STREAMING OF MEDICAL IMAGES USING JPEG 2000 INTERACTIVE PROTOCOL

Rita Noumeir, Jean-François Pambrun
Department of Electrical Engineering
University of Quebec, École de technologie supérieure
Montreal, Canada
Rita.noumeir@etsmtl.ca

Abstract— Access to all relevant information at the diagnostic decision moment improves the quality of care. With the deployment of the Electronic Health Record (EHR), information resides in different distributed systems. When conducting a diagnosis based on medical images for example, the physician needs to compare old images with current ones, while old images may reside in a different system: they need to import images for visualization which leads to a problem related to persistency management and information consistency. Since image streaming promises a solution for avoiding image import, we describe here how JPEG 2000 Interactive Protocol (JPIP) can be used to enable streaming of medical images directly from EHR connected image archives to visualization workstations. Moreover, we describe JPIP implementations in order to visualize a large image and present measurements of bandwidth efficiency improvements.

Keywords- Medical Imaging, Electronic Health Record, Image Communication, Image Streaming, JPIP, JPEG2000.

I. INTRODUCTION

The Electronic Health Record (EHR) enables informed health decision by making available all relevant prior diagnostic information; and this, independently from the geographic location of the point of access or the institution where the information was initially gathered. Prior diagnostic information is very varied and includes observations, laboratory results and images. The deployment of EHR is expected to improve the quality of care by enabling more informed decision; it is also expected to improve the efficacy and efficiency of the overall healthcare system by improving productivity and by reducing the duplication of information gathering.

EHR brings a big challenge because it is not a single system that can be provided by a single manufacturer. It is a virtual system that results from the cooperation of several heterogeneous distributed systems for providing ubiquitous access to the diverse diagnostic information related to a specific patient. Interoperability is therefore essential. Interoperability in healthcare has been very difficult to achieve; it is costly and frequently requires specific integration interfaces despite the existence of medical standards for many decades now. Even though, standards are necessary, alone they are not sufficient. They enable interoperability within a limited scope, for a specific clinical domain or a specific function. To close this gap, Integrating the Healthcare Enterprise (IHE) provides a process for building a detailed framework for the implementation of standards [1]. IHE started in 1998 and was sponsored jointly by the Radiological Society of North America (RSNA) and the Healthcare Information and Management Systems Society (HIMSS). Currently, several other associations sponsor IHE. IHE has expanded over several clinical domains and benefits from broad international support. IHE defined recently an architectural infrastructure for enabling documents sharing between multiple enterprises [2]. This is known as the Cross-Enterprise Document Sharing Integration Profile (XDS). XDS lays the basic framework for deploying regional and national EHR by addressing the needs for the registration, distribution and access across health enterprises of patient's documents. As medical images constitute important information of the patient health record, XDS has been extended to include images. As the result of an extensive investigation effort of several design solutions [3], the Cross Enterprise Document Sharing for imaging (XDS-I) is published as part of the IHE Technical Framework.

The deployment of XDS-I as the framework for sharing images within the EHR is taking place in many countries including Canada, USA, Japan and several European countries. But, several difficulties have emerged, such as the need to compare old images with current ones. In fact, to conduct the interpretation, the radiologist usually compares the current images with prior ones that may have been acquired in a different enterprise. With the EHR, the radiologist knows about the existence of those priors and can access them. However, comparison is conducted within a single software application that offers specific operations for medical imaging interpretation, such as a synchronized navigation between two different image sets. This application is thus required to have access to both image sets. Presently, most medical imaging applications assume images are under their complete control: all images are identified and managed in a single consistent way. This assumption does not hold when foreign images need to be imported into the system, as identification schemes are different between several enterprises and may result in identification that is not unique. Patient and order identifications are such examples.

Also, importing foreign images into a local application creates another major problem related to persistency management and information consistency. Image import is basically image duplication. How can foreign images be identified as such so they can be deleted or discarded at the end of the process? Moreover, how to propagate information correction to the duplicated instance?

