilize IntServ/RSVP flows for the linking of the two ATM LANs.

Although in RFC 1821 [10] similar topological observations are made, 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. We do not think that it is necessarily sufficient to constrain on one of the possible communication patterns. However which of the communication patterns are really worthwhile being investigated is in our opinion an open issue which 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 not. If ATM will really play a significant role in the LAN environment or for residential users, then both patterns, CP1 and CP2, have to be taken into account. If ATM will only be a 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 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. The latter depends on the fast introduction of a standardized ATM API (Application Programming Interface) and how it is accepted by application programmers which are used to TCP/IP applications and APIs, but have mostly no experience with native ATM-mode applications. Also, TCP+UDP-based applications have the advantage of providing support for heterogeneity. Besides the slow deployment of an ATM API, another argument often raised against the vision of ATM on the desktop is that there is a large gap between the services demanded by applications and the services provided by ATM. Here a mapping of cell-level guarantees and services to something more meaningful for applications like packets or frames would be needed. There seems to be at least one layer of abstraction missing - may be a perfect gap for IntServ to fill out and continue IP's dominance on the desktop.

Another point, only seldom considered, is that the interaction approaches should also be made dependent on what the purpose of the internetwork is: global internetwork, private internetwork with centralized administration and control (of network engineering and protocol usage), or private internetwork with distributed management by independent organizations but on a scale that is still moderate. The global case will probably be dominated by Internet technology due 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 centralized case.

5. Interaction Models

5.1 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 SVCs (Switched Virtual Circuit) or PVCs (Permanent Virtual Circuit) 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, which is of course principally possible and inexpensive in terms of invested effort, but 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 makes any use of the features provided by ATM like the VC model (which allows for a presorting of flows on the data link layer), bandwidth management, or traffic management (traffic control and scheduling in hardware). Instead it duplicates these functions in the IntServ architecture (in software, which is usually less efficient).

This represents the most hostile approach to ATM and neglects the capabilities of ATM signalling as far as possible and therefore neglects many of the good features of ATM, but on the other hand avoids its complexities as well.

There is much less implementation complexity in the traditional model compared to other approaches. Despite of its obvious deficiencies it must be seriously analyzed 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.

5.2 ATM Subordination Model

The ATM subordination model is an extension of the traditional model but 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 how 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 which 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. IP Switching approaches might be regarded as examples of this approach if they are extended to use RSVP as an explicit mechanism to set up VCs. This approach is also called ’ATM under IP’ [17], showing literally the active subordination of ATM to the Internet. However most approaches take the more pragmatic passive subordination of ATM which regards ATM signalling as fixed. Furthermore, the extensions to ATM signalling and the changes to the IntServ architecture may be regarded as steps towards each other.

The IETF favors the passive ATM subordination model [14] since they view ATM as an important link layer technology, whose QoS capabilities should be utilized by an Integrated Services Internet. However, the IETF does not consider a more integrated solution of the QoS architectures as we will present one in the next subsection. The reason is that most people active in the IETF expect ATM to be solely a WAN solution, and may be *the* WAN solution presenting the backbone of a future Internet, however 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 desired, and RSVP/IP is at the edges of the network, where its ease of use is well desired. So they should not be considered as opponents but rather as partners, though still on a provider-user basis.

5.3 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. Hence, it may also be called peer model or integrated model since it requires an integrated fashion of interworking between ATM and Internet. However, this model will only become interesting if ATM is successful on end-systems. If that happens the kind of interaction provided by the peer model is probably necessary to be supported, hence the peer model is important, as it 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.

The peer model introduces the need of a much tighter integration between ATM and Internet, the Internet is no longer just using ATM but they really need to interwork. This possibly leads to the fact that the QoS of ATM can no longer just be ordered through ATM’s QoS interface, but the traffic management of both world’s has to be integrated on a lower level. Since ATM’s QoS architecture seems more powerful than the IntServ architecture the fundamental problem is the mapping of ATM’s QoS on IntServ’s QoS, e.g., how to simulate ATM’s accurateness and QoS reliability on IntServ’s rather crude and unreliable QoS provision.

There are a number of people (many of which are telecommunications’ people) who think that ’the time has come to view ATM not solely as a link layer’: ATM is about to have a standardized API, it offers routing, addressing and session services. These are all elements which distinguish it from traditional link layer technologies and allow it to compete with IP as an end-to-end solution. However at the moment ATM on the desktop is practically non-existing and therefore, if ever, the peer model will play a role in middle-term future, when there are more ATM end-systems available as for example ATM-connected workstations, video-servers, set-top boxes, or tele/videophones.

