TOWARDS AN INTEGRATED SOLUTION FOR MULTIMEDIA COMMUNICATIONS

Laurent Mathy[1]
Guy Leduc[2]
André Danthine
Olivier Bonaventure

Université de Liège
Service de Systèmes et Automatique
Institut d'Electricité Montefiore, B28,
B-4000 Liège 1, Belgium
mathy@cs.ubc.ca,
{leduc, bonavent}@montefiore.ulg.ac.be,
danthine@vm1.ulg.ac.be

This work has been partially supported by the Commission of the European Union under the RACE M1005 project ACCOPI: "Access Control and Copyright Protection for Images".

1. Introduction

During the last few years, we have been witnessing tremendous changes in the communication environment. Those changes are not only due to the availability of higher data rates in the communication networks but also to the appearance of new application types such as multimedia applications.

After a brief introduction to multimedia communications, we will present the research activities in multimedia we have undertaken during the last few years. They focus on two aspects of multimedia communications: a new semantics for expressing Quality of Service (QoS), and a support for group communications.

2. What Is Multimedia Communications?

As we are entering the "telecommunication century", people wish to communicate as easily and naturally as possible and therefore multimedia is emerging as a powerful communication tool. Indeed, multimedia allows one to communicate by using jointly different information types (i.e. different media) such as video, audio, text, still images, and so on.

Moreover, beside the simultaneous use of several media, "natural interpersonal" communications also benefit from the possibilities offered by group communications, that is communications involving more than two peers. Such multimedia group communications among computers allow geographically distant people to communicate as if they were in the same meeting room or around the same dinner table.

---

[1] At the University of British Columbia (UBC), Vancouver, Canada, since November 95.

[2] Maître de recherches F.N.R.S (Senior Research Associate)

2.1. Heterogeneity in Multimedia Communications

Each of the different media making up a multimedia message having its own characteristics, it has also its own requirements. Those requirements may differ greatly from one medium to another. For instance, some media are said to be continuous[3] (e.g. video, audio) while others are said to be discrete (e.g. text, still images). While continuous media are characterised by stringent communication performance (e.g. time constraints), discrete media are often characterised by the need to avoid any error during communication.

From such a diversity comes the need, for a multimedia application, to be able to describe the requirements of each of the media it is using to communicate. This is done through the use of Quality of Service (QoS) parameters. Those QoS parameters allow for a description of the traffic characteristics required for the transfer of a given medium. The following performance QoS parameters allow for a complete traffic characterisation:

- The throughput, indicating the data exchange rate.
- The transit delay, indicating the elapsed time between the sending and the reception (by another user) of a piece of information.
- The delay jitter, indicating the maximum variation of the transit delay. In other words, the delay jitter indicates the irregularity in the pieces of information delivery pattern.
- The error rate, indicating a measure of the degradation suffered by the information during transmission.

2.2. Integrated Multimedia Communications

In the past, to deal with that heterogeneity in the media, several independent specialised networks were developed to transport the different specific information types. Among those specialised networks, we find [3]:

1. The Public Switched Telephone Networks (PSTN) designed for classical two-way voice conversation.
2. The Packet Switched Data Networks (PSDN) designed to transport computer data in the public domain. Such networks may be based, for instance, on the X.25 protocol suite.
3. The legacy LANs such as 802.3, 802.5 and FDDI designed to transport computer data at high speed and low cost in the private domain.
4. The Community Antenna TeleVision networks (CATV) designed for the distribution of television signals. With voice and image, it is a multimedia based network, but as the voice and the image signals are integrated they are considered as a unique medium. Most of these networks do not allow any interactivity.

A straightforward and apparently appealing way to achieve multimedia communications, is to transport each specific medium of a multimedia communication on its specialised network type as depicted on figure 1.

