CORBAMed and DHE: Middleware Service Approach in Healthcare Information Systems

Dongsong Zhang , Center for the Management of Information (CMI), University of Arizona Tucson, AZ 85721
Tel.: (520) 621 1860 , zhangd@u.arizona.edu

Ralph Martinez
Department of ECE , University of Arizona Tucson, AZ 85721 , Tel.: (520) 621-6174 , martinez@ece.arizona.edu

ABSTRACT

The healthcare organizational structure is often naturally distributed. In order to improve the quality of care, the adoption of standards to allow the effective and robust networking of the various centers for clinical, epidemiological, administrative information management purposes has been widely recognized as an urgent and strategic need in the healthcare community. This paper analyzes the status and challenges that today's healthcare information systems have, and introduces two middleware service frameworks for information systems, namely CORBAMed and DHE, in detail. The middleware service can address heterogeneous problems and significantly assist interoperability and integrity of information systems by providing common services and a set of standard interfaces that enable different applications to interact with each other.

1. INTRODUCTION

1.1 Characteristics of Current Healthcare

Healthcare is experiencing rapid improvements due to its changing focus and tremendous achievements in computing and communication technology. Concurrently, the healthcare organizational structure is often naturally distributed, being geographically dispersed centers at different levels of complexity and scale. The various structures operating over the territory are heterogeneous in organizational, logistic, and clinical perspectives. Healthcare is driven to move quickly from single-provider, proprietary, mainframe systems toward multi-providers, scalable, component-based systems. Patients are provided with healthcare services by multiple institutions. The patient information is widely dispersed — a patient's record is held in medical records departments of multiple hospitals. Therefore, ubiquitous lifetime patient information needs to be obtained from many locations and be efficiently integrated. Electronic medical records are commonly used (Forslund, 1997). They can be easily maintained and distributed across a healthcare network.

1.2 Challenges of Developing Healthcare Information Systems

The distributed nature of healthcare determines the need of powerful tools to access patient medical records over a wide area and assemble various data from different sources. There is a strong desire to create and apply some common standards, which are assertions, not realizations of requirements, for the architecture of HIS. The adoption of them allows the effective and robust networking of the various centers for clinical, epidemiological, and administrative information management purposes. As a result, physicians will be able to maximize the utilization and sharing of resources to make the right judgment. However, it is far from trivial to make it a reality because of several challenges:

- Size/Scalability. Usually the size of the overall healthcare network is huge.
- Representationheterogeneity. Same patient information may be represented in disparate ways within different systems, which sets the barrier for data integration. There is no consensus on data formats.
- Platform and programming language heterogeneity. Individual healthcare information systems are often platform-dependent. They are usually developed using different programming languages on heterogeneous platforms, which causes incompatibility problems.
- Lack of computing interoperability. Hospitals use assorted devices, instruments and systems that collect and maintain patient information. It lacks of computing interoperability among them.
- Real-time capabilities. Gathering and integrating distributed patient information must be quick enough to be useful for diagnosis or treatment.
- Security. Today's healthcare providers are concerned with collecting relevant patient information from multiple resources. The security and limitations on data access and disclosure must be taken into consideration during patient information sharing. The security will be the key to the smooth flow and share of information among the providers in healthcare community.

Developing advanced healthcare information integration techniques and standards will address the above problems and greatly benefit the healthcare community. First, it helps assemble longitudinal medical records of a patient from multiple sources by specifying a set of information formats. Second, it promotes interoperability and enables the reliable sharing of medical information between healthcare organizations. Also, it supports remote synchronous and asynchronous healthcare consultation. Finally, it reduces the cost of healthcare information management. Therefore, there is a strong need for a distributed computing infrastructure accessible to a large variety of platforms, and a set of standard interfaces that links distributed data repositories together while maintaining integrity. The development of a fully-interoperable, component-based and open-standard framework for HIS will make healthcare organizations transcend the limitations and redundancies created by today's automated systems, and encapsulate the pieces of existing systems so that all of the system components could work together seamlessly. This paper will introduce two middleware frameworks, CORBAMed and DHE that bridge the gap between different platforms and applications by providing standardized interfaces and services.

