Real-Time Telediagnosis of Radiological Images through an Asynchronous Transfer Mode network: The ARTeMeD Project

Raffaele Noro, Jean-Pierre Hubaux, Reto Meuli, Ricardo Norberto Laurini, and René Patthey

The ARTeMeD project aims to solve problems of interactivity and real-time in teleradiology. It integrates personal multimedia facilities and patient data access in a common platform that allows radiologists to collaborate from remote sites through a suitable communication support. ARTeMeD is based on asynchronous transfer mode (ATM) network technology and an optimized manipulation of medical image material. The ARTeMeD system opens interesting perspectives in the area of collaborative teleradiology.

KEY WORDS: client, server, ATM, graphical interface, real time

IN THIS article, we describe the ATM Real-Time Medical Diagnosis (ARTeMeD) project under way in Lausanne at the Telecommunications Laboratory of the EPFL and the University Hospital. In this project, we are exploring solutions in telemedicine[1] that allow interactive, real-time remote diagnosis both in pathology and in radiology. To contribute effectively in this area, we analysed existing systems and identified ways of improving their performance. The project involves studying how medical image processing maps onto a network technology. The fast transmission of high-resolution images, the capability of remote manipulation of region of interest (ROI) (that is a user-selected portion of an image with high perceptual quality or resolution), and live audio and video communication are addressed in ARTeMeD, as well as the synchronization requirements and the multiparty support for collaborative work.

The results of the project include an architecture, the ARTeMeD system, which englobes the aforementioned aspects. The key issues of ARTeMeD are the optimization of storage, access, and retrieval of medical images; the use of high-speed networking technology; and the deployment of suitable image compression algorithms.

This article is organized as follows: in the next two sections, we describe the architecture. We then discuss the organization and manipulation of medical image material. Then, we present our prototype implementation, performance measures, and a comparison to a pre-existing standard solution for teleradiology. Some conclusions and suggestions for further work are discussed as well.

SYSTEM ARCHITECTURE

The ARTeMeD project has evolved by taking into account the needs for real-time and interactivity in teleradiology. These needs have been analyzed in terms of system requirements and system design. The resulting architecture is an integration of desktop conferencing facilities and a new method for storage/retrieval of medical images over a suitable communication support (Fig 1). The communication support considered here is based on asynchronous transfer mode (ATM) network technology.[2] This particular subsystem provides two access points to the application level: one for the exchange of medical data and one for multimedia personal communication.

A typical scenario in which our system can be used is as follows: images are stored in a server and several radiologists want to access and browse the images on their screen. Collaboration of radiologists over the same subject is allowed; in this case, the same images are distributed by the server to the participants of the session. The conferencing subsystem interconnects all the participants. At a given time, only one radiologist has control over the images: his or her interactions with the server produce the images that are distributed. Such a situation is sketched in Fig 2 for the case of two radiologists.

It is worth noticing that the control of images flows from one client to the server (one-to-one connection), images flow from server to all the clients (one-to-many), and conferencing interconnects each client with each client (many-to-many).

The contribution of the ARTeMeD project is mainly in the medical part: We propose an innovative client-server approach for optimizing the re-

From the Telecommunication's Services Group TCOM, Ecole Polytechnique Fédérale de Lausanne; and the Centre Hospitalier Universitaire du Canton de Vaud, Institut de Radiologie, Institut de Pathologie, and Office Informatique, Lausanne, Switzerland.

Address reprint requests to Raffaele Noro, EPFL-DE-TCOM, 1015 Lausanne, Switzerland.

0897-1889/97/1003-1039$3.00/0

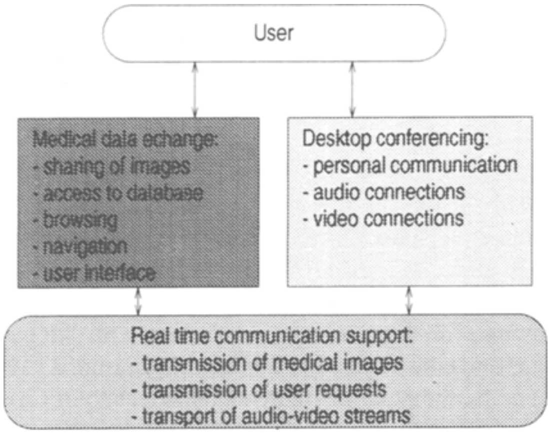

Fig 1. Integration of desktop conferencing and medical data exchange over communications support for ARTeMeD.