However, such a solution suffers from several major drawbacks, among which we find:

- The use of several networks multiplies the costs of network equipment, network maintenance, network operation, and so on.
- The proliferation of customer devices. This not only increases customers' costs but also the number of addresses through which a customer is reached (e.g. phone number, host address, etc.). It also increases the number of equipments on the desk.
- Although a multimedia message has to be interpreted as a whole, its components (the different media) are here transferred independently of one another on different networks. Any control to be exercised on the multimedia communication has thus to be done by the customers.
- Different media cannot share resources, which induces very poor resource usage.

---

[3] Here, the term "continuous" refers to the user's impression of the data, not necessarily to its internal representation.

- Each specialised network has been built to meet the requirements of a given medium, thus offering inflexible performance.

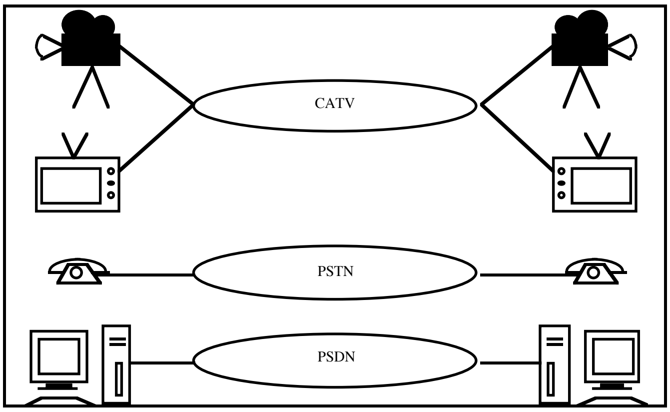

Figure 1: Multimedia Communications using Specialised Networks.

Fortunately, due to the advances in digital network technology (e.g. ATM), computers' processing power and digitalisation techniques, it is now possible to envisage to build a single network able to transport the different information types at the same time. This single network is said to integrate the services (i.e. the different information types). Such a network not only solves the problems discussed above for the previous solution, but also offers a better adaptation to changing media requirements (due to changes in coding/compression schemes).

Moreover, as in an integrated solution, information is digitized, the way to store that information is independent of its type. This not only avoids the simultaneous use of special purpose storage devices (e.g. video recorder with TV set, voice recorder with telephone set, data storage device with workstation, etc.), but this also allows the storage of multimedia information as a single document in storage devices associated with the multimedia workstation.

Figure 2 depicts the new solution using an integrated network allowing integrated multimedia communications.

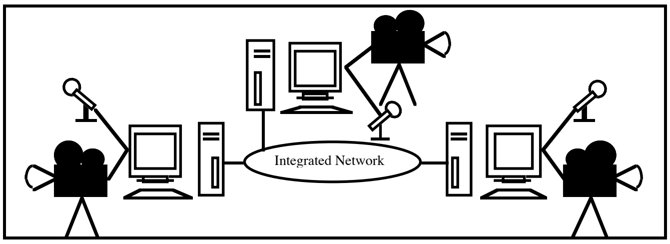

Figure 2: Multimedia Communications in an Integrated Network.

3. Issues in Integrated Multimedia Communications

A research axis of our department is oriented toward the achievement of an integrated solution for multimedia communications in a corporate environment.

In previous sections, we have seen that two important issues in multimedia communications are the QoS and the group communication supports. It is along these two directions that we have concentrated our efforts.

3.1. Quality of Service support

The QoS support is needed to meet the performance requirements of multimedia communications.

Although the QoS concept has been introduced in the OSI Reference Model, the OSI transport service offered by TP4 provides what is called best-effort QoS. This is called best-effort because nothing happens if the performance, corresponding to the QoS values selected by the transport service users, are not achieved by the service provider. Therefore, there is here no strong relation between the QoS "support" offered by the transport service and the performance it actually achieves. The QoS values specified by the transport service users are thus appearing as nothing but a wish. This best-effort semantics (or even no-effort at all) is a common characteristic of the transport services offered by transport protocols that were mainly designed for file transfers (e.g. TCP), which only require reliability as we have already seen. It is clear that such a semantics is not suited to multimedia communications at all.

