eHealth2015 – Health Informatics Meets eHealth
D. Hayn et al. (Eds.)

doi:10.3233/978-1-61499-524-1-211

Prefetching of Medical Imaging Data Across XDS Affinity Domains

Emmanuel HELM[a,1], Andreas SCHULER[a], Oliver KRAUSS[a] and Barbara FRANZ[b]
[a]*FH Oberösterreich, Forschungs & Entwicklungs GmbH, Austria*
[b]*FH Oberösterreich, Fakultät für Informatik/Kommunikation/Medien, Hagenberg, Austria*

Abstract. Prior studies as well as medical imaging data are crucial for a radiologist to diagnose a patient. In this paper the radiological workflow is analyzed from a patient's perspective in order to gain knowledge on how possible existing prefetching strategies still can be applied in connection with a standardized distributed health information system conforming to architectures defined by IHE and ELGA. As a result an adaption to such architectures is proposed and further evaluated in a testing environment. Although the approach presented works in terms of prefetching relevant prior studies together with medical imaging data, additional research has to be carried out on how to apply intelligent search strategies in order to narrow retrieved results concerning their possible utilization for a specific diagnosis.

Keywords. Health Information Exchange [L01.700.253], Software Design [L01.224.900.820]

1. Introduction

State of the art medicine is highly dependent on modern imaging techniques. With the technological shift from film-based to digital radiology new practices and business models evolved. Medical images can be stored, processed and examined beyond organizational and geographical borders. With the introduction of the Austrian electronic health record (ELGA), patients and their doctors will be able to access clinical reports and associated medical images [1].

While loading images from a local image archive is fast, a decentralized wide area network like ELGA combined with large image data leads to longer transmission times. The workflow of radiologists does not take waiting times for loading image data from remote storages into account. To eliminate those waiting times, mechanisms to load reports and images in advance, so-called prefetching techniques, have been developed and established. Knowledge-based systems utilizing information retrieval techniques, artificial neural networks and inductive decision trees have been developed to prefetch the correct images. These approaches presume availability and accessibility of the images and do not consider architectural constraints and limitations. [2] [3] [4]

To provide both availability and accessibility *Integrating the Healthcare Enterprise* (IHE) defines in its *Technical Frameworks* for the *Radiology* and *IT-Infrastructure* domains how to design and implement cross-enterprise information sharing architectures

[1] Corresponding Author: Emmanuel Helm, FH OÖ Forschungs & Entwicklungs GmbH, Softwarepark 11, 4232 Hagenberg, E-Mail: emmanuel.helm@fh-hagenberg.at

for medical documents and images [5]. This paper shows how IHE-based systems like ELGA could be used to enable the prefetching of medical imaging data.

2. Methods

In the research work carried out, an in depth analysis of the radiology workflow was conducted via direct observation in practitioner's offices accompanied by interviews with the personnel to verify the findings. The treatment of 5 different patients was accompanied by members of the research team at 3 different local radiologist's offices. Two of those radiologists were analyzed and the workflows documented in the course of a bachelor thesis [6]. The selection of the radiologists was based on availability and the fact that they had to be customers of our research project partner CAS GmbH.

The analysis was done not only to gain insight in the workflow and the persons involved from a patient's perspective, but also to identify possible timeframes for an automated prefetching mechanism to preload required data ahead of diagnosis. Table 1 shows an overview of the different radiological use cases observed during the patients' treatment. The intention behind the selection of the patients was to cover a wide range of different examinations concerning modality and anatomy.

A literature research was conducted in order to compare and align the findings according to available research results as well as with the requirements given by the *Scheduled Workflow (SWF)* defined by *Integrating the Healthcare Enterprise (IHE)* [7]. Whilst the SWF is used to maintain consistent exchange of information between the different actors in radiology, the process analyzed by the research team covers the radiology workflow from a patient's perspective. This patient-centered analysis aligns with the way *Radiology Information Systems* (RIS) work, always linking emerging data to the respective patient. However there are intersections between SWF and the analyzed radiology workflow that enable verifying the correctness of the latter.

In addition to the process analysis described, the technical requirements for access, storing and sharing of clinical documents in combination with imaging data in context of a distributed healthcare environment were identified. This was primarily done relying on the IHE-profiles Cross Document Enterprise Sharing (XDS), XDS for Imaging (XDS-I), Cross Enterprise Community Access (XCA) and XCA for Imaging (XCA-I) as well as the publicly available information concerning the architecture of ELGA in Austria, which at its core is based on these profiles [1, 5, 7].

