Multimedia group communications: towards new services

1996 Distrib. Syst. Engng. 3 197

(http://iopscience.iop.org/0967-1846/3/3/005)

Distrib. Syst. Engng **3** (1996) 197–210. Printed in the UK

Multimedia group communications: towards new services

A Mauthe†, D Hutchison‡, G Coulson§ and S Namuye‖

Computing Department, Lancaster University, Lancaster LA1 4YR, UK

Received 24 July 1995, in final form 21 July 1996

Abstract. Interpersonal communication among a group of users employing different media types is becoming more and more widespread in computing and telecommunications. Group communication places a variety of new requirements onto the underlying communications architecture and although many existing protocols and services do offer some limited support for multicast group communication, these new requirements make it difficult to find efficient and comprehensive solutions. The impact of multimedia group communication on the communication system and the way in which existing systems, international standardization bodies and researchers cope with these challenges is the subject of this paper. First the characteristics and requirements of multimedia group applications are discussed and illustrated by examples of existing group applications. Subsequently a survey of the kind of support available in today's communication system is presented. In addition the ongoing discussion about the standardization of group communication within ISO and ITU and the direction these efforts take is briefly summarized. Further, some selected examples of research projects which deal with different communication and protocol related aspects of multimedia group communication are presented which give an indication of the trends in this area.

1. Introduction

In recent years computers have developed very rapidly from simple processing machines to sophisticated communications systems employing multiple media. A wide range of new applications employ audio and video to convey information or to support communication among human users. Many of these interactive multimedia applications operate in a group environment. The main characteristics of group applications employing multimedia components are: multiple senders and receivers, high data volumes, high data rates and time-dependent data values. Thus, there is a need to support real-time transmission at high data rates to multiple heterogeneous receivers.

A number of communications protocols and networks already provide group support to a certain extent. For example ethernet, DQDB, IP and XTP offer multicast ($1 : N$) data transmission. However, although multicast is a fundamental basis for group communication, it is by no means sufficient for the support of multimedia group applications. Different communication topologies including support for many-to-many ($M : N$) communication is also required. Further, new management protocols and services are needed to fully support applications such as groupware and conferencing. The area is developing fast and more and more protocols are now offering at least basic multicast and group enhancements. Certain trends can be identified but it is still too early to determine which approaches fulfil real user requirements best and will therefore succeed. International standardization organizations, namely ISO and ITU, are also addressing these issues and are currently discussing how group communication should be accommodated in new and existing standards. Last, but not least, there are various universities and research institutes working on different aspects of multimedia group communication and its support in the communication subsystem.

In this paper we present a survey of multimedia group communication, including an analysis of user requirements, a study on the support offered in today's systems, an outline of standards activities in group communication, and a discussion of research projects in this area. The paper is organized in seven sections. In section 2 a few selected examples of multimedia group applications are presented. Their characteristics and requirements are then systematically discussed in section 3. The main purpose of this discussion is to provide criteria according to which group communication protocols and services can be assessed. Section 4 then presents a survey of existing group communication support provided in today's communication protocols and services. Standardization efforts, and the discussion and trends in this area are outlined in section 5. Some selected examples of research projects dealing in particular with communication aspects of multimedia group

† E-mail: andreas@comp.lancs.ac.uk
‡ E-mail: dh@comp.lancs.ac.uk
§ E-mail: geoff@comp.lancs.ac.uk
‖ E-mail: namuye@comp.lancs.ac.uk

applications are discussed in section 6. A summary and discussion of the results of the survey conclude the paper.

2. Multimedia group applications

In general, multimedia group applications can be divided into two classes, *dissemination services and interactive applications*. Dissemination services distribute information among a group of users, for instance traditional broadcast services (e.g. radio and TV programs), broadcast services for specific user groups (e.g. registered students), or for certain domains (e.g. a company). Interactive applications are applications where direct interaction between a group or sub-group of participants is the main purpose of the communication, e.g. conferencing applications, remote teaching, but also public hearings and panel discussions. In this section we concentrate on interactive multimedia group applications. This is because the characteristics of these applications and the requirements they place onto the underlying communications system are much more diverse than those of dissemination services. To illustrate their requirements and the implications they have on the underlying communication system we discuss a few selected examples of existing systems and system scenarios. A more detailed study of multimedia group applications can be found in [49].

Within the **MOST** project (Mobile Open Systems Technologies for the utilities industry) at Lancaster University a prototype collaborative multimedia application to support field engineers of the UK power distribution industry has been developed. The application provides support for information exchange and collaboration among a group of engineers working at different points of the electrical power supply network [14]. The main components are the conference manager, the audio communication tool, and the Geographic Information System (GIS). The latter is used to view geographical information (e.g. network diagrams). Further, a tool to control the audio communication and multiple windows to display maps and other relevant information (e.g. the history of operations) is offered as part of the user interface. Maps and the history of operations can be shared among multiple users. The distributed conference manager (or group coordinator) is responsible for group management tasks like establishing and maintaining user groups, and for the coordination of activities of any module operating in conference mode. A mobile enhanced version of the ANSAware platform [2] is used in the implementation of the prototype. The application is being used as part of a field trial by a regional electricity board [18].

TeamWorkStation-1 (TWS-1) is a locally remote co-authoring/argumentation system. It provides a small group of 2–4 users with a shared workspace which they can simultaneously see, point at and draw on [33, 34]. Each user can run application programs and manual tools concurrently in the shared workspace. Distributed Macintosh computers are connected by a data network (LocalTalk), voice network (telephone), a specially developed video network (based on NTSC and RGB) and

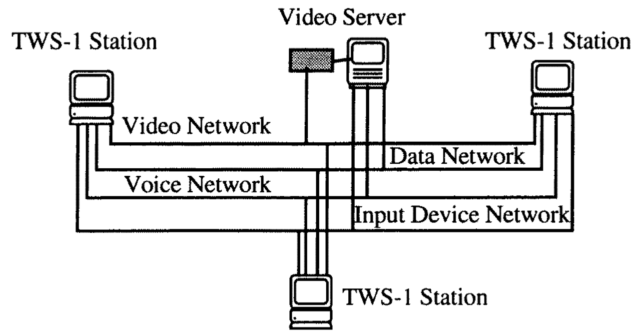

Figure 1. Network architecture of TWS-1.

an input device network. Figure 1 shows the network architecture of TWS-1.

On each workstation there is a shared screen containing a shared drawing window and windows with video images of every participant. A window on the shared screen is transmitted to all other participants and is therefore accessible to all members of the collaboration. For the shared drawing window a translucent video overlay technique is employed (i.e. the images of the users writing on the board are overlayed; users get the same impression as writing concurrently on the same sheet of paper). A video server which gathers and processes the shared screen images also controls the video network.

In [60] a virtual workspace for *virtual teleconferencing* is described. A virtual conference room is created using computer graphics in real-time. Images of objects that can be touched and moved with a special data glove and joystick are constructed together with the images of remotely located conference participants. The user interacts with the system through the data glove. Special shutter glasses are employed to derive the user's eye and head position. Further devices are a special 70-inch display, and a graphic workstation. The computer graphic images are locally generated based on the user's head position. Head position and data glove information are transmitted to each site which holds a *virtual world data base manager*. Since every user has a copy of this, and image processing is done individually at each site, the transmitted data volume is relatively low. For data communication TCP/IP sockets are employed. Additionally, an in-house telephone network is used for the communication of voice.

Insufficient support for multipeer communication is the prime reason for the use of proprietary and very specialized communication architectures in today's systems (as illustrated in the above examples). With emerging high-speed networks such as ATM an integrated service utilizing a single network is becoming possible. However, the necessary infrastructure for this is not everywhere available and heterogeneity in networks and end-systems will remain. This is especially true in the case of mobile computing and user mobility.

3. Group application requirements

In [48] a set of characteristics describing group applications is presented in detail. Various aspects of group communication requirements have also been studied by other authors (e.g. [13, 27, 59]). However, the nature and extent of support required from the communication system is still controversial. Where and how certain services should be provided is also an issue of discussion. There are no general answers to these questions; nor are there any generally recognized criteria by which proposed solutions can be assessed. In the following subsections we outline some requirements and characteristics of group applications which are particularly relevant to the underlying communication service.