2. THE STATE OF HEALTHCARE INFORMATICS

The advent of HIS began in the 1960's. Originally, the were mainframe-based systems that primarily focused on supporting accounting and administrative functions in the hospitals. In the 1970's, department information systems, such as radiology information systems and pharmacy management systems, were developed to support administrative and information tracking functions at clinical departmentlevel.

Since the 1990's, healthcare informatics has begun to support the clinician and focus on the patient needs. It tries to incrementally construct a homogeneous and consistent set of information common to the whole organization that optimizes the cycle of healthcare activities. The healthcare community has increasingly recognized the urgent need for interoperable information systems. Obviously, centralized mainframe systems are not able to meet these requirements. Instead, distributed object technology seems ideal for this purpose because of the flexibility and simplification it provides.

There have been increasing efforts toward establishing patient data-interchange standards and designing distributed system architectures for healthcare, such as Health Level 7 (HL7), Digital Imaging and Communications in Medicine (DICOM), and Electronic Data Interchange for Administration, Commerce and Transport (EDIFACT), etc. These standards are widely used. However, most of them focus on health data exchange and only support part of patient information. These standards or protocols neither provide truly open systems that enable users to select best-of-breed applications, nor exchange data in a meaningful way. In order to achieve an open plug-and-play environment, industry-standard middleware architectures and interfaces should be well defined and adopted.

3. MIDDLEWARE SERVICE APPROACH IN HEALTHCARE INFORMATION SYSTEMS

3.1 Middleware

In order to develop enterprise-wide information systems, it is imperative to solve heterogeneity and distribution problems. To help solve these problems, vendors are offering distributed system services that have standard programming interfaces and protocols. These services are called "middleware services" because they sit in the middle of the platform layer and application layer.

Basically, middleware is a package of connectivity software that bridges the gap between application programs and heterogeneous hardware and operating system dependent platforms. It is extensively used in the heterogeneous, distributed environment through providing a functional set of Application Programming Interfaces (API) or a protocol to allow an application to get access to any popular platforms. A primary advantage of using object-oriented middleware is that any object can be inexpensively updated while the involved technologies remain intact and unchanged. The implementations of a middleware service must be able to run on multiple platforms. Otherwise, it is only a platform service. A middleware service has distributed nature: it can be accessed remotely or it can access other remote services and applications. Generally, the middleware services may include several components such as information management, computation, communications, control, system management and presentation management (Bernstein, 1996).

HIS architecture is dominated by the extensive use of object-oriented middleware functionalities. Today the most widely publicized middleware initiatives are the Open Software Foundation's Distributed Computing Environment (DCE), Microsoft's COM/DCOM, and Object Management Group (OMGTM)'s Common Object Request Broker Architecture (CORBAâ1). CORBA is a distributed object framework that intends to enable interoperability between objects in a heterogeneous, distributed environment. Its architecture separates the object interfaces from implementations, while a set of CORBA services facilitate the management of the distributed object infrastructure (Orfali & Harkey, 1998).

3.2 CORBAMedä2

3.2.1 What is CORBAMed?

One of the leading efforts toward developing industry standards for healthcare information system is CORBAMed, which is the OMG's Domain Task Force on Healthcare. The mission of CORBAMed is to improve the quality of care and reduce cost by the use of CORBA technologies for interoperability, and to define standardized object-oriented interfaces between healthcare related services and functions (Forslund et al., 1999; OMG, 1997). Some important features of CORBAMed are described as follows:

- Opens the access to a larger range of software components.
- Provides a high degree of interoperability between disparate computer systems so that medical records can be integrated.
- Improves the quality of components and reduces the software cost by increasing the competition between healthcare providers.
- Is fully object-oriented.

