PARMED-Information system for long distance collaboration in medicine*

Harald Kosch, Laszló Böszörményi[1] and Renata Słota[2]
[1]Institute of Information Technology, University of Klagenfurt, Klagenfurt, Austria
[2]Institute of Computer Science, University of Mining and Metallurgy, Cracow, Poland

Abstract

This paper describes the PARMED project initiative. The aim of PARMED is to provide medical stuff with a large virtual database of image and video data distributed over several medical centers. Clients can access the PARMED system by an active client software which provides graphical interfaces to query the virtual database. Incoming video streams and image data is post-processed by the active client software to allow analysis of this data for diagnostics or virtual teaching. Quality of service and security aspects are crucial for such kind of distributed system and will be addressed with priority.

1 Introduction

One of the most important contributions of Europe to the way people around the world are thinking about the revolution in information and communications technologies is the name now attached to it - the "information society."

In an already irreversible process, the nations have traveled well down the path blazed by new electronic technologies. These provide a single means of processing huge quantities of information in digital form and then storing, retrieving and communicating them.

At the same time, advances in networking, storage technologies and software have fueled interest in providing multimedia services yet on personal computers [1]. The advantages of digital storage of video, images and other multimedia components are numerous. Within the last ten years, various image processing, transmission, and archiving systems have been developed for medical applications. These have been focused in the areas of Radiology and Pathology, yet they are now finding their way into such areas as Cardiology, Neurology, Orthopedics, and Surgery. Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET), all kinds of imaging devices, as UltraSound Imaging (USI), and Computer Assisted Tomography (CAT) scanning systems have provided new windows to viewing the body and have become integral to the diagnosis and treatment of the disease [2].

*We gratefully acknowledge the contribution of Lionel Brunie, Marian Bubak, Witold Dzwinel, Günther Hölzl, Jacek Kitowki, Roman Krasowski, Jacek Mościński, Darin Nikolow, Janusz Otfinowski, Marek Pogoda, Markus Schordan.

The PARMED project fits into the process of constructing a helpful information society. It will provide medical stuff with a large virtual database of medical image and video data, distributed over several medical centers in Europe and maybe in a second step also outside Europe. This database can be queried via user-friendly interfaces, developed in accordance with computer scientists, and medical doctors and allows the user to formulate even complicated queries. With the help of high performance networks, yet established between the different countries of Europe, this multimedia data is brought directly to the client. At the client site we plan to provide different tools for processing the incoming multi -media informations, such as virtual teacher or data mining applications.

The paper is organized in 7 sections. Section 2 describes related works. Section 3 on the next page presents in detail the projects goals. Section 4 on page 4 gives a detailed technical description of the different project parts and the realization we attacked recently. Section 5 on page 8 discusses performance aspects and section 6 on page 9 overviews the security aspects. Finally, section 7 on page 9 concludes the paper.

2 Related work

The exchange of medical informations and collaborative work of the present still rely heavily on face-to-face interactions, group meetings, individual reflections, and hands-on experimentation. The future promises the ability for teams to conduct collaborative work using the newest instruments and computing resources through immersive electronic interaction from afar - virtual laboratory spaces. This opens the perspective to conduct the exchange.

The need for medical collaboratory has been recognized in the last years and many projects have been initiated. Actual projects rely mainly on remote single-patients diagnostics. For example in the Isar-T telemedicine project at the UHC hospital in Lille [3], methodologies are developed to perform video-based diagnostics. Thus, interactive video links between hospital's department and its counterpart at one general hospital, or from remote areas (e.g. the GSAMS (GSAMS stands for Georgia Statewide Academic and Medical System project.) project [4] in the USA) to a central in the main district town, are established. Thanks to the link, the specialists of the hospitals are able "to meet" for diagnostics or as in the GSAMS projects do "remote" diagnostics.

