Real Time Adaptive Filtering for Digital X-Ray Applications

Olivier Bockenbach¹, Michel Mangin², and Sebastian Schubert¹

¹ Mercury Computer Systems, GmbH, Lepiusstr. 70,
12163 Berlin, Germany,

² Mercury Computer Systems, Sarl, 19 Avenue du Quebec,
91951 Courtaboeuf, France
{obockenbach, mmangin, sschuber}@mc.com

Abstract. Over the last decade, many methods for adaptively filtering a data stream have been proposed. Those methods have applications in two dimensional imaging as well as in three dimensional image reconstruction. Although the primary objective of this filtering technique is to reduce the noise while avoiding to blur the edges, diagnostic, automated segmentation and surgery show a growing interest in enhancing the features contained in the image flow.

Most of the methods proposed so far emerged from thorough studies of the physics of the considered modality and therefore show only a marginal capability to be extended across modalities. Moreover, adaptive filtering belongs to the family of processing intensive algorithms. Existing technology has often driven to simplifications and modality specific optimization to sustain the expected performances.

In the specific case of real time digital X-ray as used surgery, the system has to sustain a throughput of 30 frames per second.

In this study, we take a generalized approach for adaptive filtering based on multiple oriented filters. Mapping the filtering part to the embedded real time image processing while a user/application defined adaptive recombination of the filter outputs allow to change the smoothing and edge enhancement properties of the filter without changing the oriented filter parameters.

We have implemented the filtering on a Cell Broadband Engine processor and the adaptive recombination on an off-the-shelf PC, connected via Gigabit Ethernet. This implementation is capable of filtering images of 5122 pixels at a throughput in excess of 40 frames per second while allowing to change the parameters in real time.

Keyword: Adaptive Filtering, Digital X-Ray, Cell Processor, Real Time, Medical Imaging.

1 Introduction

1.1 Background

Quality becomes one of the most important metrics when a radiologist or a surgeon examines the results from medical imaging. Medical imaging needs to provide the observer with images where the features of interest are highlighted while the interfering

noise is reduced. Over the last decades, the medical imaging community has passionately debated on the approaches and results for noise reduction and edge enhancement. Because image quality has been expressed as resolution and Signal to Noise Ratio (SNR), analyzing the physics of the imaging system and the properties of received signal from which image data is sampled has lead to implementation of spatial and temporal filters. The goal of these two categories is to take advantage of two properties of noise. Its temporal randomness is such that it zeroes itself out over time and its low frequency spectrum allows to use simple low pass filters to smooth it locally out.

Because the constant effort it to reduce the radiation dose during examination, the complexity in the implementation of this type of filter has been growing accordingly, moving from methods where the filter would be uniformly applied to a complete image set to more data dependent filters. Recently, adaptive filtering methods have been introduced^{3,4}.

Adaptive filtering can be defined as a method that takes advantage of a priori knowledge of the properties of the signal in the neighborhood of the location to be filtered. Most sophisticated methods proposed so far take advantage of the high frequency spectrum to keep the information about the underlying structures (such as edges) and use the low frequency components to identify the parameters of the noise. Most recent works use a combination of High Pass (HP) and Low Pass (LP) filters to determine the way the input data should be filtered. Such methods are extremely compute intensive since the computation requires to evaluate the output of several filters before any filtering can be applied. Even on common high-end hardware, the cost issues restrict designers to compromise and to optimize the filter for a given problem, indeed limiting the generality of the solution. Therefore, although those methods show extremely good results, they suffer from various shortcomings that prevent a wide adoption in other modalities than those for which they have been designed for^{6,7}.

Real-time digital X-ray adds a stability issue to the filtering problem. In fact, when the observer looks at a sequence of images and concentrates on certain features in the display, those features need to be filtered and displayed in a consistent way, e.g. variations in the grey level should not appear for obvious reasons. This is why the adaptiveness of the filter needs to be constrained to some basic parameters to avoid negative effects for the observer, such as flickering. Furthermore, post-processing algorithms such as automated segmentation and registration work best when previous processing steps maintain temporal consistency in pixel values.

