Efficient Exploitation of Parallel Computing on the Server-Side of Health Organizations' Intranet for Distributing Medical Images to Smart Devices

Athanasios Kakarountas and Ilias Mavridis

Computer Science and Biomedical Informatics, University of Central Greece,
Papasiopoulou 2-4, 35100 Lamia, Greece
{kakarountas,emavridis}@ieee.org

Abstract. The distribution of high-resolution medical images to smart devices, in a health organization premises, is considered in this work. The aim is to reduce network traffic and computation load (of the smart devices). Security issues are also considered. The approach takes full advantage of the parallel processing on medical images performed on the server side, exploiting GPGPUs processing power. Data are pre-calculated and modified to best fit the targeted smart device's display. The evaluation results show minimization on processing requirements at the smart device, while network traffic is reduced significantly when few actions are performed on the image.

Keywords: Parallel computing, medical data processing, medical image, safe access.

1 Introduction

The distribution of medical content from a patient's Electronic Health Record (EHR) has been under major discussion due to privacy and life critical issues. The most common approach until now was aiming to share the original medical content stored in a datacenter to registered clients, to which someone may download data and perform any kind of processing at the client-side [1], [2]. This approach is mainly targeting workstations (connected to medical equipment), desktop PCs or bulky mobile devices (e.g. laptops). However, this introduces significant increase to the network's traffic load; it offers a variety of security vulnerabilities for malicious attacks, while it doesn't take into consideration future trends for health care including the mobile smart devices revolution [3], which may be a critical tool for paramedics or professionals offering health services away from the health organization. Furthermore, this approach doesn't take into consideration the computing power that is required for proper processing and display of medical data (i.e. a medical image of several tens of MB). Finally, considering that most of the available mobile medical solutions are limited in devices with poor displaying capabilities, it is questionable if this approach actually contributes to the needs of a wireless monitoring application.

K.S. Nikita et al. (Eds.): MobiHealth 2011, LNICST 83, pp. 394–398, 2012.

This work aims in getting together several technologies, from various scientific fields, and suggest a new approach for distributing medical images on mobile smart devices. With the term smart device, we refer to a device offering wireless communication, embedding a descent processor and memory capacity, a reasonable display resolution, and finally offering an easy to use User Interface (UI) through an extensible and robust Operating System (OS). The scientific fields that get under the umbrella of the proposed approach are: Parallel Computing, Distributed Computing, Security, Image Processing, Wireless Communication and Internet Technologies.

2 Exploiting Parallel Computing on the Server-Side

The proposed approach was derived from a simple concept. The majority of the smart mobile devices lack computing power and display capabilities. Thus, there is no need to distribute medical information as is, under the constraint of no resilience in information loss, since the available hardware is inherently imposing this loss in order to display it. Issues concerning latency for downloading data and the security of the storage media (client) are also critical, introducing extra threats and vulnerabilities.

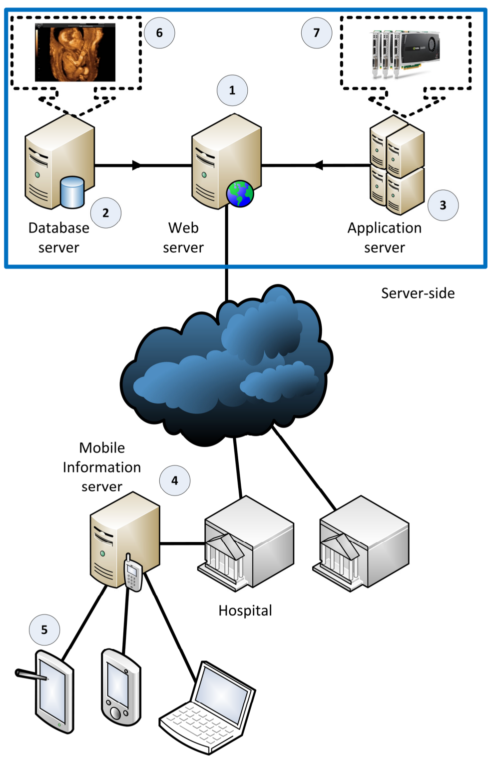

Fig. 1. The proposed approach to distribute medical image content

Thus, in terms of security, it may be said that the distribution of critical parts of data, pre-processed on the server and modified accordingly to fit the characteristics of the targeted smart device, may protect the EHR's data and expose in threat only part of it (which is a result of processing and not the originally stored data). In Fig. 1 a graphical representation of the approach is presented.

The latter mentioned framework consists of four servers, one for distributing the web content through an https connection (circled 1), a second one for serving queries and responses to and from the EHR database (circled 2), a third one for processing data in near-real-time (circled 3) and a fourth one (circled 4) for serving the modified content to the smart devices (circled 5). The first three servers are located to the server-side, while the fourth server corresponds abstractly to the hospital's telecommunication equipment. Although, this configuration has been identified, an alternative configuration may apply for various reasons (i.e. one server for all services to reduce cost).

The main contribution of this work is the migration of heavy processing load to the server side instead of the client side. This allows a wider range of smart devices to be used on medical image distribution (and other similar applications), but also moves complicated calculations, codecs and technology issues (i.e. compatibility, memory capacity etc.) from the client to the server.

2.1 Display of the Information

The smart device is registered during the establishment of a TCP connection, and from a Devices Database, the pre-characterized display resolution is set as the target of the application. The user is authenticated and then a dynamic webpage is created from which the user gets access to the medical image. Data remain in the database server and only a manipulated projection (suitable for the targeted resolution) is offered to the user. This approach makes possible the connection to the system of older smart devices with low processing characteristics.