The EUROMED group [5], financially supported by the European Commission DG 3 under the Information Society Standardization program (ISIS), focus in their 'Virtual Medical World' project on the infrastructure and networking rather than on querying medical data. Similar projects are IMPHONE in Italy [6], a system for communication of patients' images and data, or TeleMed in Denver, CA, USA [7] (creation 'virtual' patient records).

Rather than video-conferencing or using medical video and image data to do decision support for a single patient's diagnostics, we are going to provide a large virtual database of medical video and image data. Access to this database can be granted to any medical centers disposing of at least simple multimedia personal computer. Thus, we want to educate users and potential users about how to make use of collaborative resources for particular tasks. Our project should target both end users who use the collaborative tools and make use of the data and programs available and those who add data, applications, and even parts of the infrastructure as the collaboratory grows.

3 Detailed goals

Our project goals are manigfoltic : we want to construct a distributed medical collaboratory between several medical centers in order to improve education, diagnostic or simply the local infrastructure. This requires the development of an active client software which catches the queries (also allowing its refining) and which post-processes the selected videos and interfaces to the educational or diagnostic applications. The videos must be indexed and an efficient index model must be developed. Selected videos and images have to be transferred in a reasonable time delay, thus quality of service (QoS) is an important objective of our project. Now let's detail our goals.

Medical collaboratory and education A distributed medical collaboratory between several distributed medical centers is initiated. The principal goal is the realization of the potential of available multimedia telematics service, i.e. the disposition of medical video sequences and images for supporting long-distance synchronous as well as asynchronous collaboration of several medical centers for educating (potential) users. The disposition of video and image data will be enabled by content-based retrieval which relies on the multimedia descriptive information (also called indexes). These indexes will be stored in a persistent object store.

The collaboratory in PARMED is achieved by means of the following two activities:

- First, by the implementation of an active client software module for selected (medical) application(s) that supports query specification and allows image and video data stream decoding, caching and representation and further processing in the framework of a collaboratory for selected user groups.
- Second, by the education of users and potential users about how to make use of collaborative resources for particular tasks. The program should target both end users who will use the collaborative tools and make use of the data and programs available and those who will add data, applications, and even parts of the infrastructure as the system grows.

Active client software The Active Client Software (ACS) is an application program interface (API) to most of the features offered by the PARMED system. The API is planned to be platform-independent, offering possibility to use PARMED from any kind of computer in the network. Since standardization is an important feature we propose to use HTML and JAVA standards to obtain the desired versatility and flexibility. There are at least four aims of the API:

1. Browsing medical video and image data.
2. Processing of selected video and image data according to prespecified algorithms (compression, filtering). This requires specific software to be loaded by different clients. The goal is to develop a set of procedures (perhaps unique for different medical areas) to obtain the most clear pattern of a particular disease.
3. To serve as an interface for the medical education. This is the most sophisticated purpose of the client.

4. Data analysis for statistics and classification purposes. Different cases can be compared and generalized using clustering and features extraction software.

Index database Another important aspect of PARMED is the determination of the appropriate technology for storing and querying the medical video and image data. The basic idea is to add indexes to the video/image data which are stored in an persistent object store. As we are confronted not only with a huge amount of data, but also with the requirement to execute complex queries extremely fast and to guarantee an acceptable throughput for multiple users, we plan to add parallelism to the solution.

Moreover, the video data is available in a highly distributed environment. Therefore, we need a simple abstraction that can cope with the problems of persistence, distribution and parallelism in an easy and efficient way. Such an abstraction is the notion of polymorphic sets of objects [8, 9] implemented in our extension for Java = JavaSet [10]. This concept allows the definition of query expressions in a declarative manner which makes parallelization and optimization relatively easy.

Co-operative resource management Resource management and the guarantee of Quality of Service (QoS) in our distributed system are important aspects, e.g. if the bandwidth between the video store and the clients is very low, it is possible that not all of the demanded videos/images can be transfered in a reasonable time delay. In such a situation, a global intelligent resource manager must interfere to find co-operative solutions.

Video and image archiving systems Medical applications are examples of fields which generate and deal with enormous amount of data. Objectives of the proposed research is to study and to construct an efficient system for storing, retrieving and analyzing multimedia information typical for medical applications.