3.1. Application requirements: examples

Group applications range from simple message passing and talk applications to sophisticated conferencing and collaborative applications using application sharing and high-quality audio and video for information exchange. All of them have different requirements in terms of bandwidth, timing, quality, reliability, etc. The following three examples give an idea of the diverse requirements encountered.

MOST is a typical example of a wide area mobile application. It is optimized to cope with the low bandwidth and unstable environment it is operating in. In general, the number of participants is rather small; three to five at any time. With an audio encoding such as Group Special Mobile (GSM) 16 Kb s^{-1} is required. Assuming that not more than two participants speak at any time, the requirements will always be less then 32 Kb s^{-1}. The delay for audio should be around 250 ms and the jitter less then 10 ms. All shared information is held locally and only updates have to be exchanged. The data rate required for this is 300 b s^{-1}. A reliable protocol is required for the transmission of updates. With application level ordering in conjunction with the history maintained at every site, total ordering can be achieved. An important feature of this application is system awareness. For security reasons all participants have to be informed as soon as changes in the underlying communication such as the temporal disconnection of members occur.

Another example of a scientific application is the collaborative microscope application discussed in [63]. A number of scientists can jointly view and operate an electronic microscope in a local area environment. All data from the microscope are recorded as well as multicast to each participant. The scientists communicate via an audio link. The demands on the microscope output are very high in terms of quality and reliability. A high resolution video output according to the HDTV standard requires at least 34 Mb s^{-1}. A complex compression scheme cannot be used since these introduce large delays which are unacceptable for live examinations. Because of the required accuracy, a low error rate of the order of 10^{-9} is essential for such an application. With Pulse Code Modulation (PCM) encoded audio, 64 Kb s^{-1} is required for every audio channel. This can be easily reduced since usually only one participant will speak at any time. The delay for both audio and microscope output should be less then 250 ms. Bandwidth requirements of tools such as shared pointers are negligible. However, total ordering to operate these pointers is required. To avoid more than one scientist operating the microscope at any time, a token protocol should be used.

The most common scenario for group applications is tele-conferencing in a wide area environment. In general these systems have audio and video components for interactive communication among participants. With Adaptive Pulse Code Modulation (APCM) for audio, 32 Kb s^{-1} is required for every stream. Mixing and joint resource reservation for all audio streams could keep the required bandwidth at about two to three times this value although the number of participants might be much larger. The bandwidth requirement of video with conferencing quality is 0.32 Mb s^{-1} (H.261 standard encoded video for H0 ISDN). If high quality video is required this number can be 30 Mb s^{-1} and higher. Every participant adds to the overall bandwidth requirement of the application. Complex encoding schemes cannot be used since the overall delay has to be kept at around 250 ms (for video as well as for audio). Application sharing tools such as slide presentation, whiteboards, joint editors and CAD/CAM are also commonly part the conferencing application. Depending on their sophistication and the way data are stored and up-dated, their bandwidth requirements range from a few hundreds of bits per second to up to 50 MB s^{-1} (e.g. CAD/CAM). Because of their interactivity, delay requirements are in the same range as audio and video. If concurrent operations are allowed, ordering protocols are necessary to ensure data consistency. The kind of ordering is application dependent. Also, participants might assume certain roles (e.g. floor manager, observer etc) and have distinct rights which might have to be supported by the system. Last but not least, data protection and privacy have to be ensured for confidential and sensitive data.

3.2. Resource requirements

Multimedia applications in general, and multimedia group applications in particular, have very high bandwidth and processing requirements. Usually four main parameters are used to describe quality of service (QoS), and hence the quantitative resource requirements of multimedia applications:

(i) *throughput* determines the data rate needed for the transmission of a certain type of data;
(ii) *end-to-end delay* is the total delay for a data unit transmitted from the source to its destination;
(iii) *delay jitter* is the maximum variance in the interarrival time of two consecutive data units;
(iv) *error rate* refers to the number of lost or corrupted messages which can be tolerated.

Throughput requirements of discrete media are often expressed without any time reference because there are no stringent time requirements for this kind of data, and only the amount of transmitted data can be clearly determined.

Images are very often used in interpersonal communication. Depending on their use the quality required can vary considerably. This is also reflected in the storage requirements of different image types. For a 500 line vector

Table 1. QoS requirements of continuous media data.

Media	Throughput	Maximum error rate	End-to-end delay
Telephone quality audio	64 $\mathrm{Kb\,s^{-1}}$	10^{-1}	100 ms
High quality audio	1.4 $\mathrm{Mb\,s^{-1}}$	10^{-2}	100 ms
TV quality video	30–140 $\mathrm{Mb\,s^{-1}}$	10^{-3}	100 ms
Compressed video	2–10 $\mathrm{Mb\,s^{-1}}$	10^{-9}–10^{-11}	100 ms

graphic just 2.8 Kbytes are needed, a 256 colour bit map image requires 384 Kbytes (768×512 pixels), and assuming HDTV values (1250×1666 pixels) the storage requirement for a one bit map is 2.003 Mbytes. Some interesting figures about images used for medical examination are introduced in [11]. For instance chest x-rays on average require 88 Mb per examination (with 1.5 images per examination on average). A computer tomography image requires on average 126 Mb per image but 30 images per examination is a typical figure. The transmission and processing of images is not time critical but in interactive applications the overall delay should be bounded. Hence the underlying communication system has to be able to deal with burst data transfer with bounded delay.

For continuous media, in contrast to discrete media, the time aspect is essential. The QoS requirements of continuous media depend greatly on media type and encoding techniques. For digitized speech of telephone quality the throughput value is 64 $\mathrm{Kb\,s^{-1}}$. Mono-audio data of CD quality require 705 $\mathrm{Kb\,s^{-1}}$. HDTV-contribution-quality requires approximately 140 $\mathrm{Mb\,s^{-1}}$; for HDTV-distribution-quality 34 $\mathrm{Mb\,s^{-1}}$ is the rate currently under discussion for standardization [58]. On the other hand, with compression techniques data rates of continuous media can be reduced considerably. With MPEG1 encoding, for instance, a maximal data rate of approximately 1.5 $\mathrm{Mb\,s^{-1}}$ is required. Table 1 shows typical QoS requirements of audio and video.

A major problem with group applications is the potentially large number of participants which entails a large overhead in terms of connection management. Related to this are resource problems if multicast and multipeer communication is not sufficiently supported by the underlying communication system. For instance for a simple audio/video conference among six participants five audio and video connections have to be managed at each site if there is no multicast support available. This adds up to 60 connections for the whole application. Assuming PCM encoded audio with 64 $\mathrm{Kb\,s^{-1}}$ and H.261 standard encoded video for H0 ISDN with 0.32 $\mathrm{Mb\,s^{-1}}$, the bandwidth requirement of this conference is then 11.52 $\mathrm{Mb\,s^{-1}}$. On the other hand, with multicast the number of connections could be reduced to 12 and the overall bandwidth requirement is 2.304 $\mathrm{Mb\,s^{-1}}$. Every participant and every medium employed will add to the resource requirements and the complexity the system has to handle.

3.3. Topology and service requirements

The most general communication topology is $M{:}N$ communication, with unicast and broadcast as special cases. Within this general paradigm, a set of basic services can be defined.

The most fundamental service is a *multicast* service ($1{:}N$) where one data copy is transmitted from a single sender to multiple receivers. Data delivery in the multicast case is only considered to be successful if certain conditions, so called *integrity conditions* (sometimes also called multicast reliability or degree of reliability), on the number and/or identity of receivers who have to receive a 'correct copy' of the data, are met. For continuous media data, these constraints are tied to all QoS parameters characterizing a multicast data stream and do not just apply to individual data packets. In other words, as long as the QoS specified for the recipients in the identified sub-set is met data transfer is deemed successful. Quantitative conditions used in general are *quorum* (i.e. a qualified majority) and *k-reliability* ($0 < k \leq N$).

