

The importance of Java and CORBA in medicine

David W. Forslund, PhD, , James L. Cook, MD*

Los Alamos National Laboratory

Los Alamos, New Mexico

*National Jewish Center for Immunology and Respiratory Medicine

Denver, Colorado

One of the most powerful tools available for telemedicine is a multimedia medical record accessible over a wide area and simultaneously editable by multiple physicians. The ability to do this through an intuitive interface linking multiple distributed data repositories while maintaining full data integrity is a fundamental enabling technology in healthcare. We discuss the role of distributed object technology using Java and CORBA in providing this capability including an example of such a system (TeleMed) which can be accessed through the World Wide Web. Issues of security, scalability, data integrity, and usability are emphasized..

INTRODUCTION

Healthcare is undergoing a rapid change due to the enormous improvements in computing and communications technology. It is no longer adequate in healthcare to think of the electronic medical record as an analog to a paper record which is contained in a hospital's medical records department.¹ Because healthcare is rapidly taking on a distributed nature, it should be regarded as a collaborative tool that can incorporate all types of data including sound, images, and video in an integrated manner even though the sources of data and processing may be widely distributed². This requires a robust distributed computing infrastructure accessible on a wide range of platforms. The premier open architecture which can meet this requirement is based on the Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA)³. We describe here the motivation for the use of Java and CORBA in telemedicine, its benefits, limitations and an example implementation and use in the TeleMed system designed and implemented at Los Alamos. With its basis on object technology it provides support over a wide area network of proper containment and management of data and processes involved in healthcare and in other industries.

BACKGROUND

Although we will use the term *telemedicine* throughout this work, we believe the term is inappropriate because it tends to distinguish itself from the normal practice of medicine. The goal of telemedicine should be to bring advanced communications technology to bear in the normal practice of healthcare. For the purposes of this paper we define *telemedicine* to be the practice of moving healthcare to patient with communications technology. This amounts to being able to practice healthcare in a distributed

manner without regard to the location of the patient or even the physician. For example, the healthcare system needs to support synchronous and asynchronous remote consultations. That is, the physicians may need to work together in real time for diagnosis or treatment or be able to provide consultation services off-line at his or her convenience. This should be done without having to use special equipment or go to special locations.

Having data made available simultaneously over a wide area implies a protection of an individual's privacy in a way that is not normally done in medical institutions today. Clinical records need to be made available to people only if they are authorized and the protections need to be commensurate with the risk associated with a broader network accessibility. This implies secure management of a public/private key infrastructure to authenticate access and to manage access control

One of the weaknesses of most telemedicine systems being put in place is the lack of a provision for data integrity. The information generated in the telemedicine encounter are not necessarily captured in the clinical record, or worse, the information gets entered multiple times and is out of synchronization between the multiple locations, because each site manages the information differently and copies are sent around from location to location. It is important, then, that telemedicine systems fully address the management of the data they generate and minimize errors associated with copying and separate management of these copies. Basically, the principle is single entry and a provision for a unified view of the clinical encounters.

The management of medical information over a wide area cannot be done without some means of unifying the view of disparate data. For example, it should be possible to locate patient data anywhere on the network by means of a distributed master patient index as well as a means of resolving the distinction between disparate terms which are used to describe drugs, encounters, diagnoses, treatments, etc.

Once the data is managed over a wide area in a consistent manner, it should be possible with a telemedicine system to provide tools for much more automatic outcomes analysis. That is, one should be able to utilize the system itself to manage the aggregation of data to facilitate the evaluation of the telemedicine system itself as well as the determination of factors needed to improve healthcare treatments. For the

telemedicine system to be truly cost effective it must demonstrate that it can reduce the time required to reach a diagnosis and begin treatment of the patient. This should then result in improved outcomes for the same or less cost.

One of the foundational requirements for the system to be effective even if it can aggregate consistent data over a wide area is for it to be easy to use by the physician and easy and fast to access. Because a medical record can be very complex and data intensive, it must provide a mechanism for intuitive navigation through the information by a health care provider. It must provide the knowledge the caregiver needs without overwhelming her with unnecessary information not relevant to the current decision process. At the same time the relevant information must be immediately accessible to reduce the time to reach a diagnosis. Basically, the system must give the illusion of providing the entire medical record on demand to the physician while actually only delivering small parts of it.