The distributed video store acts as a server for the indexed database. This database, working as a client, contacts the video stores to retrieve a set of desired information (medical pictures or videos) which is required with the query.

4 Technical description

By combining software tools and high speed computer networks collaboratory can be implemented [11]. Success of this new form of collaboratory depends much on how this application will be tailored to the needs and wants of medical doctors. We believe that a close collaboration in place between medical doctors and computer scientists will help ensure that the best technology is applied to the most pressing problems. Following this paradigm, we developed in strong collaboration with medical doctors from Poland, France and Hungary, the following architecture of the system.

4.1 Architecture of the system

Fig. 1 on the following page shows architecture of the system by an example of four distributed video/image stores and two replicated index databases.

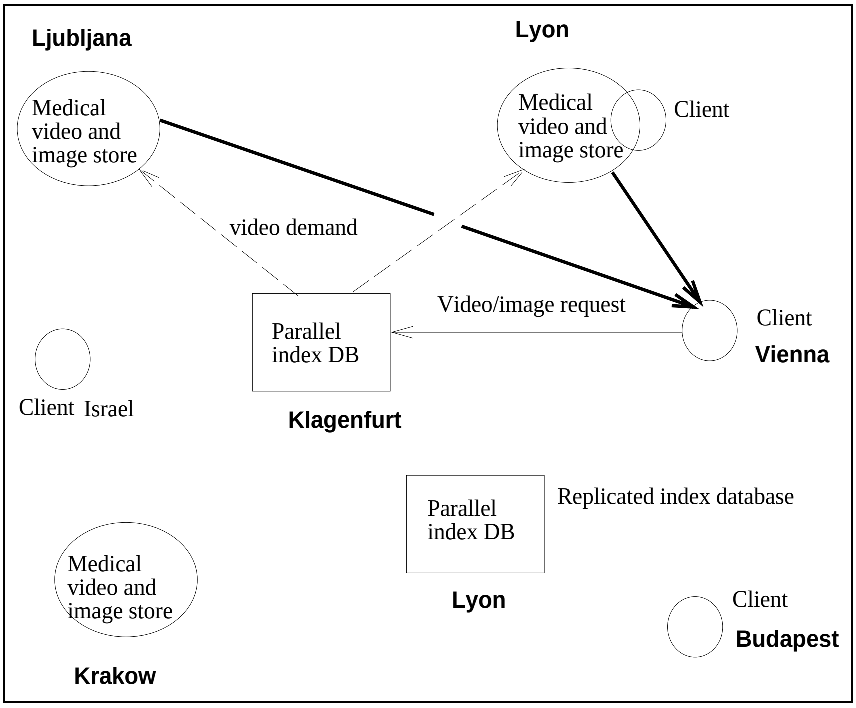

Figure 1: Architecture of the system.

The clients access to the PARMED system by an index database (DB), which is illustrated in the figure by a thin directed flesh (a query is e.g. "give me all CT-scans showing patients with scoliosis and having an age under 10 years"). The index DB searches the information on the image data (i.e. the physical location and other descriptive text) which satisfy the query. It then sends a request to the corresponding video/image sites (thin dashed fleshes). These video/image sites transfer the images to the client (thick fleshes). Besides the multimedia data complementary descriptive information could also be transfered back to the client. Notice that special attention will be paid to a balanced resource management (see section 4.7 on page 8 below).

The clients are integral part of the PARMED collaboratory. The active clients software allows the formulating of even complex queries and the post-processing of the incoming video/image and descriptive data. They are supposed to be implemented in Java. The query mask will be down-loaded from the servers and the software will guide the users during the query specification phase in a user friendly way with the aid of proposed parameters from the accompanying records to be selected.

After a successful query specification clients will intercept incoming encoded video data streams from video stores. Regardless of the actually chosen video stream format a client has to support continuous decoding explicitly. It has to be able to retrieve documents and to present them continuously while successive data arrives.