A *virtual multipeer* service ($M^*(1{:}N)$) is one way of building a general $M{:}N$ topology. In this scheme, data are transmitted from multiple senders to multiple receivers over different multicast data streams. At the receiver side the individual data streams can be distinguished, but they are managed as one entity and appear as such to the service user. The virtual multipeer service is characterized by joint multicast connection management, including joint QoS management and the validation of *multipeer integrity* conditions. Joint QoS means that QoS of different multicast connections belonging to the same virtual multipeer connection are specified and managed together. A change in QoS has to affect all multicast connections in the same way, and resources have to be shared whenever possible. Multipeer integrity conditions are the above defined integrity conditions extended to cover the multiple sender case, i.e. a virtual multipeer connection is only deemed successful if the conditions specified on receivers *and* senders are met.

In contrast to the virtual multipeer service, a *multipeer service* provides a topology where a single semantically meaningful data stream composed from different sender data streams is delivered to each service user (i.e. recipient). In multimedia group communication this can be achieved through the instantiation of *mixing filters* [66]. Integrity conditions apply to both types of multipeer connections in the same way. QoS service in the multipeer case depends on the number of senders, the sending mode, and filter techniques employed. It can vary dynamically during the communication.

Finally, a *session* is composed of various (virtual) multipeer and multicast connections. These connections are jointly established and maintained during the lifetime

Table 2. Multipeer communication services.

Communication type	Characteristics
Multicast	• 1-to-N data delivery • QoS • integrity conditions
Virtual multipeer	• M-to-N data delivery • M distinguishable data streams at receiver side • joint stream management • joint QoS • multipeer integrity
Multipeer	• M-to-N data delivery • 1 semantically meaningful data stream at receiver side • multipeer QoS • multipeer integrity
Session	• multiple M-to-N data streams • joint (virtual) multipeer connection management • QoS trading and prioritizing • session integrity

of the session. It should be possible to trade QoS between different connections and to prioritize their QoS requirements. Session integrity refers to the number and/or identity of connections which have to be present. All discussed communication types have to allow dynamic join and leave. Table 2 summarizes the different communication services.

Related to topology is ordering and the kinds of ordering semantics which can be applied to different topologies. Different ordering semantics are proposed [8, 9, 37, 40]. The main ordering relationships generally considered are *no ordering*, *source ordering* (i.e. all messages from one source are delivered in the same order to all receivers), *causal ordering* (i.e. messages from the same or different senders are delivered according to the causal dependence relationship) and *total ordering* (i.e. all messages from all senders are delivered to each receiver in the same order). No ordering and source ordering apply to all topologies whereas causal ordering and total ordering are only applicable in the multiple sender case (namely virtual multipeer and multipeer). Ordering is a pure end-to-end issue and thus has to be addressed at the transport layer or above. Moreover, causal ordering is an application level concept since there is no way to establish ‘true’ causal relationships among data messages from multiple senders.

Ordering relationships among continuous media are usually referred to as synchronization requirements. In the communication system they are supported by the temporal restrictions (i.e. delay, delay jitter) which are part of the QoS service description of a data stream.

3.4. Other aspects

Successful multipeer communication services also require an efficient *management* of groups. Groups are highly dynamic; membership, group properties and characteristics change during their lifetime. The management of groups includes membership administration, address management and the management of group properties and characteristics. Group management is an orthogonal task to communications provision since groups at different levels have to be administered. In addition, the lifetime of groups is, compared to a single multipeer communication, much longer. This has to be taken into account in the design of such a service. The services offered by group management include membership and group administration and a directory service.

The *timeliness of session establishment* is also a key issue. Long delays between the set-up of commands and their execution are annoying and can cause user acceptance problems. Although some of these interactions with the system are not real-time or time-critical in a technical sense, good performance is necessary.

Flexibility and dynamics must also be considered. For example, a change of QoS might not only be provider initiated but also be required by the user. Different participants might have *different data views* on shared objects or use different encoding standards. *Awareness* is important for many group applications. Hence, notification of communication state and operations performed by other participants is required. *Data security* and *privacy* is a very sensitive issue especially when data are transmitted over public networks or in medical and business applications where confidential data are exchanged.

Heterogeneity is traditionally a problem in distributed systems and is particularly evident in group communications systems. Users employ a wide variety of end systems with different capabilities ranging from high performance multimedia stations to low capability portable computers. On the communication side high-speed networks are encountered as well as low bandwidth mobile links. In a truly open system it has to be possible to support the communication between members of a group located in such an environment.

4. Group support in communication protocols and services

Group and multicast communication is supported by a number of protocols at different layers in the communications architecture. The provision of group and multicast protocols is currently an active area of research and many new developments can be observed in this area. Developments in the area are moving fast but certain trends can be identified as discussed below.

4.1. Group support in link layer protocols

In local area networks, group communication depends on the ability of the underlying network to broadcast messages. This is for instance the case in a ring or bus topology. A message placed on the network will eventually pass by each connected station. It is picked up by every station listening which has the appropriate multicast address set on the network interface adapter. Usually any station on the ring or bus can send a message without any restriction. Most protocols standardized in *IEEE 802.x* (e.g. token ring, ethernet, DQDB) provide multicast. Traditional ethernet is not suitable for the guaranteed transmission of continuous media data. However developments such as *isochronous ethernet* [58] do give QoS guarantees. Token ring is more suitable for the support of continuous media than ethernet but is still far from ideal in terms of QoS provision. DQDB offers asynchronous, synchronous and isochronous services and thus provides all the support needed for multimedia communication.

ATM supports one-to-many connection oriented multicast where one end-system (called the *root*) may send data cells that are multicast and received by a number of other end-systems (called *leaves*) [3]. ATM is essentially a point-to-point network; its ability to broadcast or multicast messages depends on whether or not copying functions in the switches are available. Different mechanisms to route multicast and broadcast messages in ATM networks have been proposed, namely centralized repetition in a copy network, copying before the switch by a serial copy function, and distributed repetition in the switch fabric. Many-to-many connections could, from a pure ATM switching point of view, also be supported [68]. A number of properties for this transmission mode in ATM have already been specified [42]. However, many-to-many connections over ATM are still an issue of research and not yet available. The QoS provided by ATM is considered sufficient for multimedia applications.

4.2. Group communication support in the network

4.2.1. Multicast routing mechanisms. To support multipeer communication in the network, messages have to be copied and routed to all members of the (receiver) group. To provide multicast in WANs existing routing algorithms have to be extended. Commonly, group membership protocols are used for routers to learn about members on their sub-network. There exist various different methods to propagate routing information among multicast gateways. *MOSPF* for example is an extension of the link state unicast protocol OSPF (Open Shortest Path First) [16]. In this scheme, changes in group membership are first detected by a router directly attached to the respective sub-network. The change is broadcast by this router to every other router in the same routing domain. The state of the domain topology is kept in the (unicast) link-state protocol at every router. A multicast message is forwarded according to the shortest path tree which is determined using group membership and topology information [19]. *Inter-Domain Policy Routing* (IDPR) also uses link state routing information but, in contrast to MOSPF, source specified packet forwarding is used [51].

The DVMRP *(Distance-Vector Multicast Routing Protocol)* uses variants of reverse path routing to erect a multicast tree. In this scheme a router forwards the first packet sent to a group to all outgoing links. A router on a sub-network without any members sends a prune message towards the source of the packet. This message prunes the branches where no group members reside. Hence a shortest path tree that has only leaves containing members of the group is erected. To include new potential members, pruned branches ‘grow back’ and are eventually pruned again if there are still no members. Special mechanisms to prevent packets from looping are employed [19].

Yet another way to establish multicast trees is through *Core Based Trees* (CBT). In CBT one router acts as a core from which all branches emanate. Receivers that wish to participate have to find the core and attach themselves to it [4]. *Protocol Independent Multicasting* (PIM) is, like CBT, based on the receiver initiated philosophy. Two different types of multicast tree, shared and source-specific multicast trees, can coexist [19, 51].

4.2.2. Internet multicast schemes. The *Internet Protocol* (IP) provides a scheme for multipoint data delivery. IP multicast groups are dynamic, i.e. a host can join or leave a group at any time. Further, a host can be a member of several groups. Class D addresses in IP are especially reserved to address multicast groups. They can be easily mapped onto link layer group addresses like ethernet multicast addresses. Some IP multicast group addresses are permanently assigned by the Internet authority to groups which always exist (even if they currently have no members). Other groups are transient, i.e. they are created when needed and discarded when the count of members reaches zero. If a multicast group spans multiple networks, group membership information is communicated by the *Internet Group Management Protocol* (IGMP). A host that wishes to join a group issues an IGMP request [16].