trieval and display operation times of medical high-resolution still images. ARTeMeD allows interactivity and real-time thanks to the following design features:

- Reduced amount of data to be processed, thanks to the compression and hierarchical organization of images
- Increased transfer rate for the communication support, thanks to the ATM network technology
- A simple graphical interface for the navigation of the medical images, thanks to the concept of user-defined RoI.

Some additional requirements also must be satisfied for the efficient cooperation among specialists. The latency of the system must be kept under control (i.e., the time between the generation of a request by the user and the display of the requested image on the users' screens) as well as the coordinated presentation of the same images in spatially separated sites. These two aspects often are referred to as *temporal* and *spatial synchronization,*[3,4] and are fundamental to allow interactivity in distributed environments. They are taken into account for the ARTeMeD project.

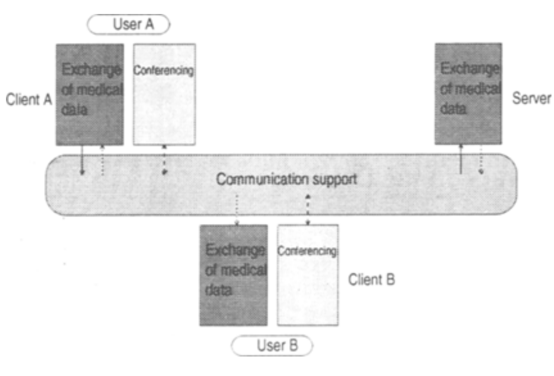

Fig 2. Server-radiologists and radiologist-radiologist interactions.

THE ARCHITECTURE COMPONENTS

Any ARTeMeD session is opened with one ARTeMeD server and one or more ARTeMeD clients. The client is expected to offer an interface to the radiologist, whereas the server is designed to inteface with the medical data stored in an attached database and also to perform some centralized activities (among others, monitoring of the participants and spatial synchronization tasks).

The basic components used to build both the server and the client of the ARTeMeD system are represented in Fig 3. The blocks labelled with 1, 2, 3, and 4 belong to the medical data exchange subsystem; the blocks 5, 6, 7, and 8 belong to the conferencing subsystem; and the blocks 9 and 10 constitute the communications subsystem.

Block 1: Decoding and Display

Each ARTeMeD client is equipped with this block. It is designed to get compressed image data from the network, to process this information, and to display the target image on the user's screen.

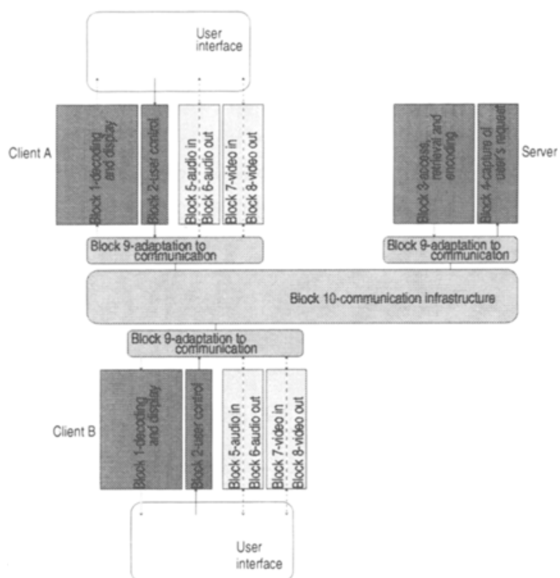

Fig 3. ARTeMeD system components.

Block 2: User Control and Graphical Interface

In the ARTEMeD client, this functional block provides the capability to the user to perform some remote actions on the server. The user actions (e.g., the definition of a RoI, the selection of a low-resolution image to be magnified, the modification of gray windows) are captured by this block and mapped onto messages to be carried over the communications support.