Special attention is paid to the post-processing of the video and image streams. We provide in this context an active client interface, which beside the querying capabilities, provide tools for further presentation of the data. First application concerning educational activities will be developed, e.g. a virtual teacher. Furthermore, statistical analysis and data mining applications shall help the doctors to perform a better diagnostics.

4.2 Diagnostics and education application

The time of examing a patient becomes more and more longer and the effective treatment is postponed. If all the medical data were stored in one multimedia database processor, it would be possible to access visual data immediately and thus the diagnostic process would become smoother and shorter and, at the same time, the diagnosis itself would be certainly much more accurate.

During the period of treatment the patient's health needs to be considered on a regular basis. The doctor examining the patient uses records of the previous examination as well as selected results of previously made visual examination. Unfortunately, descriptions provided in the patient's records are a weak point of medical documentation. In this situation having access to multimedia records from the virtual PARMED international database could complete the patient's records and allow so more reliable diagnostics in a shorter time. Moreover, the examining doctor would then have a possibility to give accurate, fully objective evaluation of changes in the clinical condition.

It is planed to integrate the local patient database of the Jagiellonian University Clinic, radiological images and videos issuing from the international PARMED database. This task also includes the design of a user-friendly input form for querying the database.

The PARMED database would also have a great impact on the education of the Clinic as presentations of patients could be made, all treatment of a given patient could be reproduced from the computer memory and changes in the clinical condition during a long-term treatment could be demonstrated.

Finally, the application of pattern recognition methods for medical images classification is studied. It is well known that automatic analysis of medical data belongs to the extremely tough tasks. It would be profitable to offer to medical doctors possibilities of grouping (clustering) images. According to the medical knowledge each of the group can be representative to a specific disease. We plan to apply different kinds of pattern recognition algorithms based for example on clustering techniques, genetic algorithms and neural networks and implement the best adapted for the PARMED video and image store.

4.3 Standardization of the index data

Index Medicus, Medline and Medical Schools Libraries have their proper hierarchical structure for the whole medicine designed to search literature. Thus, the need for standardization in medical search is obvious.

The issue of the standardization within PARMED is the creation of a multidimensional array based on the above search masks and to adapt it to our elaborated request forms. Special attention is paid that the developed standardization is compatible with the WHO coding system (contains the list of all known diseases, conditions, and their treatment methodology). This standardization process is extremely important for PARMED, as we require a common "language" platform for the descriptive information of the medical data.

4.4 Index database

The basic idea of facilitating access to the medical video and image data is to add indexes to the multimedia data and to store these indexes in a light-weighted database (hereafter called persistent object store POS). Two available persistent object store implementation in Java have been successfully tested and are integrated into the project, first the *ObjectStore* from Object Space [12] and second the *PJama* [13] implementation from the university Glasgow. These products give us the access to persistent objects, however allow not yet the formulation of queries on the index data. For this, we propose an extension of the Java language by sets (named JavaSet) which allows the formulation of database-like queries by a special select-expression [10]. At time, we are implementing the JavaSet compiler and optimizer enabling high performance processing of the translated queries.

Once the system is up and several centers are connected to the index database, we are confronted with a possible large number of concurrent access and the great volume of stored indexes. In order to manage such a system we are planning to employ parallel database technologies. This requires the extension of the persistent object store to multiple instances and the JavaSet to support distribution. This extension is subject to future goals in the project.

Two main handicaps for searching video and image data efficiently exists. First the great volume of data (especially for the video data) and second the lack of structures in their contents. In order to solve those problems, the video/image data is enriched by semantic information. This index informations, they are interactively specified by users or retrieved automatically (see section 4.6) and serves as indexes of the contents of the video or images. Such indexes are collected from all the participating video/image stores and are put into the POS. In order to provide a better flexibility when updating the indexes, or a better service for the clients, the POS will be replicated at different sites. However, most of the videos and images are not yet indexed. In strong collaboration with the participating medical doctors, we develop electronic forms for the video sequences to be filled (see the application descriptions above).