1.2 Goal of the Study

In this study, we start from a generic adaptive algorithm that can theoretically address all types of de-noising and edge enhancements endeavors. It notably consists of a set of HP and LP filters as well as a recombination function. We show that is possible to:

- identify a set of specific parameters that apply only to the filter components and that allow to cover specific Digital X-ray examinations,
- precompute the filter kernels based on these parameters so that filtering stability is insured for as long as needed,
- adapt the recombination of the filters output in such a way that an application or an operator can easily select the contribution of the different filters from simple ergonomic settings.

The idea behind this algorithm is to separate the noise from the rest of the image by trying to find a structure in the sampled data by assuming that random noise does not follow that structure and that anything without structure can be smoothly extracted from the input data. As appealing as this scheme is for areas with low frequency variations, this dichotomy does not work well in the general case where noise appears in conjunction with structured features such as tissues and bones. The separation of noise from tissue requires more sophisticated methods. We have chosen to locally track the structure of the underlying features and more specifically, their orientation. Similar work was proposed by Raupach et al¹ in the 3D case.

To give an overview, section 2 of this document describes the structure of the filter and the way the parameters influence the processing. A detailed mapping on the hardware is also described in this section. Section 3 gives a short description of the data set. For this document, we have selected static images from different X-Ray examination. All images are taken with low radiation dose. In section 4, filtered images are shown and we discuss the results with respect to the parameter used.

2 Methodology

2.1 Principles

In this section we show how to evaluate the structure of the features in the images and steer the way the filter operates. Two main properties are combined in the filtering process : the image gradient and the local orientation of structures. In areas with low gradients, the filter operates like a smoothing function reducing noise while in areas with larger gradients, the filter becomes more an identity function, preserving the high frequency information given by the edges. In the areas with larger gradients, it is important to evaluate the dominant direction in which the filter shall operate. This main direction defines the level of local anisotropy of the filter. In areas where there is no major direction, the filter turns into an isotropic filter.

In other areas, there may be fast and strong orientation changes in the main direction due to fading in and of main structures. The filter direction must smooth out the local inconsistencies of the filter direction. A local Gaussian filter can be used for that purpose.

Fig. 1. Structure of the adaptive Filter

2.2 Filter Implementation

The logical design of our implementation is based on the architecture of Westin et al⁴ depicted in figure 1. Although the structure seems very similar, the filter presented in this document defers in the type of the filters and their parameters.

The images to be filtered follow two different paths. For the Basic Filter Set, we have opted for a simple filter LP given by equation [1]. The complementary Band Pass (BP) filters are based on equation [2] and oriented along three directions, regularly distributed on the interval $[0, \pi]$. first path gives the estimation of the local anisotropy.

$$Flp(W) = \begin{cases} \cos^2(\frac{\pi W}{2W_c}) & \text{for } W \leq W_c \\ 0 & W > W_c \end{cases} \quad (1)$$

$$Fhp(W) = \begin{cases} 1 - \cos^2(\frac{\pi W}{2W_c}) & W \leq W_c \\ 1 & \text{for } W_c < W \leq \pi - W_c \\ \cos^2(\frac{\pi(W - \pi + W_c)}{2W_c}) & \pi - W_c < W \leq \pi \end{cases} \quad (2)$$

The anisotropic filters are derived from band pass lognormal filters as given in equation [3], with a 6dB of sensitivity bandwidth B and center frequency W_0 . They are oriented in the same directions at those used in the basic filter set.

$$H_p(W) = e^{-4 \ln^2(\frac{W}{W_0}) / (B^2 \ln(2))} \quad (3)$$

The processing of the anisotropic filters is followed by a relaxation step and a tensor mapping step, this path provides the information that will control the use the basic filter set..