The overall CORBAMed framework is presented in Figure 1 (ACL, 1998). It includes six major standard services that support HIS based on CORBA technologies. The first approved standard from CORBAMed was the Person Identification Service (PIDS). Its goal is to provide a standard method of locating person identifiers and their medical records across a distributed healthcare network. Lexical Query Service (LQS) that was secondly approved supports the use of multiple vocabularies in heterogeneous environment. The other four standards, namely Healthcare Data Interpretation Facility (HDIF), Clinical Observation Access Service (COAS), Healthcare Resource Access Control (HARC) and Clinical Image Access Service (CIAS), are either under assessment or to be submitted.

3.2.2 Person Identification Service (PIDS)

Every person may have dozens of healthcare providers in their life. Each healthcare organization simply assigns a unique patient identifier within its local domain to a patient. These IDs are meaningless to other systems or organizations. In a distributed environment, a patient's medical information is dispersedly located at a number of organizations. Each part of the information belonging to that patient needs to be correctly identified and collected, which means a search has to be performed in each organization in order to identify the person and retrieve the related information. However, increased specialization of healthcare providers and continuous growth of organizations exacerbate the problems with managing patientIDs.

The building block of the CORBAMed PIDS model is the ID domain. The PIDS tries to manage patient IDs by supporting the assignment of IDs within a particular ID domain, correlating IDs among different domains, allowing both attended-interactive and message-driven unattended modes for searching and matching people, and supporting federation of PIDS services in a topology-independent manner. Although it provides the confidentiality and security mechanisms, PIDS itself is not required to enforce confidentiality. Its interfaces are delineated so that "request interceptors" can enforce any policy to protect personal information. In addition, PIDS enables plug-and-play interoperability by means of a "core" set of profile elements.

3.2.3 Lexicon Query Service (LQS)

The LQS aims to provide "a set of common, read-only methods for accessing the content of medical terminologies". There are usually some typical uses of medical terminology systems. For example, using terminology services to assist information acquisition, translate coded information from one representation to another for information display, and help indexing and matching. It is a well-known fact that medical terminology systems can vary radically from a simple code-phrase pair list to a huge complicated scheme with heterogeneous representations. Therefore, it is critical to define a standard interface to any of the major medical coding schemes. The LQS specification proposes a reference model representing the common aspects of several different systems. Much of the functionality is optional, which can be implemented by a particular terminology system.

3.2.4 Healthcare Data Interpretation Facility (HDIF)

It is difficult to interpret and integrate data coming from different information systems because of the diverse forms of data. Having standard interfaces for intelligent clinical data transformations will help easily incorporate intelligent systems into HIS to interpret data appropriately. Currently, the healthcare community lacks such a set of standard interfaces. HDIF targets to offer a general-purpose infrastructure in which a variety of intelligent transformations for clinical data will be embedded, as well as standard interfaces that can conduct intelligent healthcare data transformations when it is necessary. Some CORBA services, such as Security, Trader/Naming, and Events/Notification services, could be involved in this facility.

3.2.5 Clinical Observation Access Service (COAS)

COAS is a set of interfaces and data structures related to how a server can supply clinical observations. A clinical observation refers to any information that has been captured about a single patient's medical/physical state and relevant context information. It consists of a significant portion of the information about any patient. Either medical instruments or healthcare professionals may derive this information during the examination of patients. COAS is designed as a "conduit" for medical information that offers opportunity to transfer well-defined data from standards such as HL7 and DICOM. It is desired to provide efficient ways to request and/or receive clinical observations and their associated context information, to be capable of filtering observations in terms of the observed subject, observation type and time frame, to define a query mechanism for observation retrieval, and to address the access policies that specify the range of context information associated with an observation. COAS is expected to determine a default set of observation types, which can be defined as either the format or structure of the observation, or the coding scheme being used in description of observed biological phenomenon. In order to deal with medical terms, a COAS server may specify the coding schemes or vocabularies used within an organization. COAS can make use of LQS to translate one code to another after it is received from a client. Some other mechanisms, such as publishing and subscribing to observations, will also be taken into consideration in COAS.