MBONE is a virtual network that provides multicasting facilities on top of the Internet [21]. It is composed of multiple multicast networks. There is one host in each of these networks that runs a multicast routing demon called *mrouted*. The routing algorithm commonly used is DVMRP. Different mrouted demons are connected through so-called *tunnels*. A multicast packet sent out by a client is picked up by the local mrouted demon, and it chooses the tunnels into which the packet has to be sent out according to its routing table. The mrouted on the other end of the tunnel receives the packet, examines its routing table and forwards it to any client on its sub-net that has subscribed to the group and, if necessary, to other tunnels. Figure 2 shows a schematic MBONE topology. Some pruning for MBONE is implemented. MBONE also provides a wide variety of video, audio and whiteboard applications and tools such as vat, nevot, nv, wb, etc [21].

ST-II is a stream oriented protocol, designed to provide end-to-end guarantees across the Internet. The offered service is connection oriented; in the multicast case a multi-destination connection is established along a directed multicast distribution tree. During the establishment phase the necessary network resources are reserved. This has to be agreed to a common QoS among all involved entities; i.e. the entire stream has the same QoS. New receivers have to establish a connection with the source. Further, ST-II allows the clustering of several streams into stream groups [20, 62].

RSVP is a resource reservation set-up protocol, designed as a companion protocol to IP [67]. RSVP messages do not carry any application data; they just control the packet transmission of IP. Multicast in RSVP is built

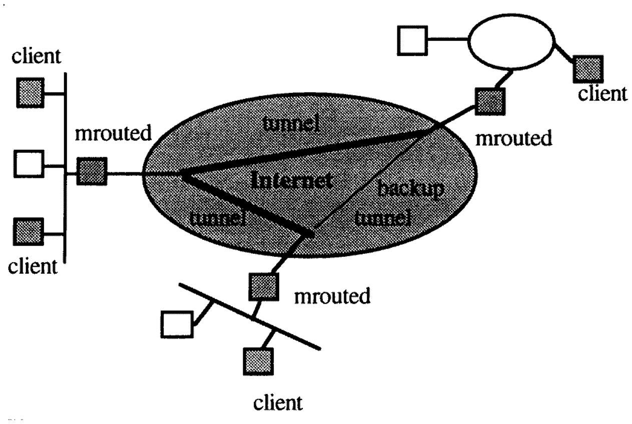

Figure 2. MBONE topology.

around existing and future multicast routing protocols. A host first has to send an IGMP message to join a multicast group. Multiple senders can send to the group. To reserve resources a receiver determines its QoS and initiates a *reservation* request. The required resources along the subnet(s) from the receiver to the sender are reserved through intermediate RSVP agents. Reservation messages are propagated towards the sender until an existing distribution tree with sufficient resources allocated for the sender group is found. The reservation model has two components, a *resource allocation* where the amount of resource is specified and a *packet filter* which selects the packets that can use the resource. With a *wildcard* filter no source specific filter is required. A resource allocation is made for all traffic directed to the receiver that initiated the resource reservation. When a source specific filter is required the receiver might specify a fixed set of senders (*fixed-filter*). With a *shared-explicit filter* receivers can specify a set of sources which they want to reserve a certain amount of resources for. For example they can request resources for one audio channel which are to be shared amongst ten explicitly specified senders [12, 50].

At present the Internet protocol architecture is undergoing a number of major changes. All proposals for new or modified protocols are subsumed under *IP next generation* (IPng), or *IP version 6* (IPv6) [32]. One of the biggest changes is the step from 32 bit addresses to 128 bit addresses. New IP addresses are also different in structure, especially as far as multicast addressing is concerned. A multicast address is identified by 1 in the first 8 bits of the address followed by the four bits `flgs` field, the 4 bits `scop` field and the 112 bits `group ID`. Group identifiers can be either permanently assigned or transient. The scope of the group is limited, i.e. transient group identifiers are only valid within a given scope. Permanent group identifiers are globally meaningful but their scope can also be bounded. The scope of an address can be intra-node, intra-link, intra-side, intra-organization, intra-community and global. The 112 bit `group ID` identifies either a permanent or transient group within a given scope [31]. So called *cluster addresses* are used to reach the nearest of a set of so called boundary routers of a cluster of nodes which are identified by a common prefix. This allows one, for instance, to chose the provider which should carry the data [32]. In IPv6 the Internet Group Membership Protocol (IGMP) has been absorbed into the *Internet Control Message Protocol* (ICMP). There are two classes of ICMP messages, error messages and informal messages. Three informal messages to exchange group information are defined, Group Membership Query, Group Membership Report and Group Membership Termination [17]. In IPv6 a 28 bit Flow Label field in the header is introduced to support QoS. Traffic is handled by the IP router according to the traffic class in the 4 bit traffic class sub-field of the Flow Label [30]. This is particularly useful for data flows such as continuous media streams with QoS specifications.

4.3. Group support in end-to-end protocols

4.3.1. Internet protocols. At the transport layer a group consists of sending and receiving processes. Most traditional transport protocols like OSI-TP4 and TCP do not support any multipeer data communication. Emerging protocols and protocol enhancements do now support multicast at the transport layer to some extent. In fact, some TCP implementations already support multicast data delivery over IP multicast. A protocol supporting *reliable* multicasting on top of IP through a so called *Single Connection Emulator* (SCE) is proposed in [61]. The SCE resides above IP and can be used by TCP. It mimics a single destination network layer interface to TCP. IP multicast is used to provide the necessary multicast functionality. As with TCP, source ordering is provided.

The *Real-Time Transport Protocol* (RTP) offers multipeer data transfer (if supported by the underlying network) for data with real-time characteristics. It was designed for the Internet and resides on top of protocols such as UDP, TCP, OSI TP1 or TP4 and ST-II [55]. For multicast data transmission a multicast address is obtained by the initiator and distributed out-of-band to the participants. RTP packets are received at the port of the underlying transport service, and encapsulated in packets of the underlying transport system, e.g. if transmitted over UDP, RTP header and data are contained in a UDP packet. The RTP header itself contains timing information and a sequence number which enables the receiver to reconstruct the timing seen by the source. RTP does not provide mechanisms to ensure timely delivery or to provide QoS guarantees. It also does not guarantee delivery or prevent out-of-order delivery. RTP is accompanied by the *Real-Time Control Protocol* (RTCP). This provides information about participants in on-going sessions. Using RTCP, each participant sends local information periodically to all other participants as options within RTP packets. RTCP necessarily has to be used to communicate over RTP [56].

4.3.2. The Xpress Transfer Protocol (XTP). XTP (Xpress Transfer Protocol) was defined as a protocol subsuming both network and transport layer functions. A connection-oriented one-to-many multicast, if supported by the underlying network, is proposed [53, 64]. XTP is connection oriented, hence source ordering is supported.

Four techniques to improve XTP's group communication support are suggested: the *bucket algorithm*, *slotting*, *damping* and *cloning*. The bucket algorithm was proposed as a control mechanism to determine how and when to update status information. Control information processing is timer based; 1-reliability can be achieved. Slotting and damping are mechanisms to reduce the number of control packets. Slotting forces a receiver to wait a random time span before generating a control packet. In case of an error a receiver sends control messages to the sender and all other group members. Other receivers damp their own control messages when they contain a larger value than that of the received control packet. Error recovery uses go-back-n. With cloning a multicast transmitter can span slave contexts from a master context. This is a way to achieve concentration. Multiple senders can send to the same multicast address. No group management was defined in XTP 3.x [53].

The new version of XTP, XTP 4.0, is defined solely as a transport protocol. Multicast is now an integral part of the protocol definition and not simply an addendum. XTP multicast provides a simplex data flow from one transmitter to an arbitrary number of receivers. Control procedures defined for unicast apply to multicast as well. Group management to determine membership in the receiver group is also defined. This allows the support of different reliability semantics†; i.e. k-reliability and full reliability (k equals the number of participants), and to determine a subset of receivers who have to receive a message. XTP 4.0 supports go-back-n or selective retransmission as its error recovery mechanism. In the reliable multicast case *all* data reported missing by any receiver have to be re-transmitted to the group with either error recovery scheme. Slotting is still used to distribute the number of control messages over time. As in the unicast case, synchronization handshakes are used for information status up-dates. No data can be transferred during a synchronization handshake [65].