On the other side of the “QoS support spectrum”, we find the concept of guaranteed QoS which, when used with a deterministic (also called hard) semantics, ensures the service users that the performance they required are going to be achieved throughout the communication lifetime. This is obviously the strongest QoS support semantics that one could ever dream of. However, achieving this semantics in the transport layer requires that this semantics be already supported in both the underlying networks and the Operating Systems (OS) supporting the applications. Unfortunately, even though several networks and OS have been designed to exhibit that performance guarantee characteristics (such networks and OS are often qualified “real-time”), their usage is not currently generalised and may never be.

Based on the observation that the best-effort semantics is not suited to multimedia communication and that the guaranteed semantics requires very specific mechanisms from the network layer, we have developed a new QoS semantics and associated negotiation scheme [2]. This work as well as the design and the specification of a new connection-mode transport service [1, 5] were achieved in the framework of an ESPRIT project called OSI95. In the OSI95 Connection-oriented Transport Service a QoS parameter is seen as a structure of three values, respectively called “compulsory”, “threshold” and “maximal quality”. Each value has its own well-defined meaning and is the result of a contract between the service users and the service provider.

The main idea behind the introduction of the enhanced QoS is that the service provider is committed to some well-defined duties, known by each side. In other words, the rules of the game are clear.

3.1.1. The Compulsory QoS Value

The idea behind the introduction of a “compulsory” QoS value is the following one: *when a compulsory value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and abort the service facility when it notices that it cannot achieve the requested service.*

No obligation of results is linked to the idea of compulsory value. The service provider tries to provide the requested service facility and, by monitoring its execution, will:

- either execute it completely without violating the selected compulsory value;
- or abort it if the selected compulsory value is not fulfilled.

3.1.2. Compulsory QoS versus Guaranteed QoS

The guaranteed QoS has a stronger semantics. When a guaranteed QoS value has been selected for a parameter of a service facility, the service provider will execute completely the service facility without violating the selected guaranteed value of the performance parameter.

The compulsory concept reflects the fact that, in some environments (e.g. a lightly loaded LAN), the compulsory QoS value may be achieved without resource reservation. Of course, the same LAN, which does not provide any reservation mechanism or any priority mechanism, may, when heavily loaded, prevent the service provider from reaching the compulsory QoS value and oblige it to abort the execution of the requested service facility.

The key point here is that although a compulsory semantics can take advantage of resource reservation mechanisms, it does not necessarily require them.

3.1.3. The Threshold QoS Value

Some service users may find that the solution of aborting the requested service facility, when one of the compulsory QoS values is not reached, is a little too radical. They may prefer getting information about the degradation of the QoS value.

To achieve that we introduced a “threshold” QoS value with the following semantics: *when a threshold value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and indicate to the service user(s) when it notices that it cannot achieve the selected value.*

This threshold QoS value may be used without an associated compulsory value. In this case, the behaviour of the service provider is very similar to the one it has to adopt with a compulsory value. The main difference is that, instead of aborting the service facility when it notices it is unable to provide the specified value, it warns either or both users depending of the service definition. If the service provider is able to provide a QoS value better than the threshold value, everything is fine.

3.1.4. Threshold QoS versus Best Effort QoS

If the threshold QoS is used without any compulsory QoS, the main difference between the threshold and the best effort is that in the former case, the service provider has the obligation to monitor the parameter and to indicate if the threshold value is not reached.

3.1.5. The Maximal Quality QoS Value

In most cases, if the service provider is able to offer a "stronger" value of the QoS parameter than the threshold, the service user will not complain about it. But it could happen, for reasons of cost or limited resources, that the service user wants to put a limit to a "richer" service facility.

To achieve that a "maximal quality" QoS value has been introduced with the following semantics: *when a maximal quality value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and avoid occurrence of interactions with the service users that would give rise to a violation of the selected value.*

It is possible to associate, with the same QoS parameter, a maximal quality value, a threshold value and a compulsory QoS value with, of course, the maximal quality "stronger" than the threshold value, itself "stronger" than the compulsory value.