3.2.6 Healthcare Resource Access Control (HARC)

The information of a patient that can be disclosed or shared by multiple healthcare organizations depends on several factors. For example, who is looking for that information, which legal rights or privilege the information-seekers have on that information, and what are the access policies of the organization that owns the information? In order to meet the requirements of the various groups involved in healthcare process, there is a need to support two groups of access control policies. One group is general policies related to the role of healthcare providers and the control of disclosure or access to patient confidential information. The other group is patient-specific access policies that are based on either the requirements of a particular patient, or the information being accessed or disclosed. CORBAMed HARC tries to define the mechanisms to support those complicated information access control in order to secure the resource. Obviously, the CORBA Security Service plays an important role in HARC.

3.2.7 Clinical Image Access Service (CIAS)

Clinical images, such as X-rays and ultrasound images, are a subset of clinical observations and a critical part of patient medical record. CORBAMed CIAS, which will be a retrieve-only service, deals with mechanisms to access clinical images and related information. CIAS does not try to replace DICOM, which is the most prominent standard for image interchange in medicine. Instead, it intends to encapsulate portions of the DICOM standard as a service wrapper so that it could give a simplified view of DICOM information model that supplies images and limited meta-data to users in the formats compatible with well-known image standards. CIAS will provide automatic image scaling and windowing to meet the needs of non-diagnostic use of medical images by general clinicians or non-image specialists. It will offer the basic services for supporting electronic patient records transmission over low and moderate speed networks.

In summary, CORBAMed framework enables interoperability and integration among heterogeneous HIS by providing a set of standard services. In each service, standard interfaces are well defined and should be implemented by applications.

3.3 The Distributed Healthcare Environment (DHE) Middleware

3.3.1 Overall Architecture of HIS

The structure of individual healthcare centers is evolving toward the integration of a set of specialized departments with their own organizational, logistic and management requirements. Each department has its individual special activities and delivers specific medical services. Thus, the hospital information system can be modeled as a federation of autonomous applications that are optimized in functionality aspects and can mutually interact through protocols and interfaces (GESI, 1997). Currently, there are a number of available applications addressing the special needs of different departments. However, they are mutually incompatible with each other and can hardly be integrated to support the overall healthcare requirements due to the lack of standards. The strongest need is therefore to enable the interoperability among different applications. Although the applications might be developed by different vendors using disparate technologies, they can be "glued" together and behave like a whole towards the environment.

In order to be capable of behaving as a single system with respect to the overall organizational functioning, these applications rely on a set of common services and must fit into a consistent and comprehensive system architecture. This architecture can be decomposed into a set of components in terms of the objectives and interactions with the rest of the system so that changing one piece of the system will remain other parts unaffected. Considering these factors, CEN/TC251/PT010 (GESI, 1997) defines a conceptual architecture framework structured in three layers as follows:

- Application layer

- Middleware Layer
- Bitways Layer: to support physical connection, and meet network and distribution requirements

Correspondingly, the components of hospital information system are grouped into three layers:

- The basic technological layer is responsible for providing the interworking services in a distributed and heterogeneous environment.
- The layer of the common services implemented through Distributed Healthcare Environment (DHE layer) provides a set of common services and interfaces to the applications in order to ensure the appropriate interworking and synchronization of them. It supports applications from both the functional and the procedural point of views.
- The layer of application: the applications are designed and optimized to support specialized healthcare-related activities.

DHE middleware is a proprietary product of an Italian firm called Gestione Sistemi per l'Informatica SrL (GESI). The overall architecture of HIS containing DHE middleware is described in Figure 2. It consists of distributed databases, a set of application services, and interface modules in order to address various needs of the healthcare organization. It provides an open infrastructure that is capable of federating heterogeneous applications through a common set of healthcare-specific components.

