

FUNDAMENTAL COMPONENTS FOR THE REALIZATION OF A FEDERATED INTEGRATED ELECTRONIC HEALTH RECORD ENVIRONMENT

D. G. Katehakis¹, S. Sfakianakis¹, M. Tsiknakis¹, S. C. Orphanoudakis^{1,2}

¹Center of Medical Informatics and Health Telematics Applications (CMI-HTA), Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH), E-mail: {katehaki, ssfak, tsiknaki, orphanou}@ics.forth.gr

²Department of Computer Science, University of Crete, Heraklion, Crete, Greece

Abstract - The Integrated Electronic Health Record (I-EHR) is a term used to describe the whole set of information that exists in electronic form and is related to the personal health of an individual. Any approach towards I-EHR focuses on the needs of professionals or citizens who want a uniform way of accessing parts of personal health record information that is physically located in disparate information sources. Any I-EHR end-user environment must provide fast, secure and authorized access to the distributed fragments of the electronic patient record (EPR) originating at multiple clinical information systems, and to deliver them in a multitude of formats. The importance of such an environment becomes apparent when used in conjunction with a number of advanced telematic services, such as medical collaboration, home care monitoring, and/ or health emergency services, to provide seamless care without visible organizational boundaries.

Keywords - Integrated Advanced Information Management Systems, Federation of Clinical Information Systems, Integrated Electronic Health Record

I. INTRODUCTION

Although each healthcare facility is autonomous and devoted to the delivery of a particular set of services, continuity of care requires that different healthcare facilities, offering complementary services or different levels of expertise, exchange relevant patient data and operate in a cooperative working environment. In this environment, diverse user groups require secure, customizable access and sharing of information residing at geographically distributed autonomous information systems. This sharing of information resources is generally accepted as the key to substantial improvements in productivity and better quality of service. Therefore the I-EHR is the cornerstone for the provision of continuity of care, and is the point of reference for any exchange of information related to the patient under consideration.

Several approaches are in use today for sharing distributed fragments of health-related information. Although the use of *message based communication* of Electronic Healthcare Record (EHR) data is used today extensively, such an approach suffers when the number of the involved autonomous Clinical Information Systems (CIS) increases. On the other hand, the use of centralized *Clinical Data Repositories* (CDR), containing data from multiple sources, usually have difficulty with data context and codification, and their complexity of design in most cases leads to extensive delays in actual implementation and use. In addition, *monolithic information systems*, that have mechanisms embedded in their structure for directly accessing host systems, are not open and therefore they are not scaleable and easily maintainable. The most promising approach, which will be examined in the rest of the paper, is

based on a federation of autonomous CISs and an underlying health care information infrastructure (HII) that facilitates the virtual view of the I-EHR, based on an open architecture that provides the framework for the reuse of public interfaces and common components.

Eventually, any successful I-EHR environment ought to be able to deliver a complete view of personal health histories to be accessed from many different in nature locations, 24 hours a day. Usually, this interaction may require the intervention of a third person (e.g. when primary health information does not exist in electronic form). Naturally, this should not happen, and the process of getting access to information should be kept as simple as possible. PC is the preferred platform, although Web based access should be supported. Handheld devices (PDA's) or Satellite TV may be required in certain cases for instant access to missing pieces of information.

II. METHODOLOGY

The federated approach towards the I-EHR is based on a set of fundamental principles that include:

- 1) *Use of Open standards* to promote interoperability among multi-vendor applications and services.
- 2) *High Quality Services* that could give value to the final system and attract people to using it.
- 3) *A Modular Architecture* to facilitate the development, maintenance and evolution of scaleable, secure, effective, and affordable systems.
- 4) *Incremental Evolution*, which enables the “monotonic” evolution of systems, building upon existing functionalities while adding new capabilities as they become available.

The above design principles guide the effort towards a reference architecture for the provision of integrated user-oriented telematic services in integrated regional health telematics networks, such as HYGEIAnet [1]. Following the Modularity principle [2], a number of components for the I-EHR have been identified and a component-based architecture has been designed and built [3]. A clear distinction has been made between generic and I-EHR specific components.

Generic components include services designed to perform:

- 1) *Patient Identification* (PID) to be used for identifying patients based on their demographic data and correlating their ids across different clinical information systems.
- 2) *Auditing* (AUD) for recording all performed interactions between middleware services and/ or end-user applications, either directly or indirectly.
- 3) *Authentication* (AUT) to control access to desired services, only for those that are properly authorized.
- 4) *Encryption* (ENC) for the secure communication of sen-

sitive personal information over and across any healthcare domain related Virtual Private Network (VPN), as well as the Internet.

5) *Resource Location (REL)* for identifying availability of related resources, such as organizations, devices, or software and the means for accessing them.