An example on the application level for the peer model taken by the ATM Forum at the moment is the VTOA (Voice and Telephony over ATM) Phase 2 work [8], 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).

5.4 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. At first glance this might look absurd today, but it could have some relevance in case that there will be a scattered set of small islands of ATM networks. For example, organizations 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.

The accurateness of ATM's QoS can only be approximated by IntServ's QoS but never be guaranteed, since the mapping from ATM QoS service categories and traffic and QoS parameters into IntServ terms seems very problematic. However this is a possible solution if QoS communication between unconnected ATM networks is needed and of course still better than delivering ATM traffic over the best-effort Internet.

From a technical point of view something similar has been developed by the Cornell University and the Connectware Inc.: Cells in Frames [13], 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.

6. 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 serve more and more the same applications due to the pertaining convergence process of data and telecommunications, they have to interwork with each other to fulfill application demands.

New and, also from an economic perspective, interesting applications like videoconferencing 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 videoconference should neither be constrained on Internet-connected participants nor on ATM-connected participants but should allow for mixed videoconferences with participants of both worlds. The ATM subordination model is not able to support such communication scenarios. Solutions for the partnership model are not existing yet due 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. On the one hand, there are technical issues, like the fast introduction of an API 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. On the other hand, also economical and political factors, for example the protection of investments, have to be taken into account.

References

[1] W. Almesberger, J.Y. Le Boudec, P. Oechslin: *Application REQuested IP over ATM (AREQUIPA)*, RFC 2170, July 1997.

[2] A.Alles: *ATM Internetworking*, White Paper, Cisco Systems Inc., May 1995.

[3] The ATM Forum: *LAN Emulation over ATM Version 1.0 Specification*, January 1995.

[4] ATM Forum: *MPOA Version 1.0*, July 1997.

[5] ATM Forum: *Private Network-Node Interface (PNNI) Signalling Specification*, February 1996.

[6] ATM Forum: *Traffic Management (TM) Specification 4.0*, April 1996.

[7] ATM Forum: *User-Network Interface Signalling (SIG) Specification Version 4.0*, April 1996.

[8] ATM Forum: *Voice and Telephony Over ATM (VTOA) to the Desktop*, May 1997.

[9] G.Armitage: *Support for Multicast over UNI 3.1 based ATM Network*s, RFC 2022, November 1996.

[10] M.Borden, E.Crawley, B.Davie, S.Batsell: *Integration of Real-Time Services in an IP-ATM Network Architecture*, RFC 1821, August 1995.

[11] R.Braden, D.Clark, S.Shenker: *Integrated Services in the Internet Architecture*, RFC 1633, June 1994.

[12] R.Braden, L.Zhang, S.Berson, S.Herzog, S.Jamin: *Resource Reservation Protocol (RSVP) - Version 1 Functional Specification*, RFC 2205, September 1997.

[13] Cells in Frames Alliance: *ATM over Anything - Cells in Frames*, Cornell University, White Paper, 1996.

[14] E.Crawley, L.Berger, S.Berson, F.Baker, M.Borden, J.Krawczyk: *A Framework for Integrated Services and RSVP over ATM*, Internet Draft, November 1997.

[15] M.Laubach: *Classical IP and ARP over ATM*, RFC 1577, January 1994.

[16] J.Luciani, D.Katz, D.Piscitello, B. Cole: *NBMA Next Hop Resolution Protocol (NHRP)*, Internet-Draft, December 1997.

[17] P.Newman, G.Minshall, T.Lyon: *IP Switching: ATM under IP*, Proc. of IEEE INFOCOM'96.

[18] S.Shenker, L.Breslau: *Two Issues in Reservation Establishment*, Proc. ACM SIGCOMM '95, August 1995.

[19] S.Shenker, C.Partridge, R.Guerin: *Specification of Guaranteed Quality of Service*, RFC 2211, September 1997.

[20] R.Talpade, M.Ammar: *VC Consumption in the MARS architecture for IP Multicast over ATM*, Position Paper at the IP/ATM Workshop at the University of Washington, November 1996

[21] J.Wroczlawski: *Specification of the Controlled-Load Network Element Service*, RFC 2212, September 1997.