3.3.2 The Applications

The applications are developed to support healthcare and medical activities, including clinical activities, and monitoring and optimization of the services provided. The classification of the application areas in the healthcare organization can be defined using a functional user-centered approach, which identifies individual users and activities that need to be carried out. Each basic application module of the system is responsible for one service/activity. This approach increases the flexibility and maintainability of the system. Generally, there are six functional application areas in healthcare center, namely patient management, medical care, nursing, medical support, ancillary services and organization and management.

3.3.3 Components of DHE Middleware

The DHE middleware comprises the basic functional infrastructure of healthcare centers. The fundamental role of DHE middleware in the HIS is to provide mechanisms to support different applications in accessing, sharing, manipulating and managing common data of the organization while maintaining the consistency and integrity of information. The DHE includes following components:

- An integrated but distributed database that contains the commoninformation. The DHE could manage five primary types of data: Activities and organizationdata describes the information related to the lifecycle of all organization activities. Resources data depicts the available resources within the healthcare center, such as personnel, instruments, beds, and drugs. Users and authorizations data contains information about the authorization of individual users of the system; Patients data handles personal, administrative and epidemiological data of the patients; Health data manages all available health data regarding patients, including clinical data.
- A set of servers that encapsulate the database and provide complete services to enable applications to enter, retrieve and manipulate the common data. Each server is self- consistent and is capable of concurrently interacting with multiple client applications.
- A set of configuration modules that allow the authorized users to configure the whole system through a GUI.
- A dictionary that assists developers and end-users to understand the overall functioning of the system.

3.3.4 The DHE Middleware Functional View

As explained, the DHE services can be directly used by applications to manage all information. To manipulate data, the DHE relies on the presence of a database management system. When the DHE server starts, it needs to select the port on which it will wait for the requests from clients. On the other side, client takes the similar procedure when it tries to contact with DHE server for the first time. The DHE server supports two types of connection modalities with the client. One is usual RPC interaction; the other is Internet-suitable interaction ? the socket is closed when a service is completed. The DHE provides APIs for both procedural (C,Visual Basic) and object-oriented (Java, CORBA, C++) programming languages so that Web applications can be developed easily with DHE.

Individual applications must interact with each other through DHE services and must refer to the data managed by DHE for consistent information. With respect to the use of DHE services by applications, there are two main approaches: DHE federation and DHE integration. DHE federation requires creation of interfaces between applications and DHE common services that represent components external to the application themselves. DHE integration implies the utilization of DHE services directly within the architecture of the applications to support activities.

The DHE allows the developers to adopt its own security mechanisms for establishing encrypted communications over the healthcare network. One of the most important issues in distributed healthcare is how to improve the quality of care while maintaining the patient privacy and confidentiality. In the DHE, several security features are implemented (Andany J, et al., 1999):

- Encryption/Decryption modules based on private and public keys can be activated on both client and server sides.
- Digital signature can be added when validating or entering health data.
- Certification of the user activities is provided for auditing to keep track of all interactions occurred that are signed and encrypted by the user.

Sitting between platform and application layers, the DHE middleware serves as a common glue to integrate loosely scattered applications into a highly effective network with interoperability.

4. CONCLUSION

The changing focus in the business of healthcare from hospital-centered to patient-centered has dramatically emphasized the urgent need of interoperability between heterogeneous healthcare information systems. Due to the high heterogeneity on hardware platforms, operating systems, programming languages, and data standards in HIS, there must be a consensus on common interface architecture to meet the interoperability challenges and improve the quality of care. Middleware service has been recognized as an ideal approach to connect different healthcare applications. Through the investigation of two prominent information system architectures adopting middleware service, CORBAMed and DHE, we believe that the middleware service approach can satisfy the various requirements of a distributed healthcare information system because of its reaching maturity and availability.