6) *Terminology (TER)* to interpret and translate terms between different coding schemes, terminologies and internal semantics, schemas.

7) *User Profile (UPR)* to track the long-term interests of users and to maintain personalized settings and preferences.

I-EHR specific components have been identified to be the following:

1) The *I-EHR Indexing Service (IS)* used for addressing the issue of locating fragments of primary health information maintained by different CISs.

2) *Primary Health Information Access Services (AS)* for the direct access to the primary healthcare clinical systems where the complete, original (physician generated), clinical information is kept.

3) The *I-EHR Update Broker (UB)* used for the propagation to the IS of all modifications pertaining to clinical information.

The synergies of these components are depicted in Fig. 1. In the same figure, the four most important use cases that are labeled as UC1, UC2, UC3, and UC4 represent the extraction of the necessary primary health information from CISs (UC1), its propagation to the IS (UC2), a question directed by end users for the location of primary information (UC3), and direct access to primary health information (UC4).




Fig. 1: I-EHR Components and Synergies.

Generic components like AUT, TER, and REL can be used in conjunction with other services for the implementation of scalable infrastructures for telemedicine, home care, clinical messaging, etc.

A. The I-EHR Indexing Service

The IS is the core component of the I-EHR architectural approach followed in HYGEIANet. Its main responsibility is to maintain and index clinical information that is actually managed by the Heterogeneous Autonomous Decentralized Systems (HADS) of the federation. These CISs are the primary sources of clinical information, i.e. the systems that are located in healthcare facilities (hospitals, clinics, etc) and where patients' clinical data are stored. In practice, these

systems exhibit a lot of diversity: heterogeneity that spans from the software and hardware platforms to the semantics of the information they store and its internal representation. Although the heterogeneity of platforms could be tackled by a component integration technology such as CORBA, J2EE or DCOM, the diversity of local schemata ought to be handled through a federated schema and the consequent mappings.

Apparently, two extreme cases exist: In the first case, the federated schema is extremely simple in order to make it feasible for as many CISs as possible to become parts of the federation. In the second case, the federated schema carries the whole semantic information maintained by each of the participating CISs. Advantages and disadvantages exist for both approaches and the right approach lies where the definition/adoption of a federated schema is capable of supporting effective solutions to the end user's immediate needs without imposing significant constraints in dealing with the issue of incorporating new systems in the federation.

In the case of HYGEIANet, the IS follows the conceptual model depicted in Fig 2. This conceptual model has been devised in order to be generic and allow individual implementations to choose the level of detail that ought to be reached, and to support incremental scaling of the complexity of the underlying conceptual model.

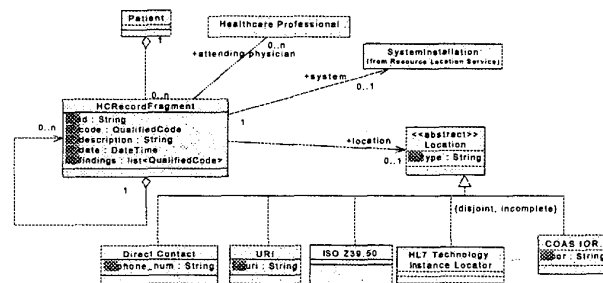


Fig. 2: The Conceptual Model of the IS.

Important entities depicted in Fig. 2 are the following:

1) *Patient*: the entity that corresponds to the client of a medical act, such as examination, treatment, or medical encounter.

2) *Healthcare Professional*: the entity that corresponds to the attending physician that has performed the medical act.

3) *System Installation*: the entity that represents the CIS where primary health information is stored.

4) *HCRRecordFragment*: the entity that contains indexing to clinical information and the context in which the corresponding medical act took place. It corresponds to a representation of the clinical findings that were the output of the communication between the subject of care (i.e. the patient) and the involved healthcare professionals. Indexed information is contained as a list of qualified codes indicating existence of specific types of clinical information without immediate knowledge of the corresponding actual values. This entity may be composite in the sense that it contains other (more atomic) *HCRRecordFragments* and thus builds a "composition tree" of clinical information fragments pertinent to the specific patient.

5) *Location*: the entity that represents a location from where detailed information about a *HRecordFragment* can be retrieved. This can be for example the CORBA IOR of a COAS server [4] from where detailed clinical observations can be accessed, an HTTP URI or just the phone number of the corresponding healthcare facility.

The generality of this model is enough to be able to incorporate different schemata, and to be deployed across different settings, under the hierarchical model of Fig. 2. In addition, provided that terminology neutrality is desired, the use of terminological services, which are implied by the use of qualified codes, is inevitable.