4.3.3. ITU-T T.120 series. The *Multilayer Protocol* (MLP), also known as the *T.120 series*, recently developed within the International Telecommunication Union–Telecommunication Standardization Sector (ITU–TSS), is designed to provide support for multimedia applications and services in a heterogeneous environment [5, 15]. The MLP can be used with various types of networks, namely public switched telephone networks (PSTN), integrated services digital networks (ISDN), circuit switched digital networks (CSDN), public switched digital networks (PSDN) and broadband-ISDN (B-ISDN). Further, support for other networks like LANs, TCP/IP, IPx on LANs, and ATM is under consideration or being developed [10]. Multipoint service is provided on top of reliable point-to-point links in a network independent manner. Groups of point-to-point transport connections are mapped together and build a Multipoint Communication Service (MCS) domain. Each domain represents effectively an independent 'conference' or group session. Within a domain different logical channels can provide (1:1), (1:N) or (N:1) data delivery. Source ordering is supported.

† Note that reliability or group reliability in XTP is comparable to the concept of integrity as defined in section 3.2.

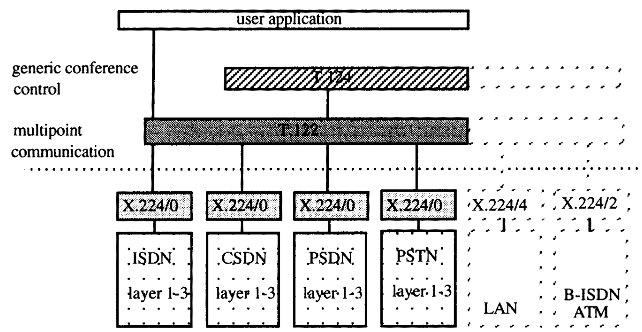

Figure 3. T.120 communication structure.

The *Generic Conference Control* (GCC) (T.124) protocol is located above MCS and provides services for managing multipeer conferences including conference establishment and termination, joining ongoing conferences, managing changes during sessions, conference database management, application database management, bandwidth control and remote actuation. Real-time and non real-time aspects are coordinated by the GCC. Apart from these protocols four application internetworking protocols are suggested: An *Audio-Visual Control* standard (T.AVC), a *Binary File Transfer protocol* (T.MBFT), a protocol to support shared whiteboarding and exchange of still images (T.SI) and a facsimile service. Figure 3 shows the T.120 communication structure.

5. Group communication standardization

Multipeer communication is currently also being addressed by the major international standardization bodies, namely IETF (Internet Engineering Task Force), ISO (International Organisation for Standardization) and ITU (International Telecommunication Union). The work carried out in this area by IETF was presented in the previous section (i.e. the IP multicast work). Recently IETF and ISO have held talks about a collaboration on some of the issues discussed in this section. However, no formal agreement has yet been reached. ISO and ITU are now, after initial differences in this area, working together on the definition of standards that support multipeer communication. Group communication issues are within the remit of various committees and working groups in ISO and ITU. Within ISO the most active groups addressing group communication issues are SC6/WG4 (dealing with network and transport layer issues) and SC21/WG1 (addressing architectural and also QoS issues). In ITU these issues are covered by ITU-T SG7.

In the discussions on multipeer transport services and protocols which have been taking place over the last two years, two major trends represented by different groups can be identified. The first group is lobbying for a quick standardization of a simple multicast solution, whereas the second is advocating a more comprehensive approach including the concurrent standardization of conceptual and architectural issues in this area. The latter approach is reflected in the Multipeer Taxonomy which was initially developed within the ECFF project of SC6/WG4 and then passed on as a liaison contribution to SC21 [38]. The aim of the Multipeer Taxonomy is to provide a layer independent terminology for group communication and a framework for the development of multipeer services. It precisely defines the main concepts required in group communication such as group, group membership, population characteristic and different multipeer topologies. This document was absorbed into the Multi Peer Communication Architecture developed by ITU-T SG7 and ISO SC21 [40].

Various topological concepts have been proposed within SC6 to solve the problem of many-to-many communication. An initial proposal to leave it to the application to deal with this problem and to provide only multicast ($1{:}N$) communication was rejected at a very early stage. To provide flexible support for many-to-many communication without introducing ($N{:}N$) connections, the concept of *group associations* was proposed [47]. With this concept connections with basic unicast and multicast topologies are grouped together. However, they are still visible to the service user as separate entities. Each of them has its own QoS and *active group integrity* (AGI) conditions attached to it, which specify the number and/or identity of mandatory receivers. Additional conditions on the *group association integrity* (IC) and the topology integrity (ATI) specify under which circumstances a group association is considered successful. In contrast to this is the proposal of an unrestricted topology where all active group members can, in general, send and receive information at the same time over the same connection [41], i.e. the basic topology is ($M{:}N$). Here, integrity conditions can be defined on the number and identity of senders and receivers. Although it is proposed that QoS should be specified for this kind of topology as well, no description of how this could be done is given. Message ordering is specified by both proposals in the same way. Ordering semantics considered are *no-ordering, local ordering* (i.e. source ordering), *partial ordering* (ordering according to user rules) *causal ordering* (i.e. potential causal ordering) and *total ordering*. Recent discussions centre around the $M{:}N$ issue since (although it is widely accepted that ($M{:}N$) is the most general topology) it has not been possible to define how QoS could be specified and negotiated for such a topology. Another point of discussion within SC6 and SG7 is whether or not the distinction between connectionless and connection-oriented in a multipeer environment is still a useful distinction.

In the most recent version of the *Enhanced Communications Transport Service* (ECTSv2.0) three different topologies are defined: simplex Transport Connection (TC) ($1{:}N$), duplex TC ($1{:}N$, $N{:}1$) and N-plex TC ($N{:}N$). To allow multiple TCs per group, all TCs end at one TSAP (Transport Service End Point). This concept is akin to the one initially proposed for group associations. A distinction is made between connectionless and connection-oriented transfer modes [36]. The problem of QoS in an N-plex conversation is still unsolved. Different solutions have so far been proposed but none of them is general enough for 'true' N-plex. All proposals would restrict this topology considerably if used together with QoS. Before the last ISO meeting in Brazil (December 1995) an initial draft of the Enhanced Communications Transport Protocol (ECTP) was produced proposing some initial concepts and protocol mechanisms to support ECTS [39].

Apart from the unsolved problem of QoS for many-to-many communication, basic multicast ($1{:}N$) communication also has an impact on QoS. The problem is that different receivers might encounter different delay, delay jitter and error rates. Further they might have different quality requirements for audio and video, for example. To solve this problem it was proposed that some QoS parameters should be viewed as *receiver-selected* (namely delay, delay jitter and error rate) whereas throughput would be a *connection-wide parameter* and the same for all participants [46]. However, with emerging filtering mechanisms [66] throughput can be different for each receiver, i.e. in certain cases it would not be a connection-wide parameter. On the other hand, certain applications might require that delay, jitter or error-rate are bounded for all participants [66]. The current position adopted by ISO is that, depending on application requirements, QoS parameters in a multicast communication can be connection-wide or receiver-selected [35].

6. Group communication research

In recent years the research community has become more and more interested in the area of group communication. The problem of multimedia group communication is tackled by various research groups at different universities and research institutes. The context in which the research is carried out differs from project to project, depending on overall project goals and objectives of the organization. Since group communication itself is a multifarious problem there are various different ways to address issues related to group communication. We are aware of several groups currently working on different aspects of group communication (e.g. [7, 29, 47, 57]). In this section we introduce some of the work relevant for multimedia group communication concentrating on research dealing with resource reservation aspects and end-to-end communication services. A full discussion of the research currently carried out in this field is beyond the scope of this paper.

