

# A Medical Imaging Informatics Appliance

James W. Cooper, Ph.D.<sup>1</sup>, Shahram Ebadollahi, Ph.D.<sup>1</sup>, Ellen Eide, Ph.D.<sup>1</sup>,  
Chalapathy Neti, Ph.D.<sup>1</sup> and Giridharan Iyengar<sup>2</sup>

<sup>1</sup>IBM T J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598

<sup>2</sup>Signal Match, Inc, 110 Washington Ave, Suite 201 Pleasantville, NY 10570

Abstract

*Objective* Improving quality and efficiency of care, reducing medical errors, and alleviating the workload of the radiologists requires efficient integration of multi-modal imaging informatics tools and techniques with existing medical records management systems.

*Methods* The Medical Imaging Informatics Appliance is designed to enable high-performance, collection-level image analytics. We demonstrate its usability through use-case scenarios.

*Results* The resulting appliance is a system which leverages the Cell broadband engine as its high-performance computing platform for carrying out image analytics. Algorithms such as image registration, are deployed on the Cell platform. Multi-modal medical record content analysis is applied to patient records to make them amenable for collection-level search and retrieval based on semantic concepts derived from RadLex taxonomy. We demonstrate the usability of the system with use-case scenarios.

*Conclusions:* The Medical Imaging Informatics Appliance provides efficient use of imaging informatics in radiologists' daily work.

*Keywords:* PACS, Imaging Informatics, High-performance Image Analytics, Decision Support

## Introduction and Motivation

Diagnostic imaging contributes substantially to the rapid proliferation of data in health-care centers. The increase has been attributed to two main factors: (1) the aging population (53% increase in the population 65 and older by 2020) and concomitant increase in disease incidence and the required care, and (2) the recent improvements in the technology of imaging which result in higher quality and therefore more voluminous data sets (A single imaging center can produce terabytes of images each year). The proliferation of diagnostic images puts extra pressure on radiologists and affects their efficiency, and reliability in reading images and handling cases. New methodologies and systems are required to assist the radiologists in coping with this trend and improving their performance.

PACS<sup>1</sup> systems were conceived in the 1980s to help radiologists store, communicate, and view diagnostic images, while systems such as RIS<sup>2</sup> and HIS<sup>3</sup> were designed to aid the hospital staff manage text records and workflow. Although the adoption of electronic health records is on the rise, those systems in their current form are not geared to provide a solution for the future needs of radiologists. There is a disconnection between components of the patient record scattered in different system. Although the electronic patient record aims at making the various systems interoperable using standards such as HL7 and frameworks such as IHE, the lack of a unified multi-modal medical record is being felt. PACS and other sub-components of the electronic health record need to be either amended or replaced with novel systems that address those needs. Today, radiologists can benefit tremendously from medical imaging informatics [1]. These techniques provide tools for managing diagnostic imaging data throughout the entire chain of content generation, storage, management, and consumption. In addition, temporal and population-based analysis of the medical diagnostic images are also highly desirable. Imaging informatics tools and techniques are being used in post-analysis or research-oriented studies and effective use of these techniques requires innovations in the following areas:

- Validation of various imaging informatics tools in the context of radiology. This allows the user (radiologist or clinician) to be able to apply such tools to data for extracting insight and decision-enabling information and have certain level of confidence about their performance and reliability. This is the subject of several studies being conducted now [2] .
- Creation of real-time applications of the informatics tools applied to medical diagnostic images. Most of the analytic routines are extremely time-consuming, which precludes them from being used in the daily routines of the radiologists. For example, the process of image registration cannot be applied on demand to images under study, because it may take up to several hours for the results to become available.