B. Primary Health Information Access Services

Access to primary health information is vital for obtaining a first hand view of the actual source for decision support. This type of information extraction can be provided to end-users either unformatted (e.g. by means of HTML), or in a structured way. In the latter case the existence of standardized interfaces (such as HL7, OMG COAS, XML, etc.) is a prerequisite for any CIS of the federation that wants to be able to deliver semantics together with fragments of information. From this point on AS will refer to any alternative means for the provision of structured and semantically consistent access to primary information.

C. The I-EHR Update Broker

The UB is another fundamental component of the I-EHR infrastructure. This component is responsible for keeping the IS up-to-date with new or updated information, and maintain a loose consistency between CISs and the IS. Each CIS is associated to a single UB that, either periodically (on pre-determined time intervals), or on demand submits properly formatted information to the IS.

A prerequisite for the UB to properly update the IS, is the existence of a mechanism that is capable of identifying modifications in the CIS managed. Also the means for accessing the modified information and identifying not only the change but also the kind of the change ought to be available e.g. if new information has been inserted, old information has been deleted or just some information has been altered.

Because the AS, in the context used, is not designed for providing update/ notification information, it becomes necessary for the UB to maintain local indexing information for all the CISs of the federation that it is responsible for. This local indexing information, called *Update Cache* from now on, is actually the partition of the IS that corresponds to the specific CIS. In this sense, the UB keeps replicated data of a subset of the IS. The goal of Update Cache is to compare new information with the existing one in order to find:

- 1) The existing modifications.
- 2) The kind of these modifications, i.e. "insert", "delete" or "update".

Whenever a modification is submitted to the IS and is committed, the UB consequently updates its own Update Cache in order to keep it synchronized.

Another required entity of the UB is the *Update Queue*, which is a queue of the changes that need be submitted to the IS. The UB, after identifying the modifications required, it inserts them in the Update Queue. Whenever it is decided that the IS be updated, update information is removed from the Update Queue and transmitted to the IS. This design allows for the decoupling of the IS update from the update of the UB itself, making it possible to impose different policies in each of them.

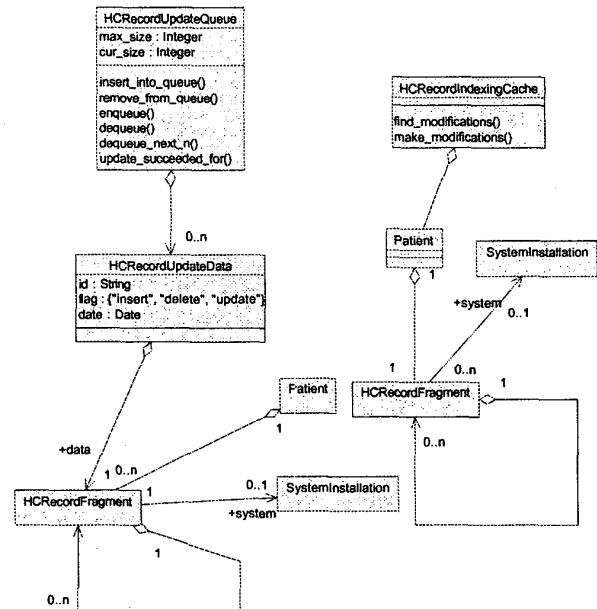


Fig. 3: The Conceptual Model of the UB.

While the IS update by the UB is quite straightforward, the update of UB itself needs more elaborate consideration due to the communication with each CIS that are in essence "legacy systems". Two alternative policies that can be employed make use of the *pull model* and the *push model*. In the pull model the UB actively asks the CIS using the AS as a gateway for collecting all available information. In the push model the UB is passively kept up-to-date, which means that the CIS sends the update information to the UB on demand. In this model an implemented event/ notification functionality of the AS could be used to support these asynchronous updates of the UB.

Both of these models are useful in practice. The pull model is more general in that the UB can extract all the modifications performed inside the CIS. In the downside, the whole data kept in a CIS is a huge volume of data that must be transmitted to the UB where the differences must be determined based on the current and the previous datasets.

In the push model, only the modifications are transmitted to the UB, so the whole process is more efficient. The disadvantage though is that there is no general way to get modifications that have to do with "deletions" i.e. with information that does not exist any more. Following the event/ notification way of communication in the AS, a client can register its

interest in receiving specific kind of information, e.g. for a specific patient or a specific type of clinical information, or information that satisfies some criteria. In all cases the notifications delivered to the client deal with new or updated information and not with information that is obsolete.

Since both of these models are needed, a combination of the two is a rational decision: the push model can be used for frequent updates and less frequently the pull model as a general check and repair method for maintaining consistency.

