

Systems for the Nineties — Distributed Multimedia Systems

Sape J. Mullender

University of Twente, Faculty of Computer Science
P.O. Box 217, 7500 AE Enschede, The Netherlands

1 Introduction

We live at the dawn of the information age. The capabilities of computers to store and look up information are only just beginning to be exploited. As little as ten years ago, practically all the information stored in computers was entered and retrieved in the form of text. Today, we are just starting to use other means of communicating information between people and machines — computers can now scan images, they can record sound, they can produce synthesized speech, and they can show two- and three-dimensional images of spatial data.

The realization that we are still at the beginning of the information age comes when we notice the vast difference between the way in which people interact with each other and the way in which people can interact with (or through) machines. When people communicate, they tend to use speech, gestures, touch, even smell; they draw pictures on the white board, they use text, pictures, photos, graphs, sometimes even video presentations. Interpersonal communication is truly multimedia communication in that it makes use of all our senses.

Compared to this, the quality of our interactions with computers is still very poor. We have a keyboard and a mouse as input devices, and a screen and a nasty beeper as output devices. More and more of the information accessible to us is stored in computers, but the quality of the presentation and the ease of accessing the information is often poorer than that of information in a book.

Multimedia research is exploring the ways in which human/computer interaction can be improved by making use of text, images, graphs, animation, sound, speech, moving images (video). The list can be extended — there are experiments with moving images that appear to move through space (*virtual reality*) and experiments with computer input through gestures (e.g., *data glove*) and output through touch (*tactile feedback*).

Architects are today starting to use three-dimensional rendering techniques (*ray tracing* and *radiosity*) so that they can see what a building looks like before it is built. Scientists are beginning to use two- and three-dimensional colour images to help them visualize scientific data. Synthesized voices warn pilots in such a way that they don't have to take their eyes off the runway — and if the aircraft is still in the clouds, a *heads-up display* can project a synthetic image of the runway that appears to be in place of its actual counterpart.

New types of information must be stored and new ways of presenting information must be explored. The challenge of multimedia research is the design of an *integrated* and *distributed*

multimedia environment. Integration means that different media and different types of information can be manipulated using the same (or very similar) interfaces, that they can be stored in *the* file system, rather than different, specialized ones, and that they can be converted from one type to another.

Thus, a sequence of renderings of a changing three-dimensional scene must be convertible into a video sequence and, similarly, it must be possible to take individual frames out of a video sequence and use them as pictures in a printed document. The interfaces for doing this must be natural and the conversions from one format to another must be automatic, triggered by the type of usage.

Multimedia systems must also be distributed. A user must be able to do the same things anywhere in the system, and access information stored anywhere in the system. The system is responsible for finding the requested data. The user need not know. Only when multimedia systems are distributed in this sense is it possible to truly share information across large geographical areas.

Both integration and distribution still require lots of study. Multimedia projects to date are mostly characterized by the research of multiple media one at a time (e.g., how to combine video communication with a computer system) with relatively little attempt to provide identical treatment of the overlapping properties of different media and the usage of multiple media in some mutual relationship (e.g., browsing through a video show with sound and animation computed on the fly)

The Huygens project at the University of Twente is just starting to investigate some of the design issues in distributed multimedia systems. Because of the expertise in the group and because of its limited size, the project is confined to addressing *systems issues* of distributed multimedia system design: solving the technical problems of getting the right data in the right places with sufficient reliability.

The next section describes the important properties distributed multimedia systems must have. Section 3 discusses research issues that need to be resolved in order to build infrastructural support for these systems and the scientific disciplines needed to do this.

2 The Challenge of Multimedia

We envision that the information systems of the future will provide much better ways for people to store and retrieve information. Currently, computer systems present and store information primarily in the form of text. Documents, such as this one, now contain structural information aiding the typesetting and sometimes very simple line graphics. Pictorial information is hardly handled and moving pictures can only be seen in a few research laboratories.

This will change. The information systems of the future will have:

- **Global information sharing and location transparency.** Almost all advanced development is done in teams. Often, these teams are spread out over several labs, sometimes in different countries. Working together in a team requires access to the same information everywhere, completely up-to-date, in the same way, independent of location. Multimedia systems must be *distributed systems*.
- **Interpersonal communication over large distances.** When communication between people involves information shared via a distributed system, an appropriate communica-

tion infrastructure must involve the distributed system. Ideally, people, separated by a large geographical distance, should be able to communicate much as if they are in the same room, in front of the same computer screen. This allows them to talk, see each other's facial expressions and gestures and to manipulate the information stored by the system.