---

<sup>1</sup> PACS: Picture Archiving and Communication System

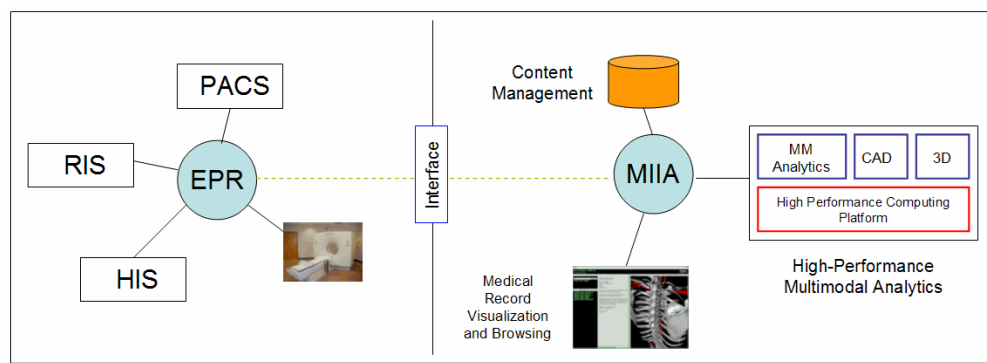
<sup>2</sup> RIS: Radiology Information System

<sup>3</sup> HIS: Hospital Information System

- Advanced management of multi-modal medical record content is required to combine the various pieces of multi-modal data and make them available for population and temporal studies, decision support and treatment planning.

To address the latter two requirements we describe the Medical Imaging Informatics Appliance (MIIA), a system that enables real-time, collection-level analytics for improved diagnosis and patient care.

The MIIA is a novel system for integrating imaging informatics with current PACS/RIS systems in use in the radiology department. The figure below shows an electronic health record system enhanced by the MIIA.



**Figure 1. Enhancing the electronic medical record system with medical imaging informatics appliance. The appliance serves as an intermediary for image-enabled medical records with high-performance analytics and multi-modal content management capabilities and thin client support for visualizing and browsing the image-based records.**

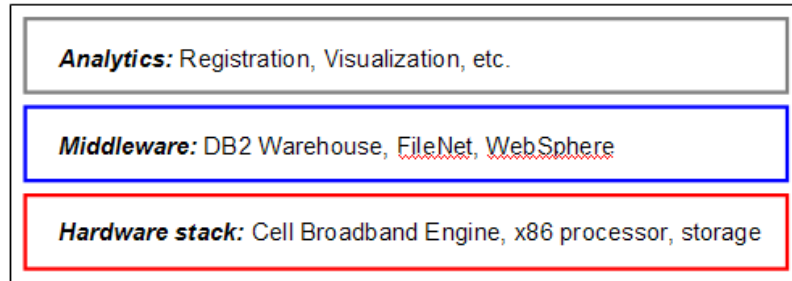
In the following, we describe the architecture of the MIIA and the details of its components. Later, we provide an example of using the MIIA in the context of radiology workflow.

## **An Appliance for Enabling Imaging Informatics**

One way to view the appliance is according to its levels of abstraction, as shown in Figure 1. The tight coupling of hardware together with enterprise grade middleware comprise a platform which is well suited for showcasing high-performance analytics.

The MIIA as shown in Figure 1 consists of four main components, high-performance multi-modal image analytics, multi-media content management, appliance management, and a visualization and interaction user interface. Through an interface the MIIA is connected to the EPR (Electronic Patient Record) or the

sub-components of it. In the following we elaborate on the components of the MIIA.



**Figure 2. Diagram of the MIIA according to levels of abstraction.**

### **Multi-modal Image Analytics**

Diagnostic images are acquired using different imaging modalities, such as X-ray, CT-scan, MRI, etc., where each modality is suited for illustrating and capturing a different aspect of the patient's anatomy. Radiologists use variety of tools to perform information extraction, or reconstruction of multi-modal images to read them better and to make more informed decisions. A variety of tools and software are used to perform these tasks. Most of the image analytics routines are compute-intensive and are not suitable for use during image reading and diagnosis. For example, registering a brain MR image to a baseline for the purpose of assessing the progression or regression of a tumor, is a very time consuming process. The radiologist has to make a request for this computation to be carried out and wait for the results so he can assess the state of the patient's health.