We have implemented all filters in the spatial domain with a convolution function. The filter kernel size has been set to 15×15 . It is initialized with the values depending on the type of the filter, the central frequency W_0 and the cut off frequency W_c . The output of the filters is represented with tensors.

The tensor relaxation step is implemented as a simple convolution with Gaussian filter. Its kernel size is again 15×15 .

In the general case, the tensor mapping steps requires the evaluation of the Eigenvalues of the tensors. However, since we are in the two dimensional case, we have used the direct mathematical solution for the Eigenvalues instead of using more complex iterative methods. The output of this step at location x is driven by parameters $(\alpha, a, b, c, \lambda_1)$ defining the trade off between resolution and stability, as given in equation 4.

$$o(x) = c(1 - \lambda_1') + \frac{\lambda_1'^b}{\lambda_1'^{a+b} + \alpha^b} \quad (4)$$

Fig. 2. System Block Diagram

Where :

$$\lambda_i = \frac{\lambda_1}{\lambda_{i_{\max}}} \quad (5)$$

and where $\lambda_{\max}$ is the largest Eigenvalue of the tensor. λ_1 is the largest eigenvalues of the tensor while the parameters **a**, **b** and **c** are user definable and arguments to the tensor mapping step. It is worth noticing that an output value larger than 1 increases the high-frequency components of the input image and hence acts as an edge enhancement feature in the filtering process.

The final adaptive filter step then merely consists of a linear combination of the output of the basic filters weighted by the control tensor resulting from the tensor mapping step.

2.3 Hardware

The system hardware prototype is based on a Mercury Cell Technology Evaluation System (CTES) which consists of a BladeCenterTM¹ hosting one Dual Cell-Based Blade (DCBB)². For convenience, we connected the DCBB to a visualization server via Gigabit Ethernet (GE), intended to run the control part of the visualization application.. The relationship between these components is shown in Fig. 2.

The visualization can be considered as a standard PC and runs an off-the-shelf Windows XP³ operating system. It provides all the required means for network communications and supports the drivers for the selected GPU board. In addition, it also enables us to run Amira⁴ 3.1.1 for real-time visualization of the reconstructed slices and for volume rendering.

The DCBB is a board built around a dual Cell Broadband Engine (CBE) architecture and designed to be collocated with other boards, potentially other DCBBs, in a Blade-Center. DCBB boards feature two CBE processors in Symmetric Multi Processor

¹ International Business Machines Corporation, Armonk, New York, USA.

² Mercury Computer Systems, Inc., Chelmsford, Massachusetts, USA.

³ Microsoft Corporation, Redmond, Washington, USA.

⁴ Mercury Computer Systems, Inc., Chelmsford, Massachusetts, USA.

(SMP) mode, clocked at 3 GHz (Fig.3). Each CBE processor has 512 MB of Rambus (XDR) memory directly connected to its pins. A sound description of the CBE processor has been given by Pham et al⁸.

2.4 Algorithm Mapping

We have implemented the major part of the processing on one of the CBE processors of the DCBB. It receives the image flow over Gigabit Ethernet at a throughput rate of 30 frames per second. The images are encoded as 16 bits integers. They are first converted to floating point values on the PPE before being distributed over the SPEs.

The processing step on the DCBB are implemented in such a way that they produce for output the norm of the filtered images, the LP component and the weighted sum of the HP filters along with important resulting values such as the largest Eigenvalues. The output images are sent over the network to the visualization server where the smoothing and the adaptive combination take place.

Fig. 3. Dual Cell Based Board

The application running on the visualization server captures the user settable values for α , a , b and c in equation [4], perform the computation of the resulting images accordingly. Finally the images are sent to the monitor and other post-processing steps.