Table 1. Overview of the use cases analyzed at the radiologist's office, observed during the patients' treatment.

Use case	Modality	Anatomy
UC.01	radiology	shoulder
UC.02	mammography	mamma
UC.03	sonography	abdominal
UC.04	ct-scan	knee
UC.05	mri	head

3. Results

The application of the methods described allow for the definition of a conceptual architecture capable of prefetching previous medical findings relevant to a current radiological examination and compliant to the requirements defined by IHE and ELGA.

3.1. Radiology Workflow

From an organizational point of view seven major steps can be identified analyzing the radiology workflow. Figure 1 shows those seven steps and the involved persons.

The process starts with the patient arranging an *Appointment*. In all of the 5 use cases the patient calls the clinic and the doctor's receptionist asks a few questions regarding the type of examination. This is necessary because certain examinations always take place at the same time slots. In traditional radiology information systems this is the time when prefetching from local storage systems (e.g. slow tape robots) is initiated.

At the appointed date the patient arrives at the clinic and registers for the examination. This *Patient admission* is handled again by the doctor's receptionist and it includes the handover of the referral. The referral is the basis for the medical treatment as it defines which examinations should be conducted. At this time the social insurance data of the patient is checked and the treatment relationship officially starts.

While the patient waits for the *Examination* to start the according entries are added to the DICOM Modality Worklist. Depending on the examination, radiological technologists (RT) may have to reconfigure the modalities, e.g. consider the weight and height of a patient for MRI scans. After the examination the RT checks the quality of the results and applies image processing techniques to improve them if necessary. The results are then saved in the local Picture Archiving and Communication System (PACS) and the Radiology Information System (RIS) is notified.

In the observed settings the radiologist starts the *Diagnosis* by scanning the barcode of the patient's file. The diagnostic workstation immediately loads the current examination results and prefetched preliminary findings if available. The studies are presented depending on the radiologist's hanging protocol. Based on this information the radiologist takes a dictation of the clinical findings.

The next step is writing the *Report*. A secretary either listens to the dictation and writes it down or corrects the automated speech to text software. To finish the record the radiologist has to check and *Attest* it.

The final step is to transmit the report to the referring physician. Directed digital transmission and the manual delivery by the patient are common methods of choice.

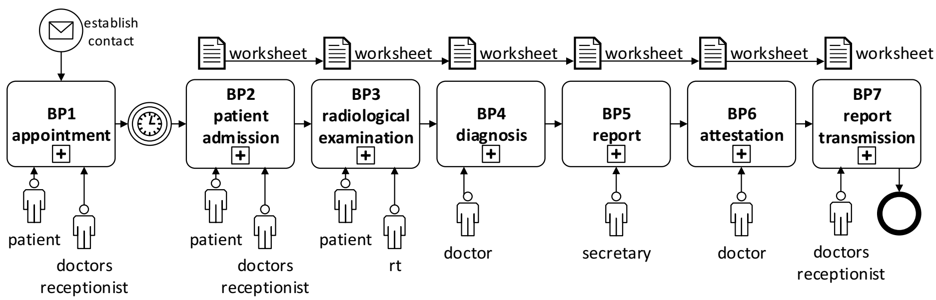

Figure 1. Workflow of a radiological examination.