VIRTUAL PATIENT RECORD REQUIREMENTS

Considering these above mentioned requirements for distributed healthcare, we realize that we need to be able to manage a "virtual" patient record that is assembled from multiple locations dynamically. Although this is highly desired, there are many obstacles to being able to create a "virtual" patient record in healthcare. One important requirement is the ability to map the disparate data and process representations that are prevalent in healthcare into a common structure. This would enable very different systems with their own internal data structures to communicate effectively. This approach means more than simply a messaging infrastructure, but a common way to describe the processes as well as the information in healthcare. It should rely on infrastructure usable in a variety of disciplines and capture the processes of the "business" of healthcare. Sometimes this concept is referred to as "common business objects". The actual internal layout of data is not particularly relevant to this approach.

A second requirement is high accessibility of the information while simultaneously providing high data integrity. To provide this, low-latency networks are needed (more important than high bandwidth) so that locations can be queried quickly and data references transferred. When quality networks are not available, technology for database replication and consistency management must be provided. A combination of these technologies is required for proper support of the virtual patient record. A corollary of this requirement is the support for distributed transaction processing so that data can be updated consistently between multiple data sources.

As data is merged to produce the medical record "on the fly", the maintenance of privacy for the end user becomes even more important. Although this can be complex, the infrastructure is being put in place to-

day to manage this problem in areas besides healthcare. Appropriate authentication as well as authorization to limit access to certain types of data are required. The data also needs to be encrypted at least between the source and the destination with a manageable risk associated with its interception. Finally, audit trails of access are required so that data access and modification is non-repudiable.

A method of locating patient information is also needed to support the virtual patient record so that locations can be "automatically" queried to see if they have relevant patient data. Although a central repository may seem to be the most logical solution, it is, in fact, not desirable because of issues associated with manageability and ownership of the information. An Internet-based "DNS-like" solution in which there are peers which exchange relevant data pointers is more resilient to system down time and allows for the local management of patient privacy in the institution in which the knowledge is created.

To achieve a truly usable virtual patient record will require industry cooperation on an unprecedented scale and the use of infrastructure support relevant to a broad-base of industries. Signs of this new level of cooperation is already apparent in the various consortia that are seeking to improve healthcare.

THE ROLE OF THE INTERNET

The rapid acceptance of the Internet can provide a substantial impetus to meeting these above requirements. It provides standard communication mechanisms (e.g. TCP/IP, HTTP, IIOP), platform independence, low cost distribution methods (e.g. Java) and a high level of connectivity. It supports various distributed services including standard naming mechanisms (e.g. URN's, i.e., Universal Resource Naming⁴), security (e.g., DCE, certificate servers, etc.), tools for distributed data abstraction and encapsulation (CORBA). At the same time, there is a rapid evolution and acceptance of multimedia technologies (internet-base audio and video), as well as scalable storage systems (e.g., object relational databases) which support a variety of data-mining capabilities.

The requirements we have outlined to support a virtual patient record are formidable. Is there any infrastructure being put in place which could simultaneously meet all of these requirements? Probably not. However, there is an approach within industry which provides many of the mechanisms needed to support a virtual patient record. It is based on the work coming from the OMG consortium. At the first level, now traditional object oriented analysis and design (OOAD) methodology is used to provide the right level of abstractions and definitions of data and processes involved in healthcare. This work is being spearheaded by the healthcare task force in the OMG known as CORBAmed. It is insufficient at this point of technology to simply describe the data that needs to be exchanged between healthcare facilities. The processes must be encapsulated and common definitions must be agreed upon, not simply in the aca-

demic sense but in the truly practical sense of commercial applications.

As powerful as the OOAD methodology is, the wide area management and access control of objects is much more important. It is this work of the OMG that can provide the greatest benefit to the healthcare industry. As part of the infrastructure support of CORBA, a variety of services are being defined for the management of the objects being accessed over the network. The strength of the OMG comes from the fact that it is an industry-led consortium representing a very broad commercial market and thus can leverage the infrastructure investment across many industries. The current membership consists of over 700 companies and institutions.

The CORBA architecture⁵ separates the object interfaces from implementations. Besides providing for language and platform independence, this approach enables an open specification supporting proprietary implementations which facilitates competition and object reuse. Also part of the Object Management Architecture of CORBA are the various CORBA services, such as Naming, Event, Query, Transaction and Persistence, Security, Versioning, and Lifecycle. Each of these are important to the management of the distributed object infrastructure of value to healthcare. In addition, there can be a number of facilities, some of which would be special to healthcare such as Medical Objects and a Master Patient Index, although there may be common services underlying similar facilities in other domains. Many of the components of CORBA are commercially available from a variety of vendors. The Object Request Broker is the core of this technology which negotiates connections between locations and provides the support for the various services and facilities. The interoperability of the ORBs with the Internet Inter-ORB protocol (IIOP) has been demonstrated between almost all of the commercially available ORBs. Efforts are now underway to develop the components for the various vertical markets that will be managed by these ORBs.