4.5 Data acquisition system

The system should be equipped with data acquisition input interface. It consists of a hardware-software part for digitalization Roentgen pictures and of a pure software part for collecting textual medical information about the patients' records which are analyzed and put into the index database (the range of information will be determined after medical consultations). Completed informations are stored in the data storage system. It is important to obtain informations about the patients from different sources. It may be a medical instrument which post-process the digital pictures at the output of our hardware-software input interface for digitizing photos. Algorithms for image processing, quality correction, detail extraction and compression are at time in implementation.

4.6 Storing and accessing the video/image data

From the system/hardware point of view the data archiving systems based on the three-level hierarchical mass storage comprising besides of hard disks, also juke-boxes and tape libraries, supported by HSM software is applied.

First, the field of distributed video store acting as a server for the indexed database is studied and requirements for accessing the video/image archiving system is specified. Then, two archiving systems are to be constructed: a three-level storage management (HSM) system [14] consisting of hard disks cluster, magnet-optical and tape libraries, and a smaller system with a hard disks cluster. One of the problems to be tackled in this approach is the monitoring of the HSM system in order to obtain QoS information. In this approach we are not be concerned with scheduling of requests of access to the archived data. This task is handled by internal mechanisms of the HSM software. Based on the monitoring results, the implementation of the specialized archiving system for medical data will be attempted including request scheduling and QoS monitoring.

The optimization of I/O capabilities of the system is in the second important issue. In this context, we plan to design an efficient and user friendly graphical interface specially adapted for medical applications. The interface retrieves images and movies together with accompanying text information, enable to filter images to enhance most valuable information, show in a compact way a full state of the patient and treatment. The software environment should be flexible and easily modified by the users according to their needs.

4.7 Resource management

A very important aspect in the proposed highly distributed environment is that resources have to be allocated and scheduled end-to-end for each client/server occurrence. The resource-management has to be distributed to be failsafe and to avoid large communication efforts.

A "co-operative resource sharing" is in development. The idea is to split the resource-manager into two parts. One part is doing the basic management-work of the resources as part of the operating system to avoid occupation from more than one application. The other part is settled in the application area and communicates with the applications (see above for the storage system). The goal is to give "balanced resources" to the applications, e.g. if the computing power in the pipe of a client/server-connection is very poor then the communication power in most cases also doesn't have to be very high. Balancing also includes the division of the resources to the various applications. An application should be able to determine how resources can be shared between other applications. In many cases it is not practical (and often not possible) to increase the number of resources if a new additional application is started. So it is necessary to have a good policy to treat applications in the same priority class more or less equivalently.

5 Performance considerations

The standard bandwidths between the different locations of the project allow correct working with the medical data. For example, to down-load in Klagenfurt medical images with 4000 x 4000 pixels of 4 bits from Kraków (representing a volume of 3 MBytes), only some tenths of seconds are required.

Even in difficult resource availability situations (e.g. a low network bandwidth), our intelligent resource manager (see section 4.7) examines the overall situation and balances the resources between the clients in a way that the quality of services can be maintained the most as possible. For example, if a reduced bandwidth from the video/image store to the client is detected, the resource manager tries to maintain the estimated transfer rate by first-aid measures (e.g. supplemental caching). If this is not possible, only most important medical data is transfered (e.g. those with the highest hit-rate) or the the video stream itself is manipulated (e.g. less refreshments). The client is entirely informed of those measures and can on his hand propose further measures.

6 Security Aspects

Security is a critical feature of the system. The transfered data is highly sensible (patient, doctors and institutes data) and should be best protected. It is clear that if the underlying communications facility is Internet, the communicated data is exposed to all security threats presented in open networks.

Several projects have addressed this problem e.g. [15] and proposed objectives and measures to avoid the misusing of patients and doctors data. In the context of the PARMED project we have to deal with more severe problems as information might issue from different countries, each having different security requirements. Furthermore, we use a WWW based approach, which needs supplement security solutions, too. In this context we rely on the EUROMED project [16], they develop at time a security solution for the WWW-based telemedicine project. Our project supplementary contributions to the security aspect will be an integration of different country security requirements for transferring medical data.