One possible solution for all the previously stated problems consists in avoiding image import. This is achievable with image streaming. Image streaming can also provide tremendous gain in bandwidth when viewing large images or large image sets, by only streaming the data necessary to fulfill the user's task at the best screen resolution. This can be implemented with JPEG 2000 Interactive Protocol (JPIP). In this paper we describe how JPIP can be used in the context of EHR to enable streaming of medical images directly from imaging sources to image processing workstations. We also describe JPIP implementations in order to visualize a large image, such as a digital mammography image, and present measurements of bandwidth efficiency improvements.

II. ARCHITECTURE FOR INCLUDING IMAGES IN THE EHR

A. XDS Architecture

Within care delivery organizations multiple systems exist, each of which may produce, store or retrieve different clinical information. The XDS architecture enables patient's information, from separate care delivery systems, to be shared in the form of documents. A shared document is a very broad concept that represents a unit of health information being shared in a standard format.

The architectural model is based on a central registry that holds metadata describing every published document. It also responds to queries about documents meeting specific criteria. The registry does not store the document itself. However, it maintains information about the location from which documents may be retrieved. Therefore, the architecture is based on one or multiple distributed document repositories. A repository stores documents in a persistent manner and responds to document retrieval requests. Systems that produce information relevant to patient's continuity of care, such as radiology reporting systems, publish information as documents. Systems that are interested in accessing the patient's record query the registry for documents meeting certain criteria. Within the response to a query, the registry includes a reference to the document address, enabling the document consumer system to retrieve the document from its repository.

In order to share a set of images, a Digital Imaging and Communications in Medicine (DICOM) manifest that contains references to a set of DICOM instances is published. With this solution, the manifest is published and not the images. When a consumer retrieves the manifest, it needs to decode it to get the list of referenced images that are specified by a Universal Identifier (UID) along with the application title where to retrieve it. The consumer can then issue a DICOM transaction, such as retrieve (C-MOVE) or Web Access to DICOM Persistent Objects (WADO), to retrieve the images.

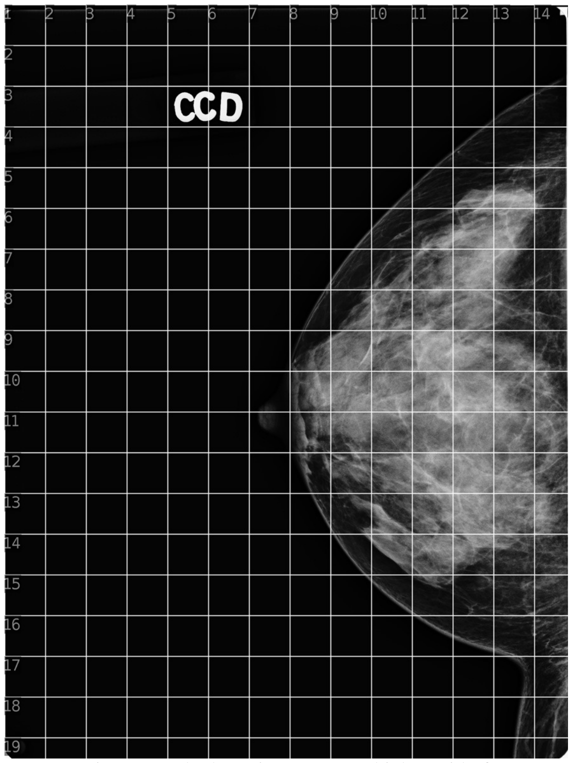

Figure 1. The large image over which a grid of 256 x 256 pixels is drawn.

B. Delivering images using JPIP

JPEG 2000 Interactive Protocol (JPIP) is a client/server standard image streaming protocol [5]. It allows a client application to request only portions of a JPEG 2000 image that are necessary to fulfill the client's viewing needs. JPIP streaming relieves the client application from importing the image into its environment eliminating thus the problems of persistency, consistency and reconciliation. It also results in an improvement in bandwidth efficiency when viewing images in a client/server environment. This improvement is very important in medical imaging as medical images are either large images or very large image sets. JPIP can be used as part of the XDS-I framework as follows [6]:

1. An imaging workstation (XDS-I consumer) queries the registry for a specific patient and for imaging priors that are relevant for the imaging case at hand. The registry responds with a list of documents, each representing a set of images that are available from an imaging archive.