2.2 Data Processing

The available data (circled 6 in Fig. 1) in the EHR is fetched from the database server to the processing servers. In the proposed framework, those servers are capable for parallel processing in General Purpose Graphic Processing Units (GPGPUs), since medical information consists, apart from text, of multimedia content such as image, video, sound in their appropriate container formats. The solution of GPGPUs allow multiple processing using several kernels, allowing thus pre-processing of information and appropriate modification to fit the targeted display. The GPGPU architecture that is considered in the proposed Framework is the CUDA NVidia [4],[5]. An example of pre-processing is illustrated in Fig. 2, where the possible actions (move around, zoom in-out and lighten-darken) of the user are pre-calculated and are available before the action takes place.

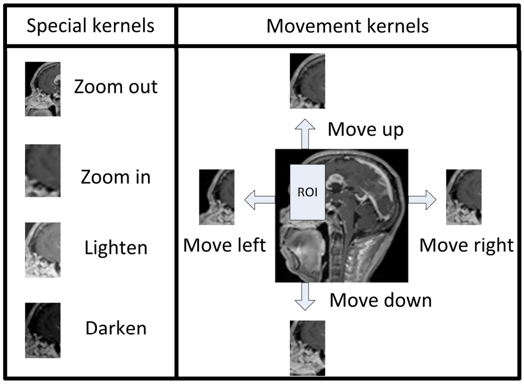

Fig. 2. An example of kernels executed on GPGPUs allowing pre-calculation of actions and modification for the targeted display

2.3 Security Issues

Security is always a hot topic of discussion and debate in such applications. The nature of the proposed framework guarantees that the actual data (except text) never gets out of the server boundaries. In contrast, a projection of data is offered to the user, to allow best viewing of the Region of Interest (RoI), but never acquire the data in whole. Another security mechanism is the authentication of the user and the time to live of the offered information [6].

During the authentication phase, the user's and the patient's authentication information are requested and a temporal communication key is generated from a hash function [7]. Also the MAC address of the device is considered and checked if it is registered in the hospital's registered devices. The derived message digest is used to create a temporal communication session. If the session remains inactive for a pre-defined time period, it ends and further access to the data requires the establishment of a new session. The communication in application level is exploiting the HTTPS protocol.

3 Evaluation of the Proposed Approach

Although the proposed approach is considered for a Web-based application, it may be installed as a wireless service at a governmental building, i.e. hospital. Such a scenario was considered, implementing a case of continuous access requests to a large collection of DICOM files [8]. An Apache server (on a Linux Ubuntu 10.10 computer with 2 GB RAM), embedding mySQL and PHP was responsible for distributing content through the web and a handheld mobile phone (Samsung Nexus S embedding Android 2.3) was used as the client device. This scenario was considered to evaluate under hard conditions the traffic load, the data processing time and the display latency. The network was based on WiFi 802.11g and a wireless 54 Mbps dedicated connection was the communication medium. As it is observed in Table 1, the performance (processing time) achieved with the proposed Framework is significantly increased compared to a typical implementation.

Table 1. Evaluation of the proposed Framework, by accessing DICOM images in real-time over an 802.11g wireless connection

Approach	Network Traffic (Mbps)	Processing Time (s in average)	Time to display after action (s)	Initial download latency (s in average)
Typical	50,0	2,3	0,3	11,9
Proposed	1,0	0,4	0,2	0,3

4 Conclusions

The proposed approach takes full advantage of the parallel processing to data performed on the server side, exploiting the GPGPUs processing power. Data are pre-calculated and modified to best fit to the targeted display. The network traffic is decreased, and in cases of in-hospital installations, the performance is significantly increased. Critical issues of security are also considered and addressed, offering a good degree of security in communication and a high degree of security for the EHR. Finally, it is the first time in the technical literature that a parallel processing architecture is becoming part of a EHR system.

References

1. Zhang, X.M., Zhang, N.: An Open, Secure and Flexible Platform Based on Internet of Things and Cloud Computing for Ambient Aiding Living and Telemedicine. In: Int. Conf. Computer and Management 2011, pp. 1–4 (2011)
2. Kannoju, P.K., Sridhar, K.V., Prasad, K.S.R.: A New Paradigm of Electronic Health Record for Efficient Implementation of Health Care Delivery. In: Second International Conference on Intelligent Systems, Modelling and Simulation 2011, pp. 118–120 (2011)
3. Smartphoning it in. Harvard Medical School, Harvard Health Letter (November 2010)
4. Franco, J., Bernab, G., Fernndez, J., Acacio, M.E.: A parallel implementation of the 2d wavelet transform using cuda. In: Euromicro Conference on Parallel, Distributed, and Network-Based Processing (2009)
5. Matela, J., Rusňák, V., Holub, P.: Efficient JPEG2000 EBCOT Context Modeling for Massively Parallel Architectures. In: Data Compression Conference 2011, March 29-31, pp. 423–432 (2011)
6. Michail, H.E., Kakarountas, A.P., Goutis, C.E.: Server Side Hashing Core Exceeding 3 Gbps of Throughput. Int. J. Net. Sec. 1, 43–53 (2007)
7. Michail, H.E., Kakarountas, A.P., Milidonis, A.S., Goutis, C.E.: A Top-Down Design Methodology for Implementing Ultra High-Speed Hashing Cores. IEEE Trans. Dep. Sec. Comp. 6, 255–268 (2009)
8. OsiriX sample DICOM images, http://pubimage.hcuge.ch:8080/