Research in the ***Tenet Group*** at ICSI, Berkeley focuses on the design and development of real-time communication services and network support for continuous media applications. A suite of protocols which provide QoS guarantees for real-time packet switching networks has been designed and implemented. Earlier work was concerned with *channel groups* and *half-duplex real-time channels*. Channel groups are abstractions by which the user can describe relationships between channels to the network [25]. These relationships would for instance allow resource sharing among related channels, which could be considered during routing, indicate establishment relationships, etc. Multipoint-to-multipoint connections can also be emulated using channel groups. A further concept introduced in this context is that of *target sets*. A target set is a set of tuples indicating the address and QoS specification of every member in the set. Senders always establish multicast channels to a target set. An (M:N) semantic can be achieved if multiple senders address the same target set [23].

The concept of *half-duplex real-time channel* for multiparty interactive multimedia applications is discussed in [59]. This abstraction exploits the fact that in applications such as lectures or seminars usually only one participant is allowed to speak at a time. Hence it is sufficient to reserve resources for only one communication channel, if the network can offer guarantees in both directions of the same multicast connection. Application level mechanisms like floor passing ensure that there is only one sender at a time. More recent work concentrates on *resource sharing*. With resource sharing, related connections can share resources in a controlled manner. In this case a network client specifies how traffic from related connections should be multiplexed. The client states the maximum aggregated resource requirements for the group and actually indicates to the network when to use the group rather than individual specifications. During admission of a new participant the group resource allocation is used. There is no admission test needed to admit additional members after the sharing threshold is reached [24]. *Multicast real-time channels* are now part of the Tenet Suite 2. The design of Suite 2 is still connection oriented; most of the changes for multiparty communication affect channel establishment. At the moment, channels are established from the source to all destinations. During the communication, multicast packets have to be copied in the intermediate nodes and the resource sharing mechanism requires a modification of the scheduling mechanism [6].

The ***Multipeer Broadband Transport Service*** is a transport protocol developed within the European RACE research programme. It was especially designed for distributed multimedia systems and is considering the special needs of group communication applications [28]. The two main problems addressed in this context are multipeer data transmission and relations between different data types (e.g. audio and video) transmitted over parallel transport connections. The provision of a variable multipoint transmission topology, QoS provision and negotiation, delivery semantics, voting and notification are the main design aims of the protocol. Apart from the usual simplex and duplex unicast transport connections, different formal types of *multipoint transport connections* (MP-TC) are introduced. The basic element of any multipoint transmission topology is a *connection*. Different, related connections are grouped together into *Transport Calls* (TL) which are jointly addressed at a *Transport Call End Point* (TLEP). Figure 4 shows how the different concepts are related. The transport protocol is fully specified including a list of service primitives, time sequence diagrams and state transition diagrams. Apart from the transport protocol a group management service to manage user groups, their members and properties is defined. The problem of ordering is not addressed in this context. However, since the service is connection oriented, source ordering is provided implicitly.

A Multi-user Communication Service supporting multimedia applications is proposed in [27]. This service

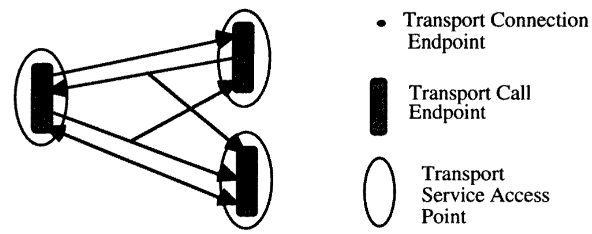

Figure 4. Multipoint transport topology (cf [28]).

offers basic and general facilities to support communication among multiple participants. A *call model* is introduced in which a call is a dynamic association between service users and the service provider. Different functions and service elements to operate on calls are defined. Apart from data exchange, other interactions such as synchronization between streams are supported. Service elements to establish a call, to add users, to add media, to remove users and to release a call are provided. Further service elements to attach and detach users to *media* are also specified. Communication can only take place among users who are attached to a medium. A generic service element supports negotiation to change the state of a call. Users can send, receive, and both send and receive data. One user is the owner of a call; they are a central instance and always involved in negotiations. Additional users have to be invited to join an on-going call. All service elements apart from 'detach media', 'remove media', 'remove user' and 'release call' are confirmed. As soon as the service provider finds out that a service request will be successful a confirm is sent to the initiator and all involved users who have already responded. Notification primitives are sent to all the users affected by state changes but who are not involved in the decision process. This is an extension of the standard confirmed and unconfirmed service elements in the OSI-RM. The service is defined independent of any implementation. It merely defines required capabilities of the communication system to support multimedia group applications.

The ***GCommS*** project at Lancaster University addresses mainly end-to-end aspects of multipeer communication [48]. The research can be divided into three closely related areas, namely end-to-end services and protocols, multipeer QoS, and scalability and reliability in multicast communication. The major components of the GCommS architecture are a multicast service (*M-Connection Service*), a multipeer service (*N2N-Connection Service*), and a communication manager (*GDC-Manager*) to provide virtual multipeer connection. The M-Connection Service is fully specified and has been mapped onto the XTP 4.0 protocol mechanism to prove its feasibility [52]. A number of tests and simulations have been carried out to gain experience with the service and to find the most suitable protocol architecture to base it on. Protocols currently being evaluated for this and the two other services are XTP 4.0 and an IP based protocol architecture using UDP, IPv6, RSVP and possibly RTP/RTCP.

The main issue of the GCommS work on multipeer QoS is the support of QoS in a heterogeneous group environment. A multipeer QoS model to support receivers with different quality requirements and capabilities has been specified. In addition, filtering mechanisms have been implemented to provide different data qualities for different receivers on a multicast stream [22]. This allows the service to provide each receiver with the quality it can handle while still utilizing the advantages of multicast. Algorithms to provide data transfer reliability in multicast communication have also been developed. Different degrees of reliability ranging from full reliability to probabilistic reliability can be supported by these algorithms. The research on communication and reliability issues is carried out in close cooperation with researchers at the Laboratoire MASI, Université Pierre et Marie Curie, Paris. The project also has links to the European COST 237 project concerned with multimedia telecommunication services, and other European research institutions.

The BERKOM (Berliner Kommunikationssystem) project was a joint Broadband ISDN (B-ISDN) trial project between various different computer manufacturers, research institutions and universities. One of the three working areas in the second phase of the project was the ***Multimedia Collaboration Service*** (MMC) which is now available as a commercial product from HP. MMC supports joint working in a distributed environment, i.e. it allows one to share applications and provides audio-visual conferencing tools. The different components of MMC are a *Conference Manager* (CM) to administrate and run conferences, *Conference Interface Agents* (CIAs) which serve as user interfaces, a *Conference Directory* (CD) to store information about user groups and conferences, an *Application Sharing Component* (ASC) for joint viewing and accessing of applications, an *Audiovisual Manager* (AVM) and the *Audiovisual Component* (AVC) [1]. The AVM offers a service to share audio and video data; it allows one to establish multiple streams between all conference participants and between multimedia applications and all conference participants for sharing audio and video data. All relations in the audio-visual communication system are controlled by the AVM [45]. The AVC manages all endpoints of data streams in the end systems. Communication streams are multicast streams ($1{:}N$). There is exactly one AVC per end system. The *Audiovisual Data Exchange Protocol* (AVXP) is used for audio and video data exchange. It resides on top of the transport layer. The *Audiovisual Control Protocol* (AVCP) is executed between the AVM and the CM. The CM opens one ($N{:}N$) group in every conference for the exchange of audio and video data between all participants and an additional ($1{:}N$) group for each shared audio-visual application. The *Source and Sink Control Protocol* (SSCP) is executed between the AVM and the AVCs. Requests from the AVM to the AVCs and event reports from AVCs to the AVM are exchanged over SSCP [45].