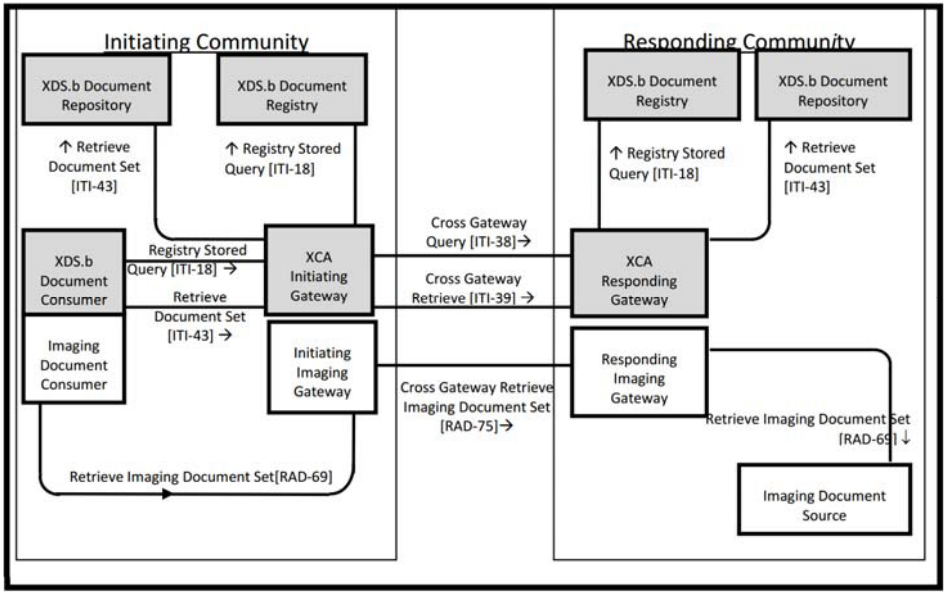

Figure 2. Architecture of a distributed health information system as proposed by [5]

3.2. Architecture

The acquired knowledge concerning the radiology workflow was used in order to examine, whether the architecture proposed by IHE and ELGA fulfills the requirements of the workflow. The combinations of the profiles outlined in section 1 as defined by IHE describe the architecture of a Distributed Health Information System (DHIS) (see figure 2). An example of the application of such a DHIS is the reference implementation provided by ELGA in Austria. Given the preconditions regarding storing and access of medical reports on an international level defined through IHE, and on a national level through ELGA, several additions to the architecture are required in order to successfully perform the workflow with focus on prefetching of prior studies ahead of diagnosis.

Attribute Name	Tag	Imaging Document Source
Study Instance UID	(0020,000D)	R
Referenced Series Sequence	(0008,1115)	R
> Series Instance UID	(0020,000E)	R
> Retrieve AE Title	(0008,0054)	R+
> Retrieve Location UID	(0040,E011)	R+
> Storage Media File-Set ID	(0088,0130)	O
> Storage Media File-Set UID	(0088,0140)	O
> Referenced SOP Sequence	(0008,1199)	R
>> Referenced SOP Class UID	(0008,1150)	R
>> Referenced SOP Instance UID	(0008,1155)	R

Figure 3. Structure of a DICOM KOS Manifest Document as described in [5].

Currently ELGA only defines how to share medical reports via XCA (see figure 2, gray actors). The addition of the XCA-I actors (see Figure 2, white actors) is crucial in order to be able to not only enable a DHIS to offer access to medical reports but to also offer the capabilities to store and retrieve medical imaging data associated with medical reports. Due to the results from the process analysis it can be concluded, that a medical report from a radiologist' point of view is of no use without the imaging data to be used in context of a diagnosis as source of supply. In order to be able to prefetch not only medical reports but in addition also any relevant imaging data an adaption to the architecture is depicted in Figure 2 and described as follows.

As specified by [7] a *DocumentSource* actor registers medical reports as a set of defined metadata that describe the document. This metadata (XDS-Entries) forms the basis for a *DocumentConsumer* actor to be able to retrieve stored documents. In case of imaging data XDS-I proposes to leverage the mechanisms used by the XDS-profile to be able to store imaging data, e.g. DICOM-documents besides CDA-reports (HL7 v3 Clinical Document Architecture). This is achieved through the use of Key Object Selection Documents (KOS), however IHE proposes not to store any DICOM imaging data in a XDS-Repository, but rather store a DICOM-document, that contains metainformation about the imaging data and where it can be retrieved [8]. Using the APND-Relationship it is possible to link an XDS-Entry from a medical report with an XDS-Entry from a KOS-Manifest Document that itself contains the study ID, series ID as well as Object ID of associated imaging data (see figure 3). In order to be more flexible the KOS-Manifest is registered in a separate submission set and afterwards linked to the medical report through the *parentDocument* relationship together with the APND *typeCode* (see figure 4). This leads to the fact, that imaging data is stored separately but at the same time is still accessible from the medical report through the metadata registered at a given XDS-Registry (see figure 4). Hence from a *DocumentConsumers* perspective it can be determined which imaging data belongs to which report, although each of them can be stored separately, e.g. in different repositories.

As a result of the adaption seen in figure 4 to the architecture defined by IHE and refined by ELGA the next section describes the steps involved in order to carry out a prefetching mechanism that is able to search and retrieve medical reports together with associated studies ahead of diagnosis.