3 Data Sets

To evaluate our implementation, we have selected two different data sets. The first one consists of two low-dose images (Fig 4a and 4b), along with a dosimeter and a resolution grid. We used those images to set the central and cutting frequencies of the filters and evaluate their effects on the noise and on the enhancement of features. With those parameters identified we have run the filter process against several low-dose cardiac sequences from which a sample image is shown is Fig 4c.

a: Low-dose dosimeter – b: Hand with resolution grid – c: cardiac sequence

Fig. 4. Samples from the data sets

4 Results and Discussion

Fig. 5 shows a typical result obtained when using this adaptive filter with three filter directions. Tests with a higher number of directions have been made without showing significant improvement in the image quality; the differences between images filtered with three and with seven direction show only minor effects in the presence of edges.

Fig. 5. Image 4-a filtered

The tensor mapping uses the following parameters $\alpha = 0.1$, $a = 0.5$, $b = 2$ and $c = 0.25$.

The image shows that the noise on the dosimeter has significantly gained in homogeneity and the contrast has improved the definition of the lines. There is no blurring at the edges and the overall impression is that the filtered image is of better quality.

Figure 6 represents a part of figure 4-b. 5-a is a part of the original image, unfiltered with a significant level of noise. Nevertheless, the features can clearly be identified by the observer. On the other hand, extracting features by automated segmentation

a : Original image b : Filtered $\alpha = 0.1$ c : Filtered $\alpha = 0.2$

Fig. 6. A part of image 4-b, unfiltered and filtered

may become challenging as one can see in the upper right corner of the resolution grid. Indeed the borders of the letters “LP/mm” are slightly blurred. Automated segmentation may have difficulties with smoother edges such as large bones.

Figures 6b and 6c represent filtered versions of the original image 6-a, with respectively the filter parameters $\{\alpha = 0.1, a = 0.5, b = 2, c = 0.25\}$ and $\{\alpha = 0.2, a = 0.5, b = 2 \text{ and } c = 0.25\}$. Both images 6b and 6c show a significant reduction of the noise. Figure 6b also shows a stronger structure of the image. Even though that structure exists in figure 6a, it has been enhanced to a point where it disturbs the image analysis. Automated segmentation is tempted to find edges in areas where the filter has enhanced them. Figure 6c demonstrates the effect of parameters **a**, which drives the contribution of the LP filters. It can be clearly seen that there are significantly less undesired enhancement of structures in figure 6c than in 6b. The first experiments with the interactive segmentation module of Amira demonstrate better identification of the features in the image 6c than on 6b.

Fig. 6d. Profiles along the black line in figure 6a

Figure 6d shows the data profile along the black line shown in fig. 6a. One can see that the level of noise is lower in the filtered images and that the contrast has been enhanced. It is also worth noticing that the resolution has not been affected.

We have also applied the filters on cardiac images. Figure 7a shows the original image taken from a sequence of 25 images. We have applied the same parameters to the filter as in the previous experiment. We obtained the images shown in 7b and 7c. It is clear to see that 7b smoothes out a significant part of the noise and enhances the dynamics of the image. What cannot be identified on those images but is clearly visible on the real-time display of the sequence is that the structure that the filter has enhanced low-density tissue that could have been taken as noise uniformities. Figure 7c shows a significant enhancement of the contrast in all structures, still without blur. The catheter is clearly visible as well as the boundaries of the bubble.

In the case of the cardiac sequence, we have loaded the 25 images into the Cell's XDR memory and performed the filtering on those images. The current implementation showed that we could sustain 41 slices/sec for images of 1024*1024 pixels. The filtered images were sent to the visualization server where the remaining parts of the adaptive combination of the filter outputs were combined. We could sustain a display rate of 50 frames per second.

Fig. 7. Image 4-c, unfiltered and filtered

The very promising and important result of this study is that it seems possible to create :

- a set of filters that defines the variables to control the adaptive filtering of a sequence of images,
- a set of parameters that make the combination of the filter outputs in such a way the resulting image combines a user settable mixture of noise reduction and edge enhancement.

