A PROTOCOL FOR DYNAMICALLY MAINTAINING AN INTEGRATED MEDICAL RECORD DATABASE FROM A SET OF DIFFERENT MEDICAL RECORD DATABASES

Emilio Hernández and Mireya Morales
Department of Computing, Simón Bolívar University, Caracas, Venezuela

Keywords: Health information systems, Medical record databases, Interoperability.

Abstract: Merging distributed medical record databases into an integrated database is an alternative approach for the implementation of the unified medical record concept. One of the main advantages of this approach is that the integrated database could be a backup for medical record databases. Additionally, the integrated database may be useful for medical research and may help improve the decisions on public policies related to health care. We propose a protocol for dynamically maintaining an integrated database that is the merge of different medical record databases, under some limitations related to security and control of the medical record.

1 INTRODUCTION

Hospital Information Systems are usually implemented as client-server or web applications (Morrison et al., 2010). These applications are suitable as a solution for single administrative domains. A National Health Service face the problem of fragmentation of the health-related information among many Hospital Information Systems. Fragmentation leads to lack of coordination, information inconsistencies and many other problems that make it difficult to define appropriate public policies. In addition, information integration is one of the keys for the improvement of the consumers health care experience.

Interoperability has been pointed out as a key for the success of a national or regional Health Information Service (Brailer, 2005; Walker et al., 2005). The efforts have mainly been focused on the definition of information exchange standards (HL7, 2010; ANSI, 2010). It is frequently assumed that cooperating data sources exchange information horizontally, probably with the help of data discovery services (Zhang and W. Xu, 2007). There are also some proposals that additionally consider the semantic level for data association (Lopez and Blobel, 2009; Hovenga, 2008).

We propose a specific scheme for interoperability, based on the integration of databases into a single medical record database, with replication of the information. This approach needs the definition of a data exchange protocol that ensures some control on the replicated information, in order to avoid inconsistencies and provide reasonable security. This interoperability scheme resembles a memory hierarchy in which an upper level integrates the databases of several institutions. In the context of a country, for example, the number of integrated databases could be much lower than the number of health centre databases (in fact, there could be a single integrated database). At this integration level the horizontal interoperability between large database systems becomes easier because there is a lower number of databases and fewer data formats.

Along with the memory hierarchy scheme, we propose a protocol for dynamically maintaining a medical record database that is actually the merge of different medical record databases, under some limitations related to security and control of the medical record "ownership".

In this paper we describe the proposed interoperability scheme and the associated synchronization protocols. Section 2 presents the components of the interoperability scheme, section 3 describes the protocols and section 5 shows the conclusions.

2 HEALTH INFORMATION SYSTEMS AS MEMORY HIERARCHIES

At the heart of a Health Information System is the medical record database. In our model, this database

Hernández E. and Morales M..
A PROTOCOL FOR DYNAMICALLY MAINTAINING AN INTEGRATED MEDICAL RECORD DATABASE FROM A SET OF DIFFERENT MEDICAL RECORD DATABASES.
DOI: 10.5220/0003173705920595
In *Proceedings of the International Conference on Health Informatics* (HEALTHINF-2011), pages 592-595
ISBN: 978-989-8425-34-8

is defined as the union of the databases located at the Local DB Level (Local Database Level). These local databases can be deployed, for example, at the computer departments of large hospitals. The local databases are integrated in a single database at the Integrated DB Level (Integrated Database Level) by providing a copy of each medical record, probably according to strict security policies. The database of the Integrated DB Level could be, in turn, a distributed database with the possibility of offering a single view.

In this sense the whole system resembles a memory hierarchy. The main difference with a computer memory hierarchy would be that the control of the original version is not at the upper level. Local DB Servers have the ”ownership” of the medical records, in such a way that they could even decide whether to prevent some medical records from being uploaded to the integrated database.

We define three levels in which synchronization software must be implemented for moving data upwards and backwards in the memory hierarchy. These levels are:

- **Access Level.** Basically the access level is composed of the end-user systems that allow medical staff to interact with the system, as well as monitoring systems that send patient information from home or from mobile devices. If the computers at the Access Level can work standalone during some periods of connectivity failure, data collected locally must be synchronized later with the Local DB Level.
- **Local DB Level.** This level is composed of the basic data centres where the data are actually stored. The combination of the Access Level and Local DB Level basically conforms to a client-server model with fault tolerance features. Local DB Servers may be located within the health centres, if they have the capacity for administering a small data centre. The data formats or the client-server software used by different Local DB Servers may vary from site to site.
- **Integrated DB Level.** This level contains, as a basic unit, a database with the information of several Local DB Servers. Local DB Servers must periodically synchronize their databases with the Integrated DB Level, using a synchronization protocol that performs data format transformations if necessary. Several Integrated DB Servers may interoperate under a secure platform, for instance a Grid Platform (Foster and Kesselman, 1999), if a higher degree of integration has to be achieved.