Figure 3 illustrates the enhanced QoS.

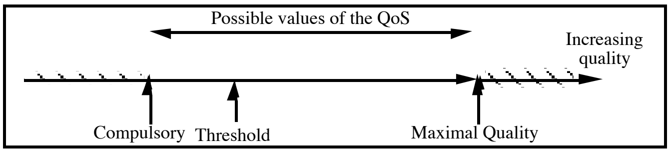

Figure 3: The Enhanced QoS.

3.2. Group Communication support

3.2.1. Group Communication Architecture

In the RACE project CIO, we developed a framework for group communication [8] as a preliminary study to the extension of the OSI 95 transport service with group communication facilities. That framework for group communication is in no way intended to relate to any particular layer of the OSI Reference Model but rather presents a general architectural model for group communications in a Multimedia environment.

3.2.1.1. Multimedia Group Communication model

To clearly understand the philosophy of that architecture, let us remind that each medium of a multimedia communication imposes its own requirements, which will be expressed in terms of Quality of Service (QoS) parameters, on the underlying network. It is easy to show that those requirements may differ from one medium to another. As an example, let us compare briefly audio and text.

While humans can cope with noise in audio information, it is much more difficult for us to accept too long or too variable delays in the delivery of audio information. Therefore, audio allows some error rates, but has stringent delay and delay jitter requirements. Moreover, the coding/compression schemes will impose some transmission rates as well.

On the other hand, what is important for text is the correctness of the message. Whether the transfer of the whole text takes more or less time or the different pieces of the text are received at regular intervals is not very important. Therefore, text will require full reliability while having rather loose performance requirements.

Moreover, a same medium does not always imposes the same requirements on the communication network. Indeed, audio can be Hi-Fi quality as well as telephone quality. In the same way, video can be black-and-white, colour, HDTV quality. Some coding schemes, called hierarchical coding schemes, provide for a "layered" medium, each level adding some quality to the previous one. We thus see that there is not only an heterogeneity in requirements from one medium to another, but that the heterogeneity can also exist “within” a given medium.

In order to cope with that heterogeneity, a multimedia communication is modelled, in our framework for group communication, as an call composed of several connections [8], each connection being “shaped” to given requirements. Figure 4 shows up a single call used for a teleconference scenario.

It is worth noting that the concept of call allows one to see, within the communication architecture, a multimedia communication as a whole, while each of the connections making up that call is related to a specific medium or flow of information.

3.2.1.2. Active Group and Topology

The call depicted on figure 4 represents an instance of group communication at a given point in time. The lecturers participating in the lecture are said to make up the active group of the call while the established connections are said to define the topology of that call.

The “look” (the active group and/or the topology) of such a call may change during its lifetime. For instance, one of the lecturers may need to establish a new connection to distribute a high-definition still picture to the other lecturers. We can also consider that a lecturer that had established a connection for overheads distribution might wish to release it when his “slide show” is over.

In the same idea, a new (or late) lecturer may join the on-going lecture, while one of the active lecturers may wish to leave the lecture before the end.

While the changes that may occur in a call are required to provide flexible group communications, some of those changes might well be considered undesirable or harmful in some cases. For instance, in our example, there is no point in continuing the lecture if either the chairman has left or the audio connections have all been released. Such situations are expressed by the Active Group Integrity (AGI) and the Association Topology Integrity (ATI) conditions.

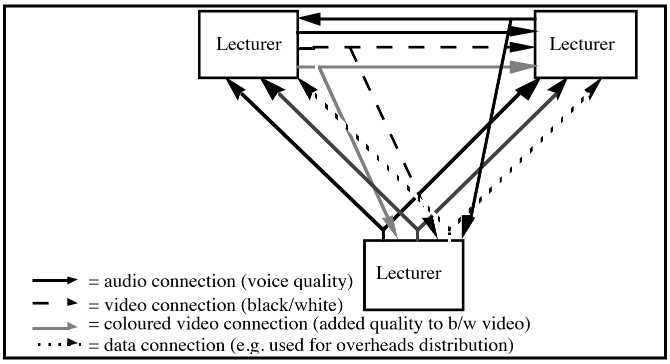

Figure 4: Model of Multimedia Communication.