A fundamental service of CORBA is the Naming Service which manages the pointers to distributed objects. This is of significance to the Master Patient Index as well as clinical lexicons and for providing the management of the variety of information sources required to support the electronic medical record. The Event service provides the foundation for telecollaboration through the electronic medical record so that when objects or data are added to the record, they can be seen simultaneously at multiple locations. The ability to fully interact with the medical record from multiple locations is sustained by the Transaction and Persistence Services which ensure data integrity both on the short term and the long term. If a multimedia data object is inserted into the medical record, these services ensure that it is preserved for viewing by others when allowed. The CORBA security service provides for full support of secure authentication, authorization, and network encryption of transmitted data while providing a non-repudiation

service so that events placed in the medical record cannot be denied at a later date and their source can be verified with confidence. The Lifecycle and versioning services provide for the long term maintenance of the distributed objects.

The popularity and ubiquity of the Internet is leading many vendors to provide WWW or HTML interfaces to their data repositories both in healthcare^{6,7} and in other fields. Although extremely useful and easy to support, there are several difficulties with this trend. The Web is primarily a publishing paradigm by design because of the stateless nature of the http protocol. Although state can be maintained on the server, the client really has no knowledge of it which can lead to confusion in the data if not a sacrifice in data security and integrity. The IIOP protocol provides support for transactions and full interactivity for Internet operations. Also because of the stateless nature of the protocol, there is only poor support for authentication and authorization. In any case, it is essentially impossible to interact with server objects in near real time which is required for full collaborative capabilities. In addition, the limitations of the URL for locating information are profound since it is typically merely a disk address and a machine name. Improvements in the identification of Internet data are in the works as result of efforts of the IETF and the URN working group. The addition of Java to the Internet may result in a substantial alleviation of some of these problems, although Java, by itself, does not deal with the interaction between the client and server and is only as powerful as the protocol underneath it. But it does substantially enhance the ability to interact with the data in a less constrained manner and is much more maintainable than CGI scripts.

The linking of Java client applications through CORBA⁸ to multiple servers provides a whole new dimension to the management of healthcare information and support for a multimedia electronic medical record. Besides supporting a multi-tiered architecture instead of a mere client/server architecture it provides major advantages in making it easy to distribute client software. The client infrastructure is provided completely by the Java Virtual Machine (JVM) or the Internet web browser, reducing the developer's task. The JVM also enables other languages to be supported by various compilers which extends Java's functionality considerably. In any case the flexible interface supported by Java allows for advanced human computer interfaces to be delivered to the physician for ease of use.

TELEMED OVERVIEW

*TeleMed*⁹ is a system developed jointly with National Jewish Center for Immunology and Respiratory Medicine in Denver, CO with almost all of the above principles in mind. It was designed with a clear separation between the underlying hardware, the services layer and the application layer, with the idea of supporting a range of application domains. Both the services layer and the application layer are designed

with open distributed object technology. CORBA/ORBs are used for all communications, and the multimedia graphical interface has support for audio objects. Multiple servers can be used and the medical records from multiple Object Oriented Databases can be combined through the “merged patient record”. The data at multiple remote locations can be accessed logically as a single data set or serially with explicit control over the location. The CORBA infrastructure has allowed us to create multiple servers which carry out various decision support functions such as comparing a subset of images to see which ones are similar. These servers can transparently take advantage of parallel architectures for scalable concept extraction techniques. In addition, the individual objects can have their own security and authentication so that different portions of the medical record can be made visible for different uses. Finally, all the objects can come from multiple locations and thus support the fully distributed paradigm. Figure 1 shows a diagram of the *TeleMed* architecture illustrating some of these features.




Figure 1: TeleMed 3 tiered architecture showing multiple persistent storage servers delivering a merged patient record to a client along with auxiliary servers for decision support.

There are a number of principles that were used to guide the development of the TeleMed software. The first one is to move data only as necessary. We pass persistent references to the data which are used to obtain the data or objects only when they are requested. This can radically reduce the bandwidth requirements for use. The second principle is to manage complex high-volume data in an understandable manner. That is, we want to give the illusion that the whole data set is available, and can be navigated through in a straightforward manner. Third, we require that the same architecture be used for real-time consultations as for asynchronous consultations. That is, users can interact with the data simultaneously or provide their comments off-line. Fourth, we seek to leverage Internet standards and use them wherever possible rather than develop our own. For example, we have developed our own CORBA security model but are replacing it with the CORBA security

commercially provided as it becomes available. Fifth, we pay close attention making the system easy to use by the non-computer professional. The protocols and procedures are captured in the client interface so that the physician is dealing with information in a way that is entirely familiar. This principle with major contributions from National Jewish Hospital has resulted in the creation of the graphical patient record (Figure 2) which shows at a glance the longitudinal medical history of a patient including various reports, lab tests as well as a variety of treatments. Through this interface we also make available data-mining services such as comparing a CT image with other CT images which are deemed to be similar based on algorithms which extract and compare image features.