Furthermore, today's multi-parallel technology can carry out the computation in real-time. There are also direct implications in the proposed architecture of this adaptive filtering. Firstly, the big burden of the processing is performed by one computer which can be embedded into the existing X-Ray cabinet and the visualization of the images only consists of a relatively light processing step. It is hence possible to get several different monitors that display the images with other filtering parameters. This makes it also possible to direct the images flow to other automated post-processing application with specific filtering requirements while keeping a visualization monitor with user friendly settings.

5 Future Work

The work on the implementation of adaptive filtering will follow several directions. The first one is a direct application of the results we have obtained today to automated segmentation and computer vision. We have archived some preliminary results with this adaptive filtering technique that show interesting aspects in this area. Figure 8 displays another experiment where we have put the tensor mapping parameters to $\{\alpha = 0.4, \mathbf{a} = 2, \mathbf{b} = 2 \text{ and } \mathbf{c} = 4\}$. The picture shows extreme structuring of the features of the image. It is also remarkable that features with very high contrast compared their neighborhood are even more accentuated. This is typically the type of filtering that computer vision applications require. Applying a simple threshold on such an image allows to extract the most important features that significantly simplifies the positioning problem in robotic surgery.

Fig. 8. Image 4-a filtered

The second axis of research consists in applying this technique to the filtering of projection data in the spiral CT applications as well as the filtering of the 3D volumes resulting from the CT reconstruction process.

On the pure technical side, the filtering process will include a multi-resolution pyramid. The expected outcome from this modification is yet another improvement in the image quality. Keeping a constant image quality, we would be capable of reducing the convolution kernel size. This processing burden reduction will have a direct impact on the throughput and make it possible to sustain higher image rates and handle larger images.

Acknowledgements

The authors gratefully acknowledge the help from the Performance Algorithm Group from Mercury Computer Systems for its support in programming the Cell Broadband Engine processor.

References

1. R Raupach, H Bruder, K Stierstorfer, C Suess, T Flohr, *A novel approach for efficient edge preserving noise reduction in CT volume data*, RSNA 2005, Paper SSQ16-03
2. J. Gurung, M. Fawad Khan, A. Maataoui, C. Herzog, R. Bux, H. Bratzke, H. Ackermann, T.J. Vogl, *Multislice CT of the pelvis : dose reduction with regard to image quality using 16-row CT*, Eur. Radiol. (2005) 15:1898-1905
3. H Öktem, K Egiazarian, J Niittylahti, J Lemmetti, *An Approach to Adaptive Enhancement of Diagnostic X-Ray Images*, EURASIP Journal on Applied Signal Processing 2003:5,430-436
4. C F Westin, J Richolt, V. Moharir, R Kikinis, *Affine adaptive filtering of CT data*, Medical Image Analysis 4 (2000) 161-167
5. L Haglund, *Adaptive Multidimensional Filtering*, PhD Dissertation, University Linköping, Sweden
6. K F Westin, L Wigström, T Looock, L Svöqvist, R Kikinis, H Knutsson, *Three Dimensional Adaptive Filtering in Magnetic Resonance Angiography*, Journal of Magnetic Resonance Imaging 14:63-71 (2001)
7. M Kachelriess, O Watzke, W Kallender, *Generalized multi-dimensional adaptive filtering for conventional and spiral single slice, multi slice and cone beam CT*, Med. Phys.18 (4), April 2001.
8. D.Pham, S. Asano, M. Bollinger, M.N. Day, H.P. Hofstee, C. Johns, J. Kahle, A. Kameyama, J. Keaty, Y. Masubuchi, M. Riley, D. Shippy, D. Stasiak, M.Suzuoki, J. Warnock, S. Weitzel, D. Wendel, T. Yamazaki, K. Yazawa, *The Design and Implementation of the First-Generation CELL Processor*, ISSCC 2005.