3.2.1.3. Enrolled Group

As we have seen in the previous section, the active participants (the active group) of a given call are not the only potential participants of that call.

Moreover, we can also consider that other calls may need to be established among (part of) those potential participants for any other purposes. Therefore, we see that, at any given time, several independent calls, whose active groups may have different memberships, may be established simultaneously.

Although those active groups may include different participants, those participants exhibit similar properties. Indeed, in our example, the potential participants are all lecturers. Therefore, it may be interesting to gather those potential participants in what we call an enrolled group and to identify that enrolled group with a group address. Such a group address makes the set of potential participants appear as a single virtual entity, which is very useful in many cases. For instance, at the time a call is being established, the group address releases the “establishing entity” from knowing either the potential participants’ individual addresses or even the number of them.

In summary, in our architecture for group communications[4], we propose to gather in enrolled groups the entities with potentially the same activities

---

[4] The model described in [8] is a little more complex than presented here, but details about that model is beyond the scope of this paper.

(each enrolled group being identified by a group address). Calls are then established among subsets of the members of a same enrolled group for the purpose of actual activity (i.e. communication).

3.2.2. QoS Support for Group Communications

In the RACE project CIO, we also studied how the enhanced QoS semantics and negotiation, designed in the OSI 95 project, can be extended to the case of multicast connections (i.e. (1→N) connections).

An interesting result from this study [6] is that, in the multicast case, QoS parameters are characterized as:

- connection-wide QoS parameters, whose scope is the whole multicast connection (and thus affect the sender and all the receivers). The throughput is a typical connection-wide QoS parameter.
- receiver-selected QoS parameters, whose scope is limited to one receiver. For such a QoS parameter, a different value may be selected between the sender and each receiver. The transit delay and the delay jitter are examples of receiver-selected QoS parameters.

In [6], we also identified, and brought some solutions to, problems of incompatibility that can arise during a full QoS negotiation among all the parties involved in a multicast connection (i.e. the sender, the receivers and the service provider).

3.2.3. Multipeer Transport Service

In section 3.2.1, we have presented a general architectural model for group communications. We now present how we particularised that general architecture in order to design a transport service providing multimedia applications with efficient communication support. This transport service, called the "ACCOPI Multimedia Transport Service (AMTS)", was developed in the RACE project ACCOPI.

Of course, we naturally chose the enhanced QoS semantics as the QoS support in AMTS.

As regards group communications, the AMTS provides the transport service users with centralised transport calls [7].

A call is said to be centralised when there are only one sender and several receivers. A centralised call as provided by the AMTS is depicted on figure 5. In this figure, there are two (1→3) and one (1→2) connections.

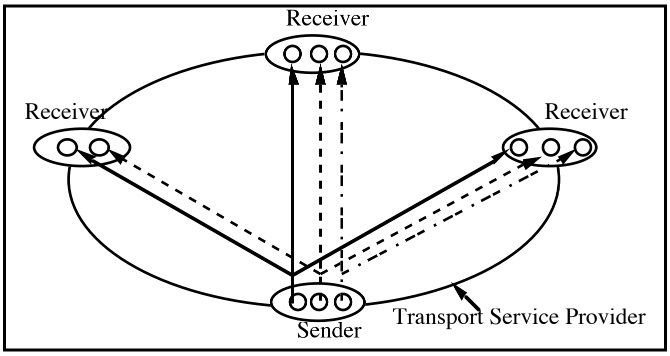

Figure 5: Centralised transport call.