3.3. Prefetching Workflow

Given the defined architecture together with the adaptions, one is able to search and retrieve prior medical reports and studies ahead of diagnosis. The process on how this is achieved can be described as follows:

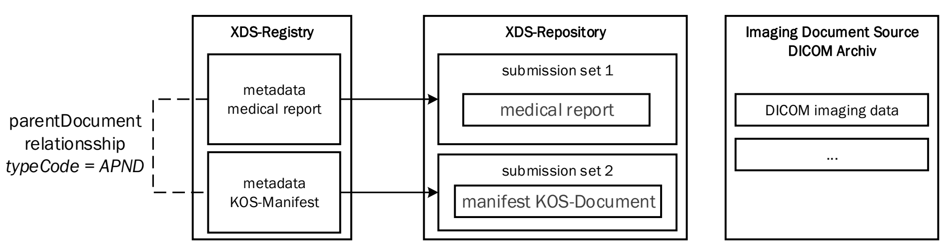

Figure 4. Concept of associating metadata of a medical report with the metadata of imaging data through the use of the parentDocument relationship.

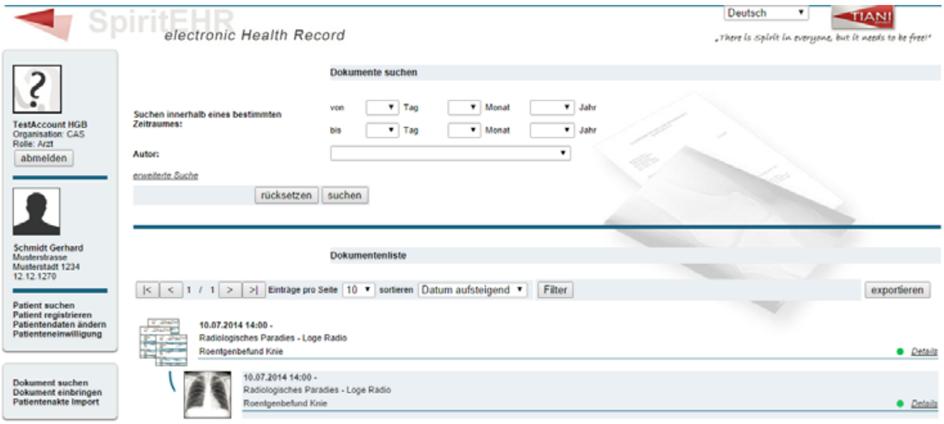

Figure 5. The test environment, showing a medical report and its relationship with associated imaging data.

During the process step *Patient Admission*, the referral is used as a basis for the search query using a *DocumentConsumer*. The *DocumentConsumer* retrieves any XDS-Metadata items complying with the desired search criteria. Whilst the result contains XDS-Metadata for medical reports, metadata entries for associated radiological images are to be retrieved by resolving the *parentDocument* relationship. This enables retrieval of the KOS-Manifest Documents using the ITI-43 transaction. After retrieving the KOS a client (RIS) is responsible for parsing the content of the KOS-document in order to determine the location of the desired images to be loaded. As the mechanism, on how to determine where to load the desired image data from, is out of scope of IHE, this paper proposes to use an additional external service that is capable of mapping a given id to a specific community in terms of XCA [7]. The external service also enables the retrieval of image data independent of different communities using the transaction RAD-69 [5]. It allows as well to migrate a given data storage to another location, as the mapping is done outside of community scope. In an ELGA context this external service could be the *GDA-Index* that already provides a national healthcare provider directory and could be extended to also provide respective storage locations (or a way to access them).

To test the presented concept a lab environment sufficient according to the requirements of IHE was set up, thus enabling the verification of the ability of a prefetching approach to be applied as part of a DHIS. This test environment contained all actors and supported all transactions for a single community as depicted in Figure 2 and was based on a TIANI Spirit EHR[1] infrastructure (see Figure 5). The consumer actors were based on Open Health Tools[2] implementations and an open source PACS (dcm4che[3]) was integrated to host the images.