In the MIIA, we provide these analytics on a high-performance computing platform, so that they can be applied to images at near real-time speed. The platform we employ for the high-performance image analytics is the Cell Broadband Engine processor [3].

#### *Cell Broadband Engine – A Vehicle for Compute-Intensive Analytics*

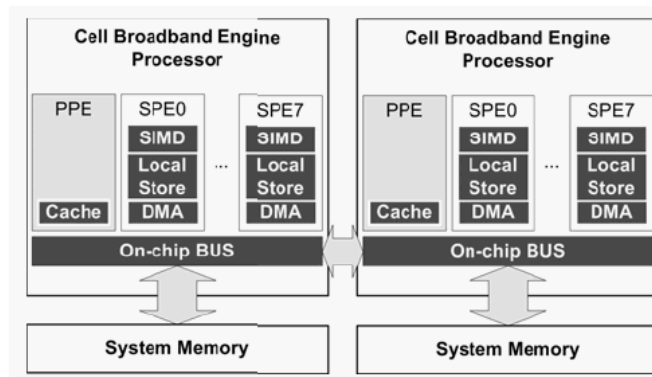
The Cell Broadband Engine (CBE), is an asymmetric multi-core processor with high peak performance: it combines eight synergistic processing elements (SPEs) and a Power Processing Element (PPE), which is a general-purpose IBM PowerPC processor. Two of these PPEs can be connected via a high speed bus as shown in Figure 3, so 16 SPEs can run in parallel. Each SPE, furthermore, has a

SIMD unit, which can perform a floating or integer operation on four data elements at every clock cycle.

To accelerate an image analytic routine on this processor, you must optimize the program to exploit the parallelism at both task and SIMD-instruction levels and use the memory bandwidth efficiently. First, at a task level, you must partition the program into multiple tasks that fit in the local store on each SPE. Unlike conventional microprocessors, each SPE does not have hardware cache memory to manage the small on-chip local store. Thus, one can view this architecture as a distributed memory multiprocessor with a very small local memory attached to a large shared memory.

At a SIMD-instruction level, moreover, you typically need to restructure frequently executed sections of the task by using intrinsics, which allow the programmer to write SIMD instructions explicitly in the code. Such restructuring is called SIMD-ization.

Finally, for utilizing the memory bandwidth efficiently, you must partition the data in such a way that each task can transfer it in a large block, such as one or multiple cache lines (128B per cache line), using a DMA engine. Pixel-wise memory access, for example, results in a very poor usage of the memory bandwidth.



**Figure 3. The high-level structure of a dual CBE system. Each synergistic processing element (SPE) has a SIMD engine, a high-speed local store, and a DMA engine.**

Certain image analytics algorithms, such as registration and 3D reconstruction are well-suited to CBE-based implementation [5,6]. Our goal is to provide a suite of such tools implemented for the CBE to power the multi-modal analytics engine of the MIIA.

## *Image Analytics – Image Registration*

The near real-time imaging analytics that the MIIA currently offers is image registration. While experienced radiologists frequently make comparisons between images taken at two different time points for a given patient to assess changes, it has been reported that it is more accurate to look at the registered (aligned) images [4].

The specific registration algorithm we have implemented on CBE is based on Matte's mutual-information-based multi-resolution linear registration algorithm [13], as implemented in National Library of Medicine's Insight Imaging Toolkit (ITK) [7] for the base implementation.

The algorithm, takes two images from a given patient, one of which is identified as "fixed" and the other as "moving," and creates a new image corresponding to the 3D affine transformation of the moving image to best align it to the fixed image.

The image registration algorithm [5,6], chooses the affine transformation, which maximizes the mutual information between sub-sampled versions of the fixed image and the transformed moving image as the optimal transformation.

Computationally expensive portions of the algorithm such as computing the mutual information between the fixed and transformed image and the gradient of the mutual information for a given transformation have been recoded for processing on the CBE as described in the previous section.