With the help of audio and video devices, multimedia systems can make this possible; they can create a communication channel through which people can communicate by seeing and hearing each other, and by manipulating information on each other's screens. Multimedia systems can make international collaboration projects an order of magnitude more efficient, just like the advent of phone has done in the first half of the century and — to a much lesser extent — the advent of facsimile has done recently.

- **Interfaces usable for the layman.** Computer systems are ever more used by people who are not computer professionals. Since, at the same time, man-machine interfaces become more and more sophisticated, a great deal of attention must be paid to making interfaces both simple, logical and natural. These interfaces should be very easy to learn, or, better still, not require any learning. Natural language will play an important rôle, as well as “natural” graphical interfaces, mouse interfaces, keyboard interfaces, etc.
- **Long-term information storage.** Advanced information systems store vast amounts of very valuable information. Many kinds of information will remain important for generations. We may expect that information currently kept in books in libraries will, one day, be stored digitally in electronic libraries. Before librarians will trust digital computer-readable storage, however, ways must be found that guarantee the integrity of the data for many hundreds of years. This is two orders of magnitude longer than current-technology media last.
- **Reliability and availability.** Given that information systems play an increasingly important rôle in the day-to-day operation of any organization, it must be obvious that these systems have to be very reliable indeed. A system is reliable when it does nothing wrong: it loses no data, it corrupts no data, it makes no errors. It is available if it also works when you want it to. Since storage media are subject to failures, reliability can only be achieved by replicating data; availability can only be achieved by replicating the access path to the data, i.e., by replicating the servers as well.
- **Full-speed and full-screen video plus high-fidelity stereophonic sound.** Video and sound form essential ingredients for interpersonal communication over large distances, but have also proved their worth in teaching, demonstrations, etc. Video and sound must be manipulatable as first-class objects — one must be able to store them and edit them, much like one stores and edits text.
- **Animation of two- and three-dimensional data.** Most simulations involve data representing space- and time-dependent values. It is very useful to be able to see these values as they change with time. Similarly, the understanding of complex three-dimensional structures can be increased by being able to view them from every angle. For example, architects can look at buildings and show them to clients before they have been built, and high-energy physicists can look at three-dimensional renderings of actual and simulated particle scatterings resulting from proton-antiproton collisions.

- **Security.** Information systems that form the basis of valuable collaborative work must be secure. This means that information stored in the system can be protected against unauthorized access. To this end, there must be a system-wide authentication system, a global security architecture, and an auditing system. At the same time, distributed systems spanning organizational and national boundaries are bound to have different security policies in different parts of the system. One organization's (lack of a) security policy should, however, never interfere with another's ability to implement a stricter or laxer one.

3 Systems Research

Multimedia systems research addresses the integration of input, storage, retrieval, and output of information in as wide a spectrum of forms as technology allows. Multimedia research is multi-disciplinary in character and encompasses research in distributed systems, telecommunication, data bases, real time, computer graphics, digital audio and video, natural language, user interfaces, psychology, education, electrical and mechanical engineering and more. Multimedia research should bring these disciplines together so that cross fertilization can occur and produce results that the individual disciplines could never bring about.

One can distinguish between research on multimedia *applications* and the *systems support* for multimedia. The applications are, of course, what multimedia systems are really all about, but the applications cannot be built without sufficient systems support. In this section, we discuss research issues in systems infrastructure necessary for multimedia applications.

3.1 Location Transparency

Multimedia systems will be distributed systems: multiple processing elements and multiple storage servers geographically distributed. An important use of multimedia systems will be for communication between people that are far apart.

Location transparency is not having to know where objects or users are. Objects have names and by uttering the name of an object, one obtains access to it. Also, different people in different organizations must be able use the same names to refer to the same objects. Global naming is the basis of global sharing.

This is not altogether trivial in environments where objects and users may migrate and where objects may be replicated for fault-tolerance and performance reasons. And it will become harder as portable computers using wireless network communication become common.

Location transparency requires system-wide agreement on naming, object invocation and security. Multimedia systems, therefore, must be based on *distributed systems* platforms and cannot merely be conceived as collections of autonomous multimedia workstations connected by fast networks and communications standards.

3.2 Storage

Storage subsystems are the heart of information systems. The storage requirements of distributed multimedia systems are extreme in a number of dimensions: size, speed, flexibility, and reliability.