7 Conclusion

The goal of the PARMED project is the construction of a medical long-distance collaboratory supported by database technology. We have started with a modest environment (with partners from a few countries), in which we want to design and implement the basic kernel of a virtual European PARMED database, available for any medical person, disposing of at least a multimedia personal computer. Our aim is collaboratory in all components of the system, starting from the input forms and reaching to the balanced resource manager.

Our long-term project aim is to turn the PARMED system into a *world-wide medical support system*. Thus, by initiating new industrial cooperations we can attract more partners into PARMED. Once a fully functional prototype can be provided and a first prototype is patented, we will integrate new medical centers, industrials and end-users in order to define a widely spanning *Europa-wide medical support system* and with partners from over-sea a *World-wide medical support system*.

References

[1] Peter Wayner. X terminal + browser + Java = Web PC. *BYTE Magazine*, 21(10), October 1996.

[2] L. Brunie, F. Leitner, F. Berthommier, P. Cinquin, and J. Demongeot. Interpretation of multimodal medical images using connexionist and variational methods. *Technology and Health Care*, 3(2):91–100, 1995.

[3] University Lille2. The Isar-Telematics project. http://www.univ-lille2.fr/isart/, June 1998.

[4] Georgia Statewide Academic and Medical System. GAGSM - the distance learning and healthcare network system. http://gagcm.gagctr.uga.edu/, June 1998.

[5] Institute of Comunications and Computers, National technical university of Athens, Athen, Greece. *The EUROMED project*, June 1998. http://euromed.iccs.ntua.gr.

[6] Hewlett-Packard Ltd. Medical Products Group. http://www.rad.unipi.it/, June 1998.

[7] Bill Siwicki (Senior Editor of TeleMed). National scientists work to create 'virtual' patient records. http://hdm.faulknergray.com/html/content/lanl.html, June 1998.

[8] L. Böszörményi and K.-H. Eder. M3set – a language for handling of distributed and persistent sets of objects. *Parallel Computing*, 22(1):1913–1925, January 1997.

[9] H. Kosch and L. Böszörményi. High Performance Sets. In *HPCN Conference 98*, LNCS 1401, pages 972–975. Springer Verlag, April 1998.

[10] M. Schordan, H. Kosch, and L. Böszörményi. JavaSet – extending Java by persistent sets. In *Proc. of the Third International Austrian-Insraeli Technion Symposium*, April 1999. to appear.

[11] David G. Kilman and David W. Forslund. An international collaboratory based on virtual patient records. *Communications of the ACM*, 40(8):110–117, August 1997.

[12] G. Landis, C. Lamb, T. Blackman, S. Haradhvala, M. Noyes, and D. Weinreb. Objectstore PSE: A Persistent Storage Engine for Java. In *Proceedings of the 1st International Workshop on Persistence for Java*, Glasgow, Scotland, September 1996. Sun Microsystems, TR 96-58.

[13] M.P Atkinson, L. Daynes, M.J. Jordan, T. Printezis, and S. Spence. An Orthogonally Persistent Java. *Sigmod Records*, 25(4):68–75, December 1996.

[14] D. Nikolow and M. Pogoda. Experience with UniTree Based Mass Storage Systems on HP Platforms in Poland. In M. Bubak and J. Mościński, editors, *Proc. Int. Conf. High Performance Computing on Hewlett-Packard Systems*, pages 189–194, September 1997.

[15] The ISTAHR group. Telematics applications programme project ISTAHR (Implementation of a Secure Healthcare Telematics Applications in Europe). http://ted.see.plym.ac.uk/istahr/, June 1998. CEC DG XIII/F.

[16] D. Polemi and A. Marsh. Secure telemedicine applications. In *HPCN Conference 98*, LNCS 1401, pages 1015–1017. Springer Verlag, April 1998.