2. The XDS-I consumer selects and retrieves a specific manifest from the Document repository.

3. For each referenced DICOM instance within the manifest, UIDs are extracted and used as values for a DICOM WADO query parameters. A WADO query is an HTTP request to a WADO server requesting a specific DICOM instance using specific query keys that are specified by the DICOM specifications. Query keys include image UIDs as well as transfer syntax. The

Consumer requests the transfer syntax to be the DICOM JPIP referenced transfer syntax.

4. The JPIP reference received by the Consumer application includes a Pixel Data Provider URL that specifies the address of the JPIP server capable of providing the pixels. The Consumer application is thus able to formulate a JPIP request to that server to allow interactive navigation.

III. RESULTS

JPIP is based on JPEG 2000 standard for compressing images. JPIP performance is tightly dependent on the way images are compressed with JPEG 2000. Moreover, JPEG 2000 compression parameters are numerous and their combination leads to a large number of possibilities. Consequently, to compress the image using JPEG 2000, understanding how the image will be requested using JPIP is essential. This requires considering how the user would manipulate the image which evidently depends on the medical image modality. We have implemented a client application that simulates the specific use case by issuing the adequate JPIP requests to a JPIP server capable of gathering information about data transfer. JPEG 2000 compression and JPIP interaction capabilities have been provided by commercial libraries from Aware Inc.

A large mammography image is used (Fig. 1). It has a width of 3540 pixels and a height of 4740 pixels; its size is 33,562,298 bytes. The image is compressed with 5 decomposition layers. Precincts are used to achieve full resolution regions of interest. The precincts size of subbands HL2, LH2, and HH2 is considered equal to 128 x 128 pixels. The size of all other precincts is considered equal to 256 x 256 pixels. To allow progressive download, the image was compressed with 10 quality layers. The image is supposed to be visualized on a screen whose width is 1920 pixels and whose height is 1080 pixels. This is the screen size of a common computer. Evidently, this size is different from the common radiology dedicated workstations screen sizes that are in use nowadays. However, screen sizes and images sizes are continuously increasing. But, the discussion here will always be valid as far as the screen size is smaller than the image size.

Clearly, the screen size is smaller than the image size; therefore information from low resolution subbands up to LL2 is enough. Quality layers are requested to be downloaded progressively: the lowest quality layer followed by a better quality layer, until all quality layers are requested. This enables a low quality initial image to be displayed very quickly, while subsequently refined until best (screen) resolution is attained. Images reconstructed with different quality layers are shown in Fig. 2. Visual quality in low-frequency regions improves with quality layers.

Peak signal to noise ratio (PSNR) is calculated for each reconstructed image. Table 1 shows the additional bytes required to transfer each quality layer.

$$RMSE\left(I,\hat{I}\right)=\sqrt{\sum_{x}\sum_{y}\left(\left(I_{(x,y)}-\hat{I}_{(x,y)}\right)^{2}\right)} \quad (1)$$