Figure 2: Web-accessible, java-based, graphical patient record showing longitudinal patient history with graphical references to remote data.

TELEMED DIRECTIONS

In order for the principles of TeleMed to be proven and shown practical for a wider audience, we have undertaken several efforts to more broadly deploy TeleMed. First, we have joined with Northern New Mexico Community College in Espanola, NM to provide TeleMed as a tool to link together 18 clinics in Northern New Mexico. It will be used to share data such as immunizations, prescriptions and health encounters between the clinics to more effectively treat these rural patients. Using Java and CORBA, we will be providing a simple electronic medical record over POTS. In addition, we have undertaken to evaluate and adapt TeleMed for us in the military both in combat situations as well as in a few Tricare regions. In order to provide the necessary linkage for wider use of the TeleMed technology, we have undertaken to coordinate efforts with other organizations for the adoption of a wide area master patient index which would enable healthcare providers to more readily locate patient information. We also anticipate the use of TeleMed as a tool to manage infectious disease information at the state level both

for the individual and aggregately. In addition, we are working with the CORBAmed healthcare taskforce of the OMG to collaborate in bringing this object technology to the healthcare community.

For more information on TeleMed please visit our Web site:
<http://www.acl.lanl.gov/TeleMed>.

CONCLUSION

Our experience with CORBA indicates that with the most recent CORBA services, it is sufficiently mature to support the requirements of the healthcare industry. Indeed, multimedia patient records can be made a reality in a scalable system usable over a wide area network, in which location information can be made transparent to the end user. Security and integration of disparate systems to make such a distributed vision a reality remain significant issues which will need continual work. The former, as much from user acceptance and ease of use as with actual protection of information. The latter involves pulling together legacy data sources in a sensible way. Finally, our experience indicates that data-mining techniques can be provided in such a way as to be readily usable by non-technical users.

Acknowledgments

This work was performed under the auspices of the United States Department of Energy and partially supported by the United States Army, Medical Advanced Technology Management Office of the U.S. Army Medical Research and Materiel Command

References

1. Dick, R. S., and Steen E. B., Eds. *The Computer-Based Patient Record, An Essential Technology for Health Care*, Institute of Medicine, National Academy Press, Washington, D.C., 1991
2. Phillips, R. L., *A Network-based Distributed, Media-rich Computing and Information Environment*, Paper presented at the Digital Media and Electronic Publishing conference sponsored by the British Computer Society in Leeds, UK, December 1994. Available at: <http://www.acl.lanl.gov/TeleMed/Papers/keynote.pdf>
3. Object Management Group, Ed., *Common Object Request Broker Architecture*, Wiley-QED, 1991.
4. Daniel, R., and M. Mealling, *Resolution of Uniform Resource Identifiers using the Domain Name System*, an Internet draft specification: draft-ietf-urn-naptr-04.txt, 1997.
5. Orfali, R., Harkey, D., Edwards, J., *The Essential Distributed Objects Survival Guide*, John Wiley & Sons, Inc., New York, 1996
6. Kohane, I. A., et al., Sharing Electronic Medical Records Across Multiple Heterogeneous and Competing Institutions. in *Proceedings of the 1996 American Medical Informatics Association An-*

nual Fall Symposium (Washington, D.C., Oct. 26—Oct. 30). ed. J. J. Cimino, American Medical Informatics Association, Hanley & Belfus, Inc., 1996, pp 608-612

7. McDonald, C. J., et al., The Regenstrief Medical Record System (RMRS): Physician use for input and output and Web browser based computing. in *Proceedings of the 1996 American Medical Informatics Association Annual Fall Symposium* (Washington, D.C., Oct. 26—Oct. 30). ed. J. J. Cimino, American Medical Informatics Association, Hanley & Belfus, Inc., 1996 p. 989
8. Orfali, R. and D. Harkey *Client/Server Programming with Java and CORBA*, John Wiley and Sons, Inc., New York, 1997.
9. Forslund, D. W., et al., Experiences with a Distributed Virtual Patient Record System. in *Proceedings of the 1996 American Medical Informatics Association Annual Fall Symposium* (Washington, D.C., Oct. 26—Oct. 30). ed. J. J. Cimino, American Medical Informatics Association, Hanley & Belfus, Inc., 1996, pp 483-487