Between the health centres at the Access Level and the Local DB Level the data will usually flow through an on-line web application. Between the Local DB Level and the Integrated DB Level the data will flow through synchronization protocols.

We are not defining interoperability protocols between Local DB Servers. This means that if a patient must be attended at a health centre far from home, the assigned Local DB Server must obtain the information related to the patient (i.e. the medical record) from the Integrated DB Level.

The Access Level includes all systems related to home or mobile monitoring of outpatients. In this case, the Local DB Level must provide the services for receiving data from the devices that collect the monitoring information.

3 THE SYNCHRONIZATION PROTOCOLS

In this section we describe the protocol for dynamically maintaining the medical record database of the Integrated DB Level by merging different medical record databases at the Local DB Level. Such a protocol needs to be defined and executed under some limitations related to security and control of the medical record ”ownership”.

In the proposed protocol there are only two allowed operations.

1. Synchronization of a Local DB Server with the Integrated DB Server.
2. Query of a single medical record to the Integrated DB from the Local DB Server. This query should be used only in those cases in which a patient is attended at a health centre located away from that in which he/she is registered, for instance, during an emergency.

3.1 Synchronization Protocol

The database synchronization protocol is inspired by well known file synchronization protocols. In our case, the protocol synchronizes records rather than files.

The synchronization protocol, depicted in Figure 1, works under the following conditions:

- Patients must be identified according to universal identity keys. This restriction may prevent the implementation of an Integrated DB Level beyond the limits of a single country or region.
- The patients must be registered just in one Local DB Server. Conflicts between different versions of a patient's medical record could be solved with the assistance of authorized users.

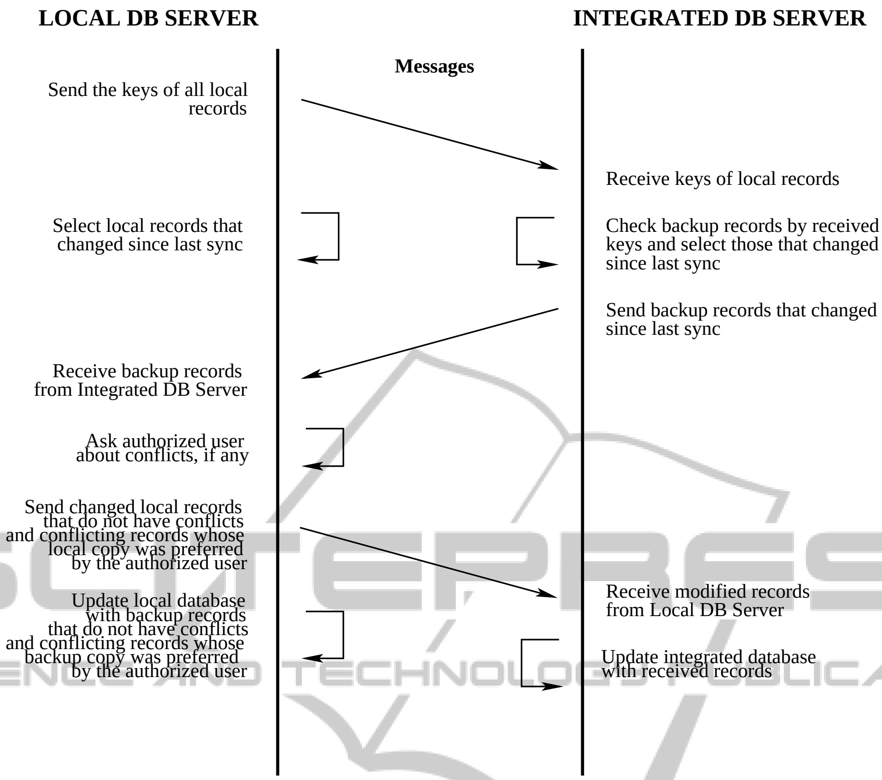

Figure 1: Synchronization protocol between a Local DB Server and the Integrated DB Server.

- The medical record is stored in the Local DB Server. The copy of the medical record at the Integrated DB Server is considered a backup.
- Different Local DB Servers may potentially have different DB formats.
- The Integrated DB records should be designed with an inclusive criteria, i.e. ideally they must contain all of the fields that could be present in any of the Local DB records.