Using the test environment, several tests were conducted, involving storage and retrieval of medical reports as well as associated imaging data. We defined use cases based on the observed ones, like mammography ‘prefetching of the first and the latest

---

[1] tiani-spirit.com/, last access: 2015-03-06
[2] projects.openhealthtools.org/sf/projects/iheprofiles/, last access: 2015-03-06
[3] dcm4che.org/, last access: 2015-03-06

reports and images of a patient for comparison'. We created, registered and uploaded the reports, created, uploaded and linked the metadata and images and then tried to load the reports (CDA, ELGA Diagnostic imaging study, EIS Full Support) and the associated DICOM objects.

We only tested the approach within a single community but from a consumer point of view neither the interfaces nor the functionality change when switching up from XDS-I to XCA-I. From the results of these tests it can be concluded, that the described approach works and using the *parentDocument* relationship enables an efficient search and retrieve mechanism of medical reports on the one the hand, and associated imaging data on the other. Figure 5 shows an example of a registered medical report associated with the imaging data based on the *parentDocument* relationship.

4. Discussion

The observation of the radiological workflow and the accompanying expert interviews showed that radiologists need previous medical findings but most of all the associated images for their diagnosis. To grant access to those images in a timely manner, IHE defined the XDS-I (cross-enterprise) and the XCA-I (cross-community) Integration Profiles. In this paper we propose the use of the *parentDocument* relationship mechanism to link the report and the images (i.e. the KOS-document referencing the images) on metadata level. This yields two major benefits. First, there is no 'hard link' to the image location within the KOS-document, so a migration of images from one storage to another is possible without updating documents. Second, since the images are not referenced directly in the medical report but in a separate, linked KOS-document, prefetching of the images is possible without loading and processing all available reports.

An important part of prefetching is the (automated) selection of the correct images given all available images. Due to constraints in time and bandwidth it is inadvisable to load all available images that are associated with a patient. To enable the selection of images in advance the Austrian PACS Procedure Code (APPC) was introduced and added to the XDS-Metadata for ELGA [9]. Thus certain parameters like modality, laterality or anatomy are included in that code and available after retrieving the Metadata. Further research is required to develop intelligent prefetching mechanisms that take all that information into account.

A major issue when separating report and image reference is the need to develop mechanisms to maintain the link in case of versioning, update or deletion of one of the documents or the respective metadata. This issue is not addressed within the IHE Integration Profiles and should be handled on an organizational level. A possible way to overcome this issue is the XDS Folder concept described in [7], but this also comes with increased effort to manage those logical containers (e.g. provide appropriate metadata).

Together with the experts of ELGA GmbH we assessed the concepts described and discussed the applicability for the Austrian electronic health record. The approach was regarded with favor and we decided to further develop it and test it together with major vendors relevant for the setup of ELGA in 2015.

References

[1] ELGA-G (2012): Elektronische Gesundheitsakte-Gesetz, BGBl. I Nr. 111/2012, 14.12.2012

[2] Wei, Chih-Ping, P. Jen-Hwa Hu, and Olivia R. Liu Sheng. "A knowledge-based system for patient image pre-fetching in heterogeneous database environments-modeling, design, and evaluation." Information Technology in Biomedicine, IEEE Transactions on 5.1 (2001): 33-45.
[3] Bui, Alex AT, et al. "Problem-oriented prefetching for an integrated clinical imaging workstation." Journal of the American Medical Informatics Association 8.3 (2001): 242-253.
[4] Sheng, Olivia R. Liu, et al. "Automated learning of patient image retrieval knowledge: neural networks versus inductive decision trees." Decision Support Systems 30.2 (2000): 105-124.
[5] IHE International, Inc., IHE Radiology (RAD), Technical Framework Vol.1 IHE RAD TF-1 Integration Profiles, http://www.ih4e.net/uploadedFiles/Documents/Radiology/IHE_RAD_TF_Vol1.pdf, last access: 21.01.2015.
[6] Eichinger, Lukas. "Änderungen im radiologischen Workflow durch die elektronische Gesundheitsakte ELGA" BSc thesis, University of Applied Sciences Upper Austria, Hagenberg, 2014
[7] IHE International, Inc., IHE IT Infrastructure Technical Framework, Volume 1 (ITI TF1): Integration Profiles, http://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_TF_Vol1.pdf, last access: 21.01.2015.
[8] Benson, Tim. Principles of health interoperability HL7 and SNOMED. Springer Science & Business Media, 2012.
[9] Fruehwald, F., et al. "APPC—A new standardised coding system for trans-organisational PACS retrieval." European radiology 20.9 (2010): 2153-2165.