MICE (*Multimedia Integrated Conferencing for Europe*) was a project between various partners from six European countries and is now followed by MERCI. Its aim was to provide a service to enable internetworking in a heterogeneous environment via multimedia conferencing technology [26, 43]. The MICE technology allows conferencing between participants located in conference rooms which can accommodate up to ten people, and at workstations. The services used are multi-way shared workspace, multi-way video, multi-party voice, and multiplexing and conference management. Audio and video data are treated independently because of the higher delay sensitivity and smaller bandwidth requirements of audio. An important element in the MICE architecture is the *Conference Multiplexing and Management Centre* (CMMC). This provides gateway functions to allow communication between users at different networks with different facilities, e.g. ISDN, EBONE, EMPB. Further, it provides a relay between different coding standards, multiplexes audio and video, reserves, controls and manages resources, ensures multicast of shared workspace data, etc [26, 44]. A conference control system based on the *CAR* conference control system offers functions to create, delete, join and leave a conference, floor control and control of conference video channels *inter alia*. Practical experience with MICE was gathered during various conferences and a series of seminars for researchers and students. Seminars were multicast using MBONE [54].

7. Conclusion

A large number of existing and future interactive multimedia systems will operate in a distributed group environment. Different forms of cooperation and communication between remote partners require dynamic, flexible and efficient communications support. These applications place particularly stringent requirements on the underlying communication system in terms of both quantities of resource and service functionality.

At the lower layers of the communication architecture, support for multicast is offered by a number of data link layer protocols, though QoS provisions offered by the most common protocols (namely ethernet and token ring) are insufficient or entirely missing. The latest developments in the Internet try to accommodate multimedia and group communication requirements. Protocols such as RSVP and RTP are a step in this direction. IPv6 provides better support for multicasting and data flows such as multimedia data streams. However, many IP networks can still not cope with the bandwidth requirements of digital audio and video. MBONE is increasingly being used to transmit audio/video conferences or other special events. However, the quality of the transmitted data is often low and does not satisfy every user.

The high resource requirements of multimedia group applications require careful management of all involved system resources. Though emerging high speed networks offer far better support for multimedia communication, the requirements of multimedia group applications might exceed even their capabilities. On the other hand there is also a large, so far mainly unutilized, potential to save resources through the exploitation of mechanisms like multicast and resource sharing. Proposals for mixing filters in RSVP are a step towards exploiting the characteristics of group communication for a better resource utilization. However, this is not the only motivation behind resource sharing: resource sharing is also an application requirement since it reflects the relationship between individual data streams belonging to the same group communication. There are numerous open issues and problems at the data link and network layer which still remain to be addressed. However, available support in existing protocols and research in this area give a clear indication of the main thrust of developments here; i.e. multicast communication and multipeer QoS will be provided mainly in the lower layers. In addition, a kind of multipeer (N:N) communication can be fairly easily supported in a packet switched network environment.

One of the biggest problems faced in the design of end-to-end protocols and services is where and how to deal with the additional complexity. For instance, where and how can it best be determined if data transmission is successful or not; or, where and how can multiple connections be managed most efficiently. At the moment this is very often left to the application. Transport protocols such as XTP 4.0 offer only (1:N) multicast. Distributed system platforms and toolkits provide hardly any group support, and then only for special purposes such as the support of distributed file systems, distributed data bases or fault tolerant applications. Research on end-to-end services and protocols indicates that the complexity of the management of multiple connections should be handled at the transport layer. The transport layer is also traditionally concerned with error detection and recovery. In a group environment this includes multipeer integrity, the provision of data transfer reliability to multiple receivers and message ordering in the multiple sender case. Further, the transport service is also the interface at which system QoS is mapped onto communication QoS. On the other hand, at the transport layer no concern should be given to application policies, e.g. join policies or participants' roles. Here only mechanisms that allow these policies to efficiently be implemented above the transport layer are necessary. Current protocol proposals very often mix these issues and become therefore very specific and overly complex. Moreover, they are often defined in an abstract way, thus their practical relevance is limited.

Above the transport layer a variety of different services are required to support different types of group application. Group management is a generic task which deals not only with membership administration but also with providing a directory service for group and membership information. Examples of other required services at this level are conference control services or services providing synchronization. BERKOM's MMC and the MICE project deal with a whole range of services necessary for multimedia conferencing.

All the addressed issues have to be taken into account in the design of new protocols and the development of standards. Whatever solution is favoured for a particular problem, deliberately simple or more comprehensive, a decision should never be made without regard to the overall system architecture. This is the only way to deal efficiently with the high complexity imposed by multipeer communication.

Acknowledgments

For research support the authors are indebted to the Engineering and Physical Sciences Research Council of Great Britain. The work on this project is funded by the EPSRC project GR/J47804 on Multimedia Group Communications Services.

References

[1] Altenhofen M, Dittrich J, Hammerschmidt R, Käppner T, Kruschel C, Kückes A and Steinig T 1993 The BERKOM Multimedia Collaboration Service *Proc. 1st ACM Multimedia Conference (Los Angeles, 1993)* (New York: ACM Press) pp 457–63

[2] ANSA 1989 *ANSA Reference Manual, Release 01.00* (Cambridge: APM Cambridge Ltd)

[3] ATM-Forum 1994 *ATM User-Network Interface, Ver. 3.1 Specification*

[4] Ballardie A J, Francis P F and Crowcroft J 1993 Core based trees—an architecture for scalable inter-domain multicast routing *Proc. ACM Sigcom'93* vol 2 (New York: ACM Press) pp 85–95

[5] Beaumont D and Boucher J 1994 ITU-T T12x recommendations and their relevance to an API for ATM networks *Technical Report* BT-Laboratories, Ipswich

[6] Bettati R, Ferrari D, Gupta A, Heffner W, Howe W, Moran M, Nguyen Q and Yavatkar R 1995 Connection establishment for multi-party real-time communication *Proc. 5th Int. Workshop on Network and Operating System Support for Digital Audio and Video (Durham, NH, 1995)* (Heidelberg: Springer) pp 255–65

[7] Biersack E and Nonnenmacher J 1995 WAVE: A new multicast routing algorithm for static and dynamic multicast groups *Proc. 5th Int. Workshop on Network and Operating System Support for Digital Audio and Video, NOSSDAV '95 (Durham, NH, 1995)* (Heidelberg: Springer) pp 243–65

[8] Birman K P and Joseph T A 1987 Reliable communication in the presence of failures *ACM Trans. Comput. Syst.* **5** 47–76

[9] Birman K P, Schiper and Stephenons 1993 Lightweight causal and atomic group multicast *ACM Trans. Comput. Syst.* **9** 272–314

[10] Boucher J 1994 *Draft Recommendation T.120: Transmission Protocols for Multimedia Data* Geneva

[11] Bowles B A and Teale R 1994 Communications services in support of collaborative health care *BT Technol. J.* **12** 29–44

[12] Braden R, Zhang L, Estrin D, Herzog S and Jamin S 1995 Resource ReSerVation (RSVP)—Version 1 Functional Specification *Internet Draft* IETF

[13] Braudes R and Zabele S 1993 Requirements for multicast protocols *RFC: 1458* Internet Networking Group, TASC

[14] Cheverst K, Friday A, Davies N and Blair G S 1994 Position statement on producing a group collaboration based application to operate in a weakly connected environment *Internal Report* Lancaster University

[15] Clark W J and Boucher J 1994 Multipoint communications—the key to groupworking *BT Technol. J.* **12** 72–80

[16] Comer D E 1991 *Internetworking with TCP/IP; Principles and Architecture* vol I (Englewood Cliffs, NJ: Prentice Hall)

[17] Conta A and Deering S 1995 ICMP for the Internet Protocol version 6 (IPv6) *Internet-Draft, draft-ietf-ipngwg-icmp-01.txt* Network Working Group