The CBE-based implementation of the registration algorithm was shown to register a pair of 256x256x30 3D MR images in one second [5], with better time improvements achieved compared to the implementation of the same algorithm on a regular high-end PC for larger volumes [6].

## **Multi-media Content Management and Retrieval**

This module of the MIIA allows for content-based retrieval of multi-media radiology records, where each record consists of a group of images and an accompanying text that provides the findings and summary created by the radiologist.

These text reports are analyzed using MedTAS [8], a powerful text analytics system that can be configured to recognize diagnoses, sites, drug dosages and Radlex terms [9]. All of the annotators run on a public domain middleware system

called Apache UIMA [10] and are derived from a sophisticated dictionary-based annotator that recognizes any single or multi-word concept. It not only annotates the text span, but includes the canonical form of that term. For example, if the trade name “Lipitor” is found, the created annotation will include a field for the generic form “atorvastatin.”

The annotations are stored in a database along with their offsets and canonical forms, so these offsets can then be used to construct highlighted documents showing the types and locations of terms that were found. Mouse-over highlighting allows the user to view the canonical form of any highlighted term. In addition, each term from the Findings branch of the Radlex ontology is stored with information about that term’s position in the Radlex tree, so that parent and child information can be computed.

The collection of multi-media radiology records can be searched for a given concept using a simple database query. The system could be also expanded to find radiology notes similar to the one that the user is currently viewing. Upon identifying those similar documents, their associated images are then displayed to the user to provide a visual reference to compare different cases.

## **User Interface**

The MIIA supports thin clients for displaying and interacting with the multi-media medical records. All image files (DICOM or ANALYZE format) are converted to display as a series of slices in a standard web browser. The user can scroll through the various images using a slider constructed using the Dojo toolkit [11]. Drug dosage displays are added to the image display pages using Ajax callbacks [12], and the Radlex ontology tree is displayed using the Dojo Tree widget.

## **Results**

We demonstrate the usability of the MIIA system through a use-case scenario, where the user (radiologist or clinician) recalls a diagnostic image of a particular patient, registers the image with a baseline image obtained from the same patient at an earlier time, search the multi-media records database of patients for a given diagnosis and recalls their images to assess their similarities with the current image to better decide on the right diagnosis and reading of the current image.

### *Description of a Use-Case Scenario*

In a typical scenario, the user selects a specific patient to study, and then sees a screen where two images from that patient's history are shown [Figure 4]. Users at this point can scroll through the different slices in the volumetric data for each image. The buttons above the images allow the clinician to:

- register the images,
- select the patient's history,
- search for other patients by diagnosis using the Radlex tree

If the user selects the "History" button, a series of images is displayed representing the patient's entire imaging history [Figure 5]. Selecting a checkbox under one of the images and the selecting the "Meds" button displays the patient's current medications as discovered by text mining of the clinical reports. Selecting an image and the "Report" button causes the display of the report, with diagnoses, site, Radlex terms and drugs and dosages highlighted. In one such report, the drug dosage phrase "Zetia, 10 mg p.o. daily" is highlighted, and when you pass a mouse over this string, the generic name "Ezitimibe" is also shown.

The reports are constructed using style-sheets such that hovering over a drug with the mouse causes the display of the generic medication name, and hovering over a diagnosis brings of the formal ICDO diagnosis name.

Since the text analysis system recognizes any term in the Radlex ontology, a user can search among all indexed reports for any report containing a Radlex term or one of its children. For example, should the user begin to enter the phrase "adeno" the system automatically provides completion prompts from which the user can select "adenocarcinoma," and then view the multi-media records that show instances of that condition.

The user can now search for any clinical report containing this Radlex term or any child term under the selected term. In this case, searching for "adenocarcinoma" or its children returns reports containing such terms as "adenoid cystic carcinoma" as well.






Figure 4. A patient display in MIIA prototype interface.