III. RESULTS

HYGEIAnet has put in place the required infrastructure for the deployment of the I-EHR at a regional level, and a large number of CISs systems have been supported, ranging from primary health care and nursing to departmental information systems for pathology, cardiology, radiology, laboratories, etc [5]. Technologies linked to the corresponding execution architecture include CORBA, DCOM, X.500, SQL/ ODBC, Java and XML.

The use of public interfaces, when available, is of paramount importance. According to Hopkins "The emphasis on well defined interfaces, separate from their implementation, is critical to the success of components in loosely coupled systems" [6].

Around the CORBA technology various standards have been emerged, with the specifications delivered by the Healthcare DTF of OMG being the most important of all [7]. Additional efforts have been studied and employed, mainly related to HL7 CDA [8] and CEN ENV 13306 [9].

Also XML has been extensively exploited to deliver clinical information in representations customized to user environments and the Human Computer Interaction (HCI) portability is handled by virtue of the employment of the Java platform (Fig. 4).

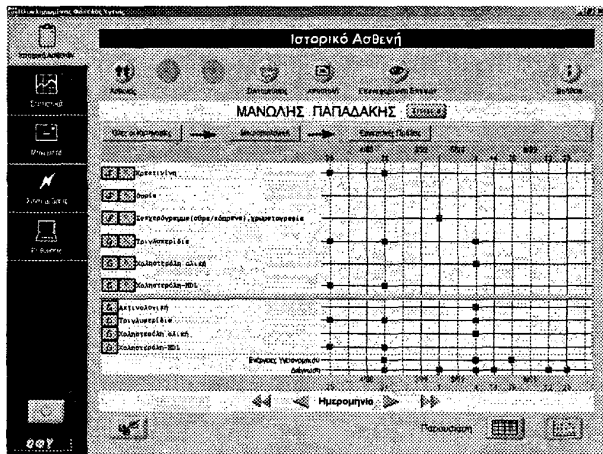


Fig. 4: The I-EHR HCI environment in HYGEIAnet.

Due to the fact that users seek selective information, following specific paths depending on their personal preferences, it is expected that the I-EHR concept will eventually lead to a uniform applications and services environment.

Important parameters that have been identified as been important from the technical point of view, and are directly linked to the acceptance of the service by the end users include integration with the overall health infrastructure, adherence to standards, as well as usability of the provided services.

IV. DISCUSSION

The I-EHR environment, as it has been developed and set up, provides a decentralized view of the patient's medical record, by dynamically composing information that resides in a variety of heterogeneous clinical information systems. Under a secure Internet/ Intranet environment, the full personal health history can be rapidly collected and composed totally transparently and sent to the authorized health professional. This is enabled by means of integrated regional health telematics networks such as HYGEIAnet, since they enable citizens to access information and services without visible organizational boundaries, providing decentralized healthcare through integrated services for seamless and personalized information delivery.

ACKNOWLEDGMENT

The work presented in this paper is funded in part by the European Commission through the PICNIC project (IST-1999-10345).

REFERENCES

- [1] M. Tsiknakis, C.E. Chronaki, S. Kapidakis, C. Nikolau, and S.C. Orphanoudakis, "An Integrated Architecture for the Provision of Health Telematic Services through Application of Digital Library Technologies", Intl. J of Digital Libraries, Vol. 1(3), pp 257-277, 1997.
- [2] C. Szyperski "Component Software: Beyond Object-Oriented Programming", Addison Wesley Longman Ltd, 1998.
- [3] D.G. Katehakis, M. Tsiknakis, S.C. Orphanoudakis: "Enabling Components of HYGEIAnet". Proc. of TEPR 2001, Boston, MA, pp. 146-153, May 8-13, 2001.
- [4] Object Management Group, Healthcare Domain Task Force, Clinical Observation Access Service (COAS). URL: <http://www.omg.org/cgi-bin/doc?formal/2001-04-06>
- [5] D.G. Katehakis, P. Lelis, E. Karabela, M. Tsiknakis, S.C. Orphanoudakis: "An Environment for the Creation of an Integrated Electronic Health Record in HYGEIAnet, the Regional Health Telematics Network of Crete". Proc. of TEPR 2000, San Francisco, CA, Vol. 1, pp. 89-98, May 9-11, 2000.
- [6] J. Hopkins "Component Primer", Communications of the ACM, vol. 43(10), pp 27-30, October 2000
- [7] Object Management Group, URL: <http://www.omg.org>
- [8] Health Level Seven, SGML/XML Special Working Group, Clinical Document Architecture (CDA). URL: <http://www.hl7.org/special/committees/sgml/sgml.htm>
- [9] Comité Européen de Normalisation (CEN), Technical Committee 251, Working Group I, Information Models. URL: <http://www.cen251.org/>