Block 3: User Commands Interpretation and Database Request

This block is situated in the ARTeMeD server. It transforms the incoming user message from the network (generated by Block 2) into a filter for the selection of images or a portion of them to be retrieved from the database.

Block 4: Access, Retrieval, and Encoding

The selection required by Block 3 is performed by the current block and the image data are generated in the format required by Block 1. Block 4 operates on the compressed images. It must supply output data that correspond to the specified compressed format, but implementation can be done in several independent ways.

Blocks 5, 6, 7, and 8: Compression and Rendering of audiovisual Streams

Audiovisual interconnections among participants are obtained by means of these blocks. They perform the two-way mapping between audio and video frames and network packets. Their functionalities are out of the scope of the ARTeMeD project, which simply integrates them beside the medical imaging part. It is worth noting that video communications can be optionally dropped from the setup with a reduced impact on the degree of interactivity of the system (if, for instance, a camera is not available or the audio quality is good enough or we can save bandwith/CPU for the medical exchange, and so on).

Block 9: Communication Interface and Adaptation

This block interfaces each user terminal with the communications support. It has two main tasks: first, it deals with etherogenity in the communications support (for instance, client and server can be connected to networks of different characteristics, such as ATM and Ethernet, and the slowest one dominates the performance of the communication), and second, it handles the paths along which data are sent (all the client-to-client, client-to-server, and server-to-client connections).

Block 10: Communications Infrastructure

This is one of the crucial components of the system. It must have a powerful and flexible communications infrastructure to support the information flow generated by such a system. High amounts of data need to be moved with strong temporal constraints. For this specific reason, ATM is, at the moment, the more suitable technology to do so. Nevertheless, in a local area network (LAN) environment (e.g., hospital network), other technologies (e.g., Ethernet or Fast-Ethernet) can be appropriate. As soon as we consider a more spread out system, however, it becomes prohibitive to obtain the same results with classic network technologies.

MEDICAL IMAGE STORAGE, ACCESS, AND RETRIEVAL IN ARTEMED SYSTEMS

Our approach to the manipulation of radiological images in ARTeMeD is organized into steps. The storage in the server's disk or database is performed off-line, directly from the acquisition device. The access and retrieval of the stored material is done on user request, and is performed in real-time during the operative session. It is important to note that the ARTeMeD system is based on a completely file transfer-free paradigm: The requested images do not need to be stored in the client's disk. Instead, they are processed by the decoding block and directly browsed on the user's screen. They take up memory, but no extra disk space. The images do not leave the server's disk, they are just displayed through a compression/decompression middlelayer onto the client's screen, giving the impression of manipulating a local copy. The remainder of this section describes the storage, access, and retrieval phases of our system; that is, the most original part of ARTeMeD.

Storage in the Database

The images coming from the acquisition device (computed tomography or magnetic resonance imaging) are pre-processed. They are organized in a hierarchy of resolution levels for the access and retrieval phases. Two (or possibly three) subsampling operations of the high-resolution images are performed, and the different levels of resolution are obtained (Fig 4). Then, for each level, all images

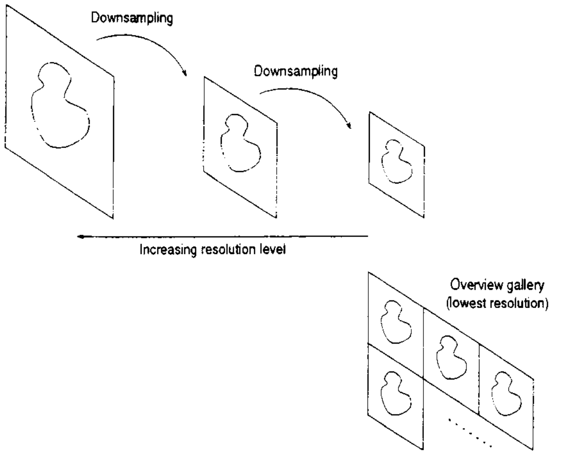

Fig 4. Scaling and composition operations in the database.

are disposed in a gallery. The galleries are logically splitted in elementary regions, which are the quanta of information for the access and retrieval phases. A compression algorithm is applied to the images in the hierarchy, reducing the amount of information to be transferred. The Motion Picture Expert Group (MPEG) has defined standards for the compression of audiovisual information. In our implementation, the MPEG-2 standard has been chosen. The elementary regions just mentioned are the slices of the MPEG algorithm.[5] Compressed slices are indexed to allow fast access to specified portions of the images.