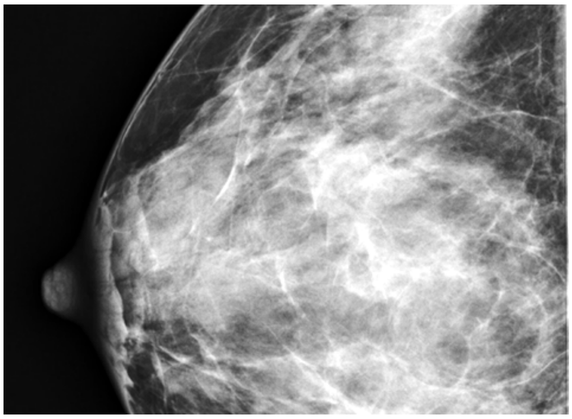

1 of 10 quality layers

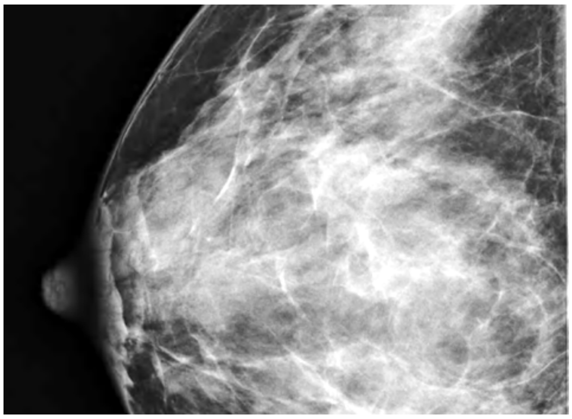

4 of 10 quality layers

Figure 2. Generated images using different quality layers.

$$PSNR\left(I,\hat{I}\right)=20\log_{10}\left(\frac{I_{\max}-I_{\min}}{RMSE\left(I,\hat{I}\right)}\right) \quad (2)$$

An image at the lowest quality layer requires 57,848 bytes only, compared to the full resolution of the image of 6,966,349. Each additional quality layer improves the quality of the image and requires additional bytes to be transferred; the total is the amount of bytes needed to display the image at the best resolution of the screen. Compared to the full resolution of the image, a compression ratio over 15:1 is achieved.

Since the best resolution of the screen is less than the full resolution of the image, JPIP requests have been generated to simulate a lens tool that is used to visit the

TABLE 1. BYTES TRANSFERRED AND ERROR MEASUREMENTS FOR 10 QUALITY LAYERS

Quality layer	PSNR(dB)	Bytes downloaded
1	44.67	57,848
2	51.81	64,580
3	57.95	63,149
4	62.01	79,832
5	63.04	31,890
6	66.73	108,365
8	68.81	34,392
8	103.83	33,811
9	Inf	249
	Total:	474,353

image completely, according to a navigation scheme that goes top down, from left to right. The regions of interest are shown in Fig. 1 as grid lines superimposed on the image. The region of interest is considered of size 256 x 256 pixels. The additional bytes needed to display full resolution regions of interest are shown in table 2. The total amount of bytes to view the complete image at full resolution is 7,014,127. This is achieved after visiting all regions of interest. It is slightly bigger than the initial image size. Of importance is the additional amount of bytes required to visualize a region of interest which is about 56 kilobytes. Moreover, one can note that many regions do not contain information of diagnostic value. These regions correspond to the background and occupy in the case of this mammography image about 60% of the whole image. These regions are not examined at full resolution and may end up not being requested at all.

IV. CONCLUSION

JPIP brings two major advantages when viewing medical images in a distributed environment, such the one encountered with EHR. The first advantage comes from the streaming capability which eliminates the need for importing foreign images into a medical image archive, avoiding thus the problems related to information consistency and persistency management of duplicated images. The second advantage comes from the significant improvement in bandwidth efficiency when viewing medical images which usually are either large images or very large image sets.

JPIP enables the client application to visualize an image very quickly with a low quality while enabling progressive refinement at a subsequent moment. It also enables the client application to visualize a large image at the best screen resolution with much less data than required when visualizing the same image lossless compressed. This additional "compression" depends on the ratio of image size to screen size. Moreover, JPIP enables the display of full resolution regions by requiring additional data whose amount is directly influenced by precincts size. In this paper, we have proposed an approach to implement JPIP in order to visualize a large image. We have also measured and presented bandwidth efficiency improvements.

While using JPIP to deliver medical images from the EHR to the radiologist's workstation appears very promising, many challenges still exist: image display applications need to integrate a JPIP client; image archive systems need to integrate a JPIP server; adoption is needed and interoperability testing is required.