Appendix

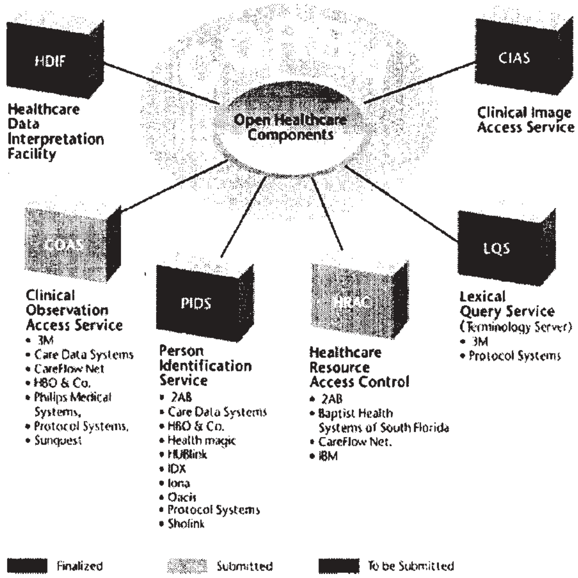

Figure 1. CORBAMed Framework (ACL, 1998)

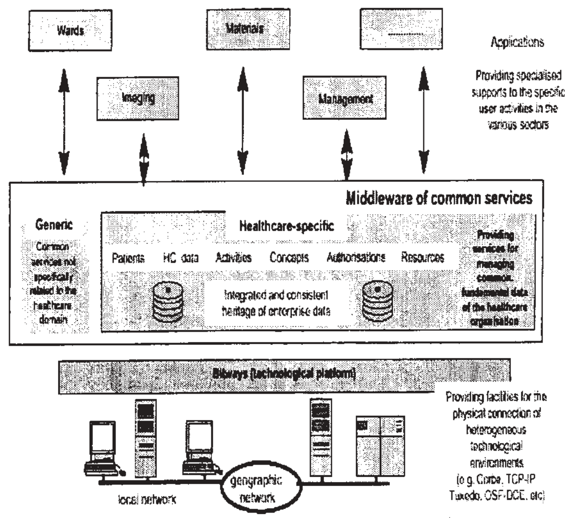

Figure 2. The overall layered architecture of healthcare inforrmation systems (with the DHE middleware) (GESI, 1997)

1. CORBA is a trademark of Object Management Group, Inc. in the U.S. and other countries.
2. CORBAMed is a trademark of Object Management Group, Inc. in the U.S. and other countries.

REFERENCES:

Andany, J., Bjorkendal, C., Ferrara, FM, Scherrer, J. and Spahni, S. (1999). Authorization and security aspects in the middleware-based healthcare information system. In: Proceedings of Medical Informatics Europe 99, IOS Press, Amsterdam, pp.315-320

ACL (Advanced Computing Laboratory, 1998). CORBA-based systems. http://www.acl.lanl.gov/telemed/Corba_Main.html

Bernstein, P. A. (1996). Middleware: a model for distributed system services. Communication of the ACM, 39(2): 86 – 98

Forslund, D. W. (1997). The Role of CORBA in Enabling Telemedicine (LA-UR-97-1010) Global Forum III: Telemedicine in Vienna, Virginia.

Forslund, D. W., George, J. E., Stabb, T., et al. (1999). Federated healthcare utilizing CORBAMed standards. Presented at RSNA Annual Conference, 11/29-12/3, Chicago, IL.

GESI (1997). The DHE Middleware Introduction (Version 1.1). http://www.gesi.it/gesi.htm

OMG (1997). CORBAMed: Healthcare Domain Specifications. http://www.omg.org/corba/cmfull.html

Orfali, R. & Harkey, D. (1998). Client/Server Programming with JAVA and CORBA. Second edition. New York: J.Wiley & Sons