A key consideration for the protocol implementation is that it needs a database subsystem that keeps track of the record versions. Every time that a medical record is modified, an associated version number field must be increased. Fortunately this requirement can be easily implemented with the help of *triggers*, which most Database Management Systems offer. A separate table indexed by the medical record key may contain the version number field and can be added to both the Local DB Server and the Integrated DB Server. A trigger must increases the version number on every record update.

The synchronization protocol is periodically launched by the Local DB Server. It starts by sending the keys of the patients locally registered to the Integrated DB Server. Only these records are going to be synchronized with the Local DB Server. Once both parts find out which of these medical records have new versions, they proceed to the information exchange. If a medical record has new versions in both sides (probably because the patient was attended in more than one Health Centre within a short period of time), an authorized user must decide what information should be stored in both the Local DB Server and the Integrated DB Server.

3.2 Querying the Integrated DB Level

If a patient is attended in the health centre where he/she is registered, the medical record is available locally and there is no need to query the Integrated DB Server. This query is needed if the patient is attended in a health centre whose Local DB Server is different from that in which he/she is registered. This may happen, for instance, in an emergency far from home. The query from a Local DB Server to the Integrated DB must meet the data acces policies defined for the Health Information System. The medical record is downloaded and a temporal medical record is created in the Local DB while the patient is attended. This record may be modified in the Local DB. During the next synchronization of this Local DB Server, the modified temporal record is uploaded to the Integrated DB Server and deleted from the Local DB. Later on, the new version of the medical record will be downloaded from the Integrated DB to the Local DB where the patient is registered, keeping a copy at the Integrated DB Level.

4 IMPLEMENTATION

The synchronization protocol must be implemented for each Hospital Information Systems, because they may have different Database Management Systems, of different database definitions. An analysis at the semantic level must be carried out in order to define table and field mappings between the Local DB and the Integrated DB. The implementation of the synchronization protocol must include the version tracking subsystem described in section 3.1. The synchronizer itself follows a client-server model in which both parts must authenticate each other, for instance by means of digital certificates.

In practice, the execution of a the synchronization protocol with several Local DB Servers could unveils the presence of many medical records for the same patient, loaded from different health centres. Once the patients are assigned their Local DB Servers, multiple medical records associated to them should progressively be removed. However, in principle, the risk of having duplicate medical records for the same patient is independent of the interoperability approach.

4.1 Prototype

As a proof of concept, a prototype of a Hospital Information System and a Synchronization Tool have been developed. The Medical Record and Interconsultation subsystems are tightly coupled. The interconsultation subsystem does not work at real-time. It is a store-and-forward telemedicine subsystem that allows auxiliary physicians in remote health centres to make interconsultations with other physicians or medical specialists, usually located in large hospitals. A database synchronizer is provided, developed with an ad hoc data exchange protocol. However, other database synchronizers could be developed following a standars data exchange protocol such as HL7 (HL7, 2010).

5 CONCLUSIONS

In this article we describe an interoperability scheme for Health Information Systems, based on the integration of databases into a single database. The data is actually replicated in the integrated database. This approach needs the definition of data exchange protocols that ensure the control on the replicated information, in order to avoid inconsistencies and provide reasonable security. This interoperability scheme resembles a memory hierarchy in which an upper level integrates the databases of several local database servers.

An integrated database centre located at the top level of the memory hierarchy will serve as backup level for local data centres. Additionally, this upper level will provide the services needed for policy making at the health authority level, and will be used for building data sets for scientific research.

REFERENCES

ANSI (2010). The ANSI ASC X12. http://www.x12.org/, accessed on July 2010.

Brailer, D. J. (2005). Interoperability: The key to the future health care system. *Health Affairs*, 10(1377).

Foster, I. and Kesselman, C. (1999). *The Grid: Blueprint for a New Computing Infrastructure*, chapter Computational Grids, pages 15–51. Morgan Kaufmann Publishers, Inc.

HL7 (2010). Health Level Seven. http://www.hl7.org/, accessed on July 2010.

Hovenga, E. (2008). Importance of achieving semantic interoperability for national health information systems. *Texto & Contexto Enfermagem*, 17(1):158–167.

Lopez, D. M. and Blobel, B. G. (2009). A development framework for semantically interoperable health information systems. *International Journal of Medical Informatics*, 78(2):83–103.

Morrison, C., Iosif, A., and Danka, M. (2010). Report on existing open-source electronic medical records. Technical Report UCAM-CL-TR-768, University of Cambridge, Computer Laboratory.

Walker, J., Pan, E., Johnston, D., Adler-Milstein, J., Bates, D. W., and Middleton, B. (2005). The value of health care information exchange and interoperability. *Health Affairs*, 10(1377).

Zhang, J. and W. Xu, D. E. (2007). System interoperability study for healthcare information system with web services. *Journal of Computer Science*, 3(7):515–522.