Multimedia systems store information ranging from unstructured data, such as plain text, to highly structured, very large complex objects, such as data representing the shape of an aircraft wing. Stored objects vary in size from a few bytes (e.g., a mailbox name) to several terabytes (e.g., a complex cartographic object representing a detailed map of France).

The Hubble Space Telescope produces a gigabyte worth of compressed astronomical observation data every day and is expected to do so for the next fifteen years. This data will be used by astronomers all over the world and, if they can get access to it, they will use long-distance multimedia communication to discuss these observations. One cannot predict where the data will be used and it does not make sense to replicate the data everywhere, nor does it make sense to store all the data in just one place.

While processor speeds have increased dramatically during the last decade, storage devices have become only marginally faster. The increasing speed mismatch has triggered research in new storage structures for file systems and data base systems which rely much more heavily on *caching data in main memory*.

There is a trend towards designing storage systems for structured files which provide better handles on managing data representation differences on different machine architectures. At the same time, there is a trend in database design to leave the inflexible relational database structure and design more flexible complex-object systems. File-system research and data-base research are thus likely to meet each other halfway, providing an opportunity for comparing notes and learning from each other.

Multimedia information is highly structured and needs to be accessed in a type-dependent way. A video sequence, for instance, must be retrieved at a high and constant rate; frames retrieved too late are no longer useful. A hypertext document, in contrast, is retrieved in the order in which the user happens to browse through it, and need not always be retrieved in real time.

The Computer Science Faculty of the University of Twente is involved in two collaborative projects on storage system design; one is a SION-sponsored project with CWI, University of Amsterdam and Free University, called Starfish; the other is a project called Linnæus, partly sponsored by Hewlett-Packard, with HP Laboratories in Palo Alto, California. These two projects are closely related, although they address slightly different goals. Starfish aims at the design and implementation of a distributed *complex-object server*, a storage server intended for storage and retrieval of small and large structured objects, possibly containing pointers to other objects. Complex objects find increasing use in database applications where, at design time, the structure of the objects stored in it are not known. Linnæus aims at extremely large, reliable and fast storage systems that can grow to capacities of tens of terabytes, an order of magnitude larger than exists today. In both of these projects, the University of Twente forms the unifying link that allows the two projects to build on each other's results.

3.3 Communication

A range of communication styles and primitives must be supported in a distributed multimedia system. For *isochronous communication*, such as needed for sound and moving images, small end-to-end latency is required with little variation; error-free transmission is less important, and retransmissions are useless because they will arrive too late.

For client/server interaction, *remote-procedure call* mechanisms have proven their worth. Modern RPC mechanisms compile communication code directly from *interface-description*

languages, making the use of RPC as simple as that of ordinary procedure call. There is a direct relationship between the granularity of client/server interaction and the efficiency of local and remote RPC implementations. High-performance RPC is essential for high-performance distributed computing.

Replication is the key instrument for fault tolerance in distributed systems. A fault tolerant service uses multiple server processes and maintains *replicated state*. To maintain the consistency of this state, a number of *broadcast* mechanisms allow a set of server processes maintaining replicated state to observe events reliably and in the same order everywhere.

Designing separate communication infrastructures for isochronous data streams, remote-procedure call, and reliable causal broadcast is relatively straightforward, given the vast amounts of literature that explain all the issues and their solutions. But integrating them is a different matter entirely. Integration is not merely being able to use all three styles of communication simultaneously over the same set of wires — it means, for instance, that one must be able to combine RPC and causal broadcast in a way that preserves causality over RPC, and it means that the marshalling mechanisms for RPC, broadcast and isochronous communication must be similar if not identical.

RPC and broadcast, when used in a multimedia context, can have real-time requirements. An example is an RPC to a window manager that draws a frame on the screen at some precise moment. Scheduling a mixture of real-time and non-real-time, and a mixture of fixed-sized and variable-sized messages is also a challenge.

The Faculty of Computer Science at the University of Twente is currently setting up several related projects on communication in distributed systems. Collaboration is starting within the faculty itself between the distributed systems group and the open systems group which addresses multimedia protocol design, specification, and verification. The BROADCAST project is a collaborative project on a European scale for which funding will be sought in the ESPRIT Basic Research Action. This project will address management of replicated and distributed objects and data over (relatively slow and unreliable) wide-area networks. With CWI there is collaborative work on multimedia document definition and transmission.