Access of Target Images in the Database

The user can specify which image(s) and resolution he or she wants (eg, the coarsest gallery or overview gallery, a part of the semi-full resolution gallery, a single image at the highest resolution). Through the graphical interface, the user sends a request to the server in the form of a geometrical area of an image and the corresponding resolution level. This request acts like a filter over the slice index, and the selected slices are extracted and aligned in an MPEG-2 compliant stream for the transmission and decompression in the client's host.

Retrieval of Images: the User Interface

The decision of which images or parts of image to analyze is left to the specialist(s). The system itself offers a minimal amount of initial information for starting inspection (the overview gallery) and a graphical interface for navigating images and increasing through predefined steps the resolution of images. The radiologist selects the images(s) needed to do the inspection or diagnosis. By using the screen's pointer, the radiologist defines the portion of the gallery he or she wants to magnify (the RoI). The request then is sent to the server (a simple click on one button of the mouse). It should be noted that, as the pointer can be moved around in the gallery, the user can generate a stream of requests with a well-defined temporal behavior. In this case, the server response is a sequence of images with the corresponding temporal characteristics.

In this way, the retrieval is interactive because the user(s) can easily control which part of the radiological folder is more or less important. Also, the user can communicate with the other participants of the session. In addition, the retrieval is real-time because only a little amount of information flows from server to clients, and no intermediate storage operations are performed, only decompression and display.

The RoI is the one of the key aspects of the ARTeMeD system. It is taken into account for the organization of the database and the graphical user interface. At the same time, it improves the interactive performances of the remote diagnosis: The amount of information and the processing time consistently are reduced for each operation.

PROTOTYPING OF AN ARTEMED SYSTEM: IMPLEMENTATION AND PERFORMANCES

The prototype currently implemented at the Telecommunications Laboratory includes most of the features discussed here (see snapshot in Fig 5). The hardware-software platform is composed of:

- A communications subsystem: ATM boards and one ATM switch connect the clients and

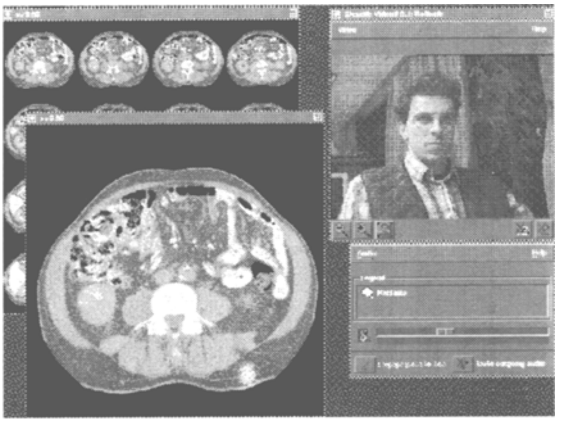

Fig 5. Sample screen from ARTeMeD system during use.

Table 1. Main Characteristics of ARTeMeD Systems

Allocated ATM link capacity: 1 to 1.5 Mb/sec
Image compression technique: MPEG-2
Levels of hierarchy: 3, that is full resolution (512 × 512 pixels), 4 and 16 times downsampled
Frame rate for interactive navigation: 5 to 10 frames/sec; transmission errors are cancelled by this refresh rate without bothering the user
Latency of the system: less than 0.5 sec from the request to the display of the image
Bandwidth occupation for conference subsystem: ≤0.5 Mb/sec
Bandwidth occupation for medical exchange subsystem: ≤1 Mb/sec

the server and simulate a wide area network (WAN)

- A conferencing subsystem: One commercial videoconferencing tool is integrated in the prototype on top of an ATM access point.
- The perceived quality is pretty good.
- A radiological browser: In the server, the images are compressed in MPEG-2 format[5,6] and stored with the corresponding index files. The client is provided with an MPEG-2 software decoder to process the data sent by the server. The system works both on Hewlett Packard (Palo Alto, CA) 9000 HP-UX and Sun (Mountain View, CA) SPARC station 20 SunOS 5.4 workstations. The graphical interface is programmed in the *X Windows System Software* of the *X Consortium* at MIT. Three resolution levels can be browsed.