The reasons that led to such topologies in the AMTS are the following:

- Each of the different media (or even each "part" of a given medium) making up a multimedia communication having its own characteristics, also have different performance requirements. Those media will often be transferred on different multicast connections, each one being tailored to offer the adequate performance. Therefore, "grouping" multicast connections together into calls, allows to handle, in the transport layer, multimedia communications as a whole. Indeed, the multicast connections carrying the different components media of a multimedia communication are no longer treated independently from one another, as it is the case with most of the current transport services.
- The centralised topologies allow the sender to control the entire call. This gives the sender the possibility to express QoS dependency relations among multicast connections of a call. A QoS dependency relation simply expresses the equality of the selected values for a given QoS parameter on all the multicast connections involved in that relation. There can only be one QoS dependency relation per parameter. It should be noted that the

set of multicast connections involved may be different from one dependency relation to another. Such QoS dependency relations on the transit delay and delay jitter can, for instance, ensure near-synchronisation, at the transport service interface, between audio and video connections.

- In order not to jeopardize performance at the transport layer, we naturally selected multicast (i.e. 1→N) connections as the basic communication scheme since such connections are the simplest for group communications. Moreover, our previous work on QoS support in group communications [6] has shown that it is really difficult to deal with QoS on connections other than multicast connections.

From this description of the AMTS, it is clear that the relationship among the connections of a call is expressed in terms of negotiated QoS values. Therefore, data concurrency between the transport connections (i.e. the ability to apply the same function to multiple pieces of data concurrently) is in no way reduced, which is a condition to achieve high-performance [4].

Therefore, we see that the AMTS provides the transport service users with a service interface well suited to the transport of multimedia communications. For more details on the AMTS, refer to [7].

4. Conclusion

For several years, we have been involved in research in multimedia communication. We have tackled, and are still considering, two of the main issues of multimedia communications, namely efficient QoS and group communication supports.

Following a “top-down” approach allowed us to first study the issues of multimedia communication from a global point of view and then to particularise and integrate the expertise we gained into an efficient multimedia transport service, which is a major achievement on the road to an integrated solution for multimedia communications.

Most of our results were considered in the European COST 237 action as well as introduced in standardisation (ISO/IEC JTC1/SC6).

Bibliography

[1] A. Danthine, Y. Baguette, G. Leduc, L. Léonard. *The OSI 95 Connection-mode Transport Service - The Enhanced QoS*. IFIP Transactions C-14, High Performance Networking, IV, Elsevier Science Publishers B.V. (North Holland), Amsterdam (1993), 235-252.

[2] A. Danthine, O. Bonaventure, G. Leduc. *The QoS Enhancements in OSI95*. The OSI95 Transport Service with Multimedia Support, Springer-Verlag, Berlin (1994), 124-149.

[3] M. de Prycker. *Asynchronous Transfer Mode: Solution for Broadband ISDN, Second Edition,*. Ellis Horwood, 1993.

[4] D. Feldmeier. *A Framework of Architectural Concepts for High-Speed Communication Systems*. IEEE Journal on Selected Areas in Communications, May 1993, vol.11, No.4, 480-488.

[5] A. Danthine, Y. Baguette, G. Leduc. *Belgian National Body Contribution - Issues Surrounding the Specification of High-Speed Transport Service and Protocol*, ISO/IEC JTC1/SC6 N7312, May 1992.

[6] L. Mathy, O. Bonaventure. *QoS Negotiation for Multicast Communications*, International COST 237 Workshop on Multimedia Transport and Teleservices, Lecture Notes in Computer Science 882, Springer-Verlag, Berlin (1994), 199-218.

[7] L. Mathy. *On the Design of a Transport Service to Support Multimedia Applications*, Proc. 4th Conference on Broadband Islands, Dublin, September 4-6, 1995, F. Williams, H. Bräke, J. Nolan, eds., 67-75.

[8] L. Mathy, G. Leduc, O. Bonaventure, A. Danthine. *A Group Communication Framework*. Broadband Islands ’94: Connecting with the End-User, Elsevier Science Publishers B.V. (North-Holland), Amsterdam (1994), 167-178.