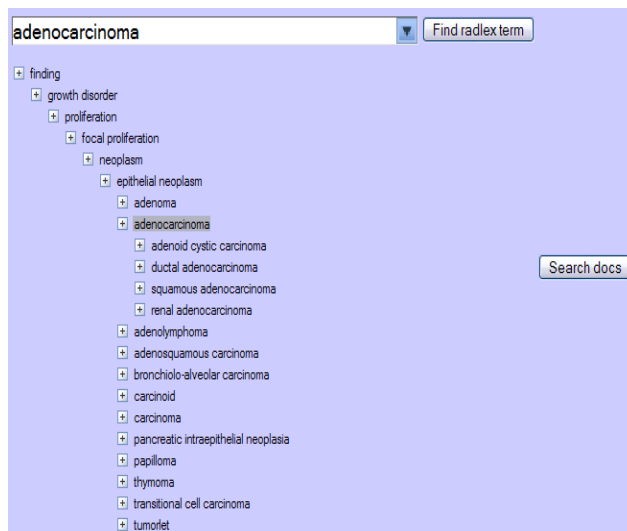


Figure 5. tree expanded only around leaf above "adenocarcinoma."

## Conclusion and Future Work

We present the concept and an early prototype of the Medical Imaging Informatics Appliance, which is designed to serve as a vehicle for enabling high-performance image analytics and multi-media medical records management. Our work on the concept of the appliance has been motivated primarily by the need for integrating high-performance image analytics with the radiology workflow. Another strong motivation was the need for system and tools for managing multi-modal medical record content and combining the various pieces of multi-modal data to make them available for population and temporal studies, decision support and treatment planning.

We plan to extend the range of high-performance image analytics offered by the appliance in the near future, starting with non-linear registration, 3D reconstruction and volume visualization.

In the future, we will add process management functionality to the appliance. This will allow the medical records, processes, analytics and visualization and browsing to be managed more effectively by specialized components.

## References

1. Andriole, Katherine P., "An Introduction to Radiology Informatics", in Business Briefings: Future Directions of Imaging, 2006
2. Cancer Biomedical Information Grid (caBIG), <https://cabig.nci.nih.gov/tools/NCIA>
3. J.A. Kahle, M.N. Day, H.P. Hofstee, C.R. Johns, T.R. Maeurer, and D. Shippy, "Introduction to the Cell Multiprocessor", IBM Journal of Research & Development, vol 49, no. 4/5, 2005, pp. 589-604
4. Erickson, B.J. et al. "Effect of Automated Image Registration on Radiologist Interpretation." Journal of Digital Imaging. June 2007. Vol 20 Number 2. pp 105-113.
5. Ohara, Moriyoshi; Yeo, Hangu; Savino, Frank; Iyengar, Giridharan; Gong, Leiguang; Inoue, Hiroshi; Komatsu, Hideaki; Sheinin, Vadim; Daijavad, Shahrokh; Erickson, Bradley, 2007 IEEE International Symposium on Biomedical Imaging, Arlington, VA.
6. Ohara, Moriyoshi; Yeo, Hangu; Savino, Frank; Iyengar, Giridharan; Gong, Leiguang; Inoue, Hiroshi; Komatsu, Hideaki; Sheinin, Vadim; Daijavad, Shahrokh, 2007 IEEE International Conference on Multimedia and Expo, July 2007, 272 – 275
7. Insight Segmentation and Registration Toolkit (ITK), <http://www.itk.org/index.htm>
8. Mack, Robert L. *et. al.* "BioTeks: Text Analytics for the Life Sciences." *IBM Systems Journal*, March, 2004.
9. see <http://Radlex.org>
10. Ferrucci D., and Lally, A. "The UIMA System Architecture," *IBM Systems Journal*, March, 2004.
11. see <http://dojotoolkit.org>
12. Hadlock, Kris, *Ajax for Web Application Developers*, Sams Publishing, New York:2007.
13. Mattes, D. Haynor, D.R., Vessele, H., Lewellen, T.K. and W. Eubank, "PET-CT Image Registration in the Chest Using Free-form Deformations," *IEEE Transactions on Medical Imaging*, vol 22 no. 1 (Jan 2003) 120-128.