ACKNOWLEDGMENT

This work was supported by a grant from Canada Health Infoway and the Government of Quebec. Thanks to Aware Inc. for providing the JPEG2000 and JPIP libraries.

REFERENCES

[1] Noumeir R., Integrating the Healthcare Enterprise Process, Int. J. Healthcare Technology & Management, v 9, n 2, 167-180, 2008

[2] IHE Technical Framework and supplements, http://www.ihe.net/Technical_Framework/index.cfm

[3] Noumeir R., IHE cross-enterprise document sharing for imaging: design challenges, Progress in Biomedical Optics and Imaging, Proc. SPIE, v 6145, Medical Imaging, p 61450Q, 2006.

[4] The Digital Imaging and Communications in Medicine (DICOM) standard, http://medical.nema.org

[5] Lima, L., Taubman, D., Leonardi, R., JPIP proxy server for remote browsing of JPEG2000 images, IEEE 10th Workshop on Multimedia Signal Processing (MMSP), p 844-849, 2008.

[6] Noumeir R., Pambrun J.F., Images within the Electronic Health Record, IEEE International Conference on Image Processing, 2009.

TABLE 2. ADDITIONAL TRANSFER SIZE FOR VIEWING REGIONS OF INTEREST

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
1	15,743	16,621	16,388	16,339	16,225	16,223	15,929	16,047	16,299	16,134	16,139	15,866	15,827	13,067
2	13,915	13,681	14,398	14,410	13,046	13,895	14,152	13,896	14,045	13,987	14,135	14,252	13,784	12,958
3	19,893	20,258	20,219	17,257	24,855	24,248	13,149	13,811	13,914	14,190	14,338	14,271	27,716	25,823
4	17,730	17,224	16,467	15,183	12,970	12,740	13,867	14,133	14,084	13,820	16,186	33,710	38,838	39,208
5	14,097	13,990	13,781	13,956	14,036	13,851	13,786	14,025	13,785	17,398	38,556	47,094	53,788	44,056
6	14,362	13,906	14,022	13,976	13,892	14,053	13,694	13,502	16,411	42,433	51,665	56,394	55,994	42,603
7	14,661	13,927	14,136	14,135	14,396	13,945	14,021	13,669	39,135	52,453	56,641	56,309	53,531	42,456
8	14,791	13,981	14,428	14,362	14,197	14,184	13,681	25,420	48,429	56,645	56,788	56,927	55,422	44,003
9	14,712	14,498	14,471	14,610	14,331	14,625	13,033	44,868	56,467	56,932	55,665	56,594	56,006	44,941
10	15,109	14,356	14,253	14,384	14,385	14,363	18,038	54,604	56,705	56,680	56,453	56,079	56,634	46,267
11	15,042	14,616	14,406	14,463	14,270	14,159	20,213	52,096	56,288	56,846	56,854	56,645	56,926	47,060
12	15,382	15,004	14,634	14,755	14,667	14,520	13,461	36,809	54,708	56,681	56,411	56,500	57,016	47,159
13	15,038	14,635	14,533	14,751	14,467	14,531	14,149	20,519	47,458	55,660	56,957	56,857	57,059	47,016
14	15,313	14,648	14,763	14,889	14,665	14,396	14,730	13,805	32,158	50,093	56,893	55,270	54,873	45,876
15	15,059	14,740	14,700	14,812	14,459	14,586	14,334	14,686	14,156	32,798	43,561	45,570	51,052	44,950
16	15,102	14,783	14,717	14,736	14,831	14,440	14,623	14,579	14,557	13,775	22,658	36,131	44,561	42,612
17	15,187	14,862	14,659	14,812	14,818	14,912	14,952	14,843	14,720	14,738	14,468	13,649	23,794	37,917
18	14,792	14,602	14,887	14,729	14,967	14,811	14,723	14,938	14,667	14,694	14,408	14,726	26,358	40,032
19	8,609	7,642	7,563	7,612	7,687	7,616	7,534	7,606	7,605	7,626	7,775	7,506	14,947	18,275