[18] Davies N, Blair G S, Cherverst K and Friday A 1994 Supporting collaborative applications in a heterogeneous

mobile environment *Internal Report, MPG-94-18* Computing Department, Lancaster University
[19] Deering S, Estrin D, Farinacci D, Jacobson V, Liu C-G and Wei L 1994 An architecture for wide-area multicast routing *Proc. SIGCOMM '94 (London)* (New York: ACM Press) pp 126–35
[20] Delgrossi L, Herrtwich R G, Vogt C and Wolf L 1993 Reservation protocols for Internetworks: a comparison of ST-II and RSVP *Proc. 4th Int. Workshop on Network and Operating Systems Support for Digital Audio and Video (Lancaster, 1993)* (Heidelberg: Springer) pp 199–208
[21] Eriksson H 1994 MBONE: the multicast backbone *Commun. ACM* **37** (8) 54–60
[22] Garcia F, Hutchison D, Mauthe A and Yeadon N 1996 QoS support for distributed multimedia communications *Proc. Int. Conf. in Distributed Processing (Dresden)* (New York: Chapman and Hall)
[23] Gupta A, Heffner W, Moran M and Szyperski C 1993 Network support for realtime multiparty applications *Proc. 4th Int. Workshop on Network and Operating Systems Support for Digital Audio and Video (Lancaster, 1993)* pp 37–40
[24] Gupta A, W Howe, Moran M and Nguyen Q 1995 Resource sharing for multi-party real-time communication *Proc. IEEE INFOCOM '95* http://tenet.berkeley.edu/tenet-papers.html
[25] Gupta A and Moran M 1993 Channel groups: a unifying abstraction for specifying inter-stream relationships *Technical Report, TR-93-015* Tenet Group, Computer Science Division, University of California at Berkeley and International Computer Science Institute, Berkeley
[26] Handley M J, Kirstein P T and Sasse M A 1993 Multimedia Integrated Conferencing for European Researchers (MICE): Piloting activities and the Conference Management and Multiplexing Centre *Comput. Networks ISDN Syst.* **26** 275–90
[27] Heijenk G J, Hou X and Niemegeers I G 1994 Communication systems supporting multimedia multi-user applications *IEEE Network* **8** 34–44
[28] Henckel L 1993 Multimedia communication platform: specification of the enhanced broadband transport service (M26a) *Technical Report* GMD-Germany
[29] Herrtwich R G 1992 The HeiProjects: Support for distributed multimedia applications *Technical Report, 43.9206* IBM European Networking Center, Heidelberg
[30] Hinden R 1994 Internet Protocol, Version 6 (IPv6) Specification *Internet-Draft, draft-hinden-ipng-ipv6-spec-00.txt* Sun Microsystems, Inc.
[31] Hinden R 1994 IP next generation addressing architecture *Internet Draft, draft-hinden-ipng-addr-00.txt* Sun Microsystems, Inc.
[32] Hinden R M 1994 IP next generation overview *Internet-Draft, draft-hinden-ipng-overview-00.txt* Sun Microsystems, Inc.
[33] Ishii H, Kobayashi M and Arita K 1994 Iterative design of seamless collaboration media *Commun. ACM* **37** (8) 83–97
[34] Ishii H and Ohkubo M 1990 Design of team work station: a real-time shared workspace fusing desktops and computer screens *Multi-User Interfaces and Applications* ed S Gibbs and A A Stuart (Amsterdam: Elsevier)
[35] ISO/ICE_JTC1/SC21 1995 *Information Technology—Quality of Service—Guide to Methods and Mechanisms*
[36] ISO/IEC_JTC1/SC6 1995 *Second Draft ECTS Definition v2.0* Geneva
[37] ISO/IEC_JTC1/SC6-N9161/III 1995 *Draft on Multipeer Taxonomy*
[38] ISO/IEC_JTC1/SC6-N9161/IV 1995 *Draft Text on the subject of 'Multi-peer Taxonomy* Beppu
[39] ISO/IEC_JTC1/SC6/WG4 1995 *Initial Draft of ETCP* Korea
[40] ITU-T/SG7 1995 *First Draft of ISO/IEC 7498—5 Multi Peer Communication Architecture*
[41] ITU-T/SG7 1995 *Liaison Statement from ITU-T SG 7 Regarding 'Multicast Transport Service (X.tms)*
[42] ITU-TS 1993 *B-ISDN ATM Functional Characteristics*
[43] Kirstein P, Handley M, Sasse M A and Clayman S 1995 Recent activities of the MICE conferencing Project *Proc. INET'95* (Reston, VA: Internet Society) pp DCA 1–6
[44] Kirstein P T, Handley M J and Sasse M A 1993 Piloting of multimedia integrated communication for European researchers (MICE) *Proc. INET '93* (Reston, VA: Internet Society) pp DCA 1–12
[45] Kückes A 1993 BERKOM multimedia collaboration service, AVC semantics *Project Report, Draft 0.1* The BERKOM Project
[46] Mathy L and Bonaventure O 1994 QoS negotiation for multicast communications *Proc. Multimedia Transport and Teleservices; Int. COST 237 Workshop (Vienna)* (Berlin: Springer) pp 199–218
[47] Mathy L, Leduc G, Bonaventure O and Danthine A 1994 A framework for group communications *Proc. 3rd Int. Broadband Islands Conference (Hamburg)* (Amsterdam: Elsevier) pp 167–74
[48] Mauthe A, Hutchison D, Coulson G and Namuye S 1994 From requirements to services: group communication support for distributed multimedia systems *Proc. 2nd Int. Workshop, IWACA '94 (Heidelberg)* (Berlin: Springer) pp 266–77
[49] Mauthe A, Hutchison D, Coulson G and Namuye S 1995 From requirements to services: a study on group communication support for distributed multimedia systems *Technical Report, MPG-95–10* Computing Department, Lancaster University
[50] Mitzel D J, Estrin D, Shenker S and Zhang L 1994 An architectural comparison of ST-II and RSVP *Technical Report* Computer Science Department, University of Southern California, Los Angeles
[51] Ramanathan R 1995 Multicast support for Nimrod: requirements and solution approaches *Internet-Draft, draft-ietf-nimrod-multicast-01.ps* BBN Systems and Technologies
[52] Rezende J, Mauthe A, Hutchison D and Fdida S 1995 M-connection service: a multicast service for distributed multimedia applications *Proc. 2nd COST 237 Workshop on Teleservices and Multimedia Communication (Copenhagen)* pp 38–58
[53] Sanders R T, Dempsey B J and Weaver A C 1992 *XTP: The Xpress Transfer Protocol* (Reading: Addison Wesley)
[54] Sasse M A, Bilting U, Schulz C-D and Turletti T 1994 Remote seminars through multimedia conferencing: experiences from the MICE project *Proc. INET '94* (Reston, VA: Internet Society) pp 251–59
[55] Schulzrinne H 1993 Issues in designing a transport protocol for audio and video conferences and other multiparticipant real-time applications *Internet-Draft* Internet Engineering Task Force; Audio–Video Working Group
[56] Schulzrinne H, Casner S, Frederick R and Jacobson V 1996 RTP: a transport protocol for real-time applications *Request for Comments, RFC-1889* IETF
[57] Simon R, Znati T and Sclabassi R 1994 Group communication in distributed multimedia systems *Proc. IEEE—Distributed Multimedia Systems* pp 294–301
[58] Steinmetz R and Nahrstedt K 1995 *Multimedia—Computing, Communications and Applications* (New York: Prentice Hall)
[59] Szyperski C and Ventre G 1994 Efficient support for multiparty communication *Proc. Int. COST237 Workshop (Vienna)* (Berlin: Springer) pp 185–98
[60] Takemura H and Kishino F 1992 Cooperative work

environment using virtual workspace *Proc. ACM Conf. on Computer-Supported Cooperative Work (Toronto, 1992)* (New York: ACM Press) pp 226–232

[61] Talpade and Ammar 1995 SCE: an architecture or reliable multicasting *Internet Engineering Task Force Internet-Draft* College of Computing, Georgia Institute of Technology

[62] Topolcic C 1990 Experimental Internet Stream Protocol, Version 2 (St-II) *RFC 1190* BBN

[63] Williams N, Blair G S, Coulson G, Davies N and Rodden T 1994 The impact of distributed multimedia systems on computer support for cooperative work *Technical Report* Computing Department, Lancaster University

[64] XTP 1992 *XTP Protocol Definition* Santa Barbara

[65] XTP-Forum 1995 *Xpress transport protocol specification* XTP Revision 4.0, Santa Barbara, USA

[66] Yeadon N, Mauthe A, Garcia F and Hutchison D 1996 QoS filters: addressing the heterogeneity gap *Proc. European Workshop IDMS '96 (Berlin)* (Heidelberg: Springer) pp 227–44

[67] Zhang L, Braden R, Estrin D, Herzog S and Jamin S 1994 Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification *Internet Draft*

[68] Zhong W D and Yukimatsu K 1994 Design requirements and architectures for multicast ATM switching *IEICE Trans Commun.* **E77-B** (11)