The current implementation is point-to-point: It does not allow multiple clients to connect to the server. This is an aspect on which we are still working. Instead, an additional feature added to the prototype is a synchronization module in the client to guarantee that the response of the server is close enough (in time) to the request generated by the user.[7]

Table 1 summarizes the characteristics and performances of the current prototype (HP version).

COMPARISON WITH OTHER SYSTEMS

The implemented prototype shows several advantages when compared with the existing systems for teleradiology. In Table 2, the comparison between the behavior of a traditional system (based on file transfer over N-ISDN links) and ARTeMeD is presented.

Even if the prototype still needs some refinements, the achievable performance encourages the deployment of such approach. We estimate, in particular, that a high degree of interactivity as well as the capability to collaborate among different users are the fundamental issues to be considered for any solution in teleradiology.

CONCLUSION AND FURTHER WORK

Current systems for teleradiology do not fullfill the requirements for a completely distributed approach to remote diagnosis. In particular, they are lacking in an acceptable perceived level of interactivity and collaborative capability (because of the bottlenecks in the communications support and in the manipulation of images compared with the amount of information that needs to be processed). The ARTeMeD project has revealed a possibility to improve such aspects. In this project, we have analysed the existing systems and we proposed a new system that, basically results from the balancing of three different subsystems in an unique, integrated architecture.

The ARTeMeD system is designed to have

Table 2. Comparison of Traditional Systems and ARTeMeD

Feature	Traditional Systems	ARTeMeD
Link capacity	128 Kb/sec	1-1.5 Mb/sec (flexible, user specified)
Disk capacity	Required, very high	Optional, reduced
Host resources	Required, consistent	Reduced (software decoder and ATM board)
User friendliness	Good	Good
Execution time of an inspection	Long (minutes)	Very short, in two or three navigations of the overview gallery, we obtain the target image (a few seconds)
Degree of interactivity	Poor	High
Collaborative capabilities	Possible; but because of link saturation, perceived quality is poor	Supported and reliable
Openess of the system	Not allowed	Taken into account

high-speed link supports (ATM) that provide service guarantees; an efficient and optimized storage, access, and retrieval of images, and the ability to conference link facilities on the same platform.

We have produced a point-to-point prototype of such system. Future work will explore, among other issues, the reliability of communications (transmission error resilience), fidelity of image material (MPEG-2 algorithms or other compression techniques have to be taylored to take in account the specificity of medical data, see reference 8), security aspects (closed communications sessions), and centralization versus distribution of the server database.

REFERENCES

1. Basso A, Verscheure O, Noro R, et al: A multimedia architecture for medical tele-imaging over ATM. CAR '96, Paris, France, June 1996

2. De Prycker M: Asynchronous Transfer Mode: Solution for Broadband ISDN, ed 3. Ellis Horwood, Bodmin, Cornwall, GB, 1995

3. Blakowski G, Steinmetz R: A media synchronization survey: Reference model, specification, and case studies. IEEE Journal of Selected Areas in Communications 14:5-35, 1996

4. Sénac P, Diaz M, Léger A, et al: Modeling logical and temporal synchronization in hypermedia systems. IEEE Journal of Selected Areas in Communications 14:84-103, 1996

5. ISO/IEC: Information Technology—Generic Coding of Moving Pictures and Associated Audio Information—Part 3: Video. DIS 13818-3 edition, 1994.

6. ISO/IEC: Information Technology—Generic Coding of Moving Pictures and Associated Audio Information—Part 1: Systems. DIS 13818-1 edition, 1994.

7. Noro R: Un mécanisme distribué de synchronisation appliqué à un cas de téléconsultation. Journée Doctorale en Informatique et Réseaux '96, Paris, France, September 1996

8. Verscheure O, Basso A, El-Maliki M, et al: Perceptual bit allocation for MPEG-2 CBR video coding. International Conference on Image Processing '96, Lausanne, Switzerland, September 1996