Noise Reduction and Image Enhancement Algorithms for Low-Dose X-Ray Fluoroscopy

Til Aach, Dietmar Kunz

PHILIPS GmbH Forschungslaboratorien
Weißhausstr. 2, D-52066 Aachen, Germany
aach@pfa.research.philips.com

Raoul Florent, Shérif Makram-Ebeid

Laboratoires d'Electronique PHILIPS SAS
22 Ave. Descartes, F-94453 Limeil-Brevannes, France

In this paper we describe several filtering techniques for the reduction of quantum noise in low-dose X-ray images, specifically X-ray fluoroscopy images. Quantum noise occurs inherently in low-dose X-ray imaging due to the very low X-ray quantum counts. As a dynamic imaging modality, both temporal and spatial filters can be applied to fluoroscopy. After describing our quantum noise model, we first develop a new motion-adaptive temporal filter based on a recursive structure. Compared to other motion-adaptive filters, our filter is much less prone to generate artefacts known as noise tails*. We then describe spatial filtering techniques derived from the concept of* spectral amplitude estimation*. These filters work in the spectral domain, and can conveniently be tailored to the statistical properties of quantum noise. The noise reduction performance of these filters is evaluated qualitatively as well as quantitatively.*

Keywords: X-ray fluoroscopy, quantum noise, temporal recursive filtering, spectral amplitude estimation.

1 Introduction

X-ray fluoroscopy is a medical imaging technique used for *on-line* monitoring of certain clinical procedures, like catheterization. The resulting moving images are viewed immediately on a monitor while the clinical procedure is being performed. Patient and medical staff are hence exposed to X radiation during prolonged periods of time. To keep these exposures to a minimum, very low X-ray dose rates are used, leaving only between 20 - 500 X-ray quanta per pixel for the imaging process. This leads to considerable degradations of image quality through *quantum noise*. An example of a fluoroscopy image is shown in Fig. 1.

This contribution describes both temporal and spatial filtering algorithms for quantum noise reduction. Toward this end, we first derive a statistical model for quantum noise. We then develop a temporal motion-adaptive first order recursive filter, which, compared to other filters of the same type, is much less likely to cause artefacts known as *noise tails*. In section 4, we describe spatial noise reduction filters based on so-called *spectral*

amplitude estimation. Working in the spectral domain, it is shown that these filters are well suited to take the statistical properties of quantum noise into account. The noise reduction performance of each filter is evaluated by a qualitative and a quantitative processing example.

2 The Quantum Noise Model

As X-ray quantum noise is by far the dominating type of noise in X-ray fluoroscopy, other noise sources can in general be neglected (*quantum-limited imaging*). Quantum noise as it originates from the X-ray beam is Poisson-distributed [1], i.e. its noise variance is identical to the mean number of absorbed X-ray quanta per pixel and frame. Its noise power spectrum (NPS) is flat. Quantum noise in the observed images has undergone filtration by the imaging system's transfer function, resulting in a spatially lowpass shaped NPS [1, 2]. Additionally, the video signal may be subjected to an intended nonlinear gain (*white compression*). The relation between noise power and signal is therefore given by a linear rise over low and medium intensities caused by the Poissonian nature, and a drop-off over higher intensities due to the white compression (see Fig. 1).

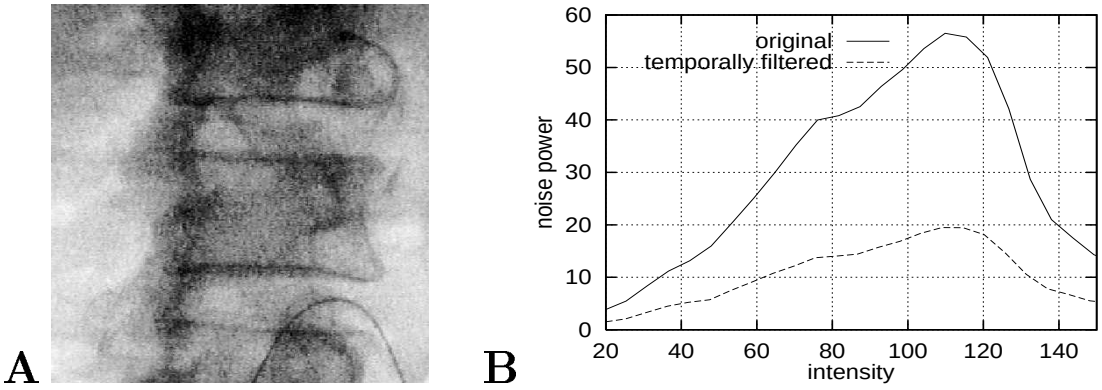

Figure 1: **A** Enlarged central part of a frame from a fluoroscopy image sequence taken during a medical intervention. **B** Noise variance as a function of image intensity for the same sequence. The linear rise of the curve is due to the Poisson distribution of noise. The decrease for high intensities is caused by the non-linear white compression. Also shown is the noise power which remains after temporal filtering.

We assume both the signal dependence of the quantum noise power and the NPS as known from image acquisition parameters and system measurements.

3 Temporal Filtering

Our temporal filter applies the following first-order recursive equation to every pixel n:

$$x(t,n) = x(t-1,n) + K(t,n) \cdot [y(t,n) - x(t-1,n)] \ , \tag{1}$$

where $y(t,n)$ is the original noisy signal of pixel n at (discrete) time t, and $x(t,n)$ the filtered signal. All equations of this section apply to every pixel, however, to simplify

the notation, we drop the pixel variable n in the following. To determine the filter gain $K(t)$, we consider a generalized weighted mean square error criterion $E(t)$, which is to be minimized. $E(t)$ is given by

$$E(t) = [x(t) - y(t)]^2 + \sum_{k=1}^{\infty} a(t,k) \cdot [x(t) - y(t-k)]^2 \quad , \tag{2}$$

where $a(t,k)$ is a measure indicating how likely it is that *no* discontinuity (i.e. no motion) has occurred between t and $t-k$. Denoting the "continuity probability" $a(t,t-1)$ between t and $t-1$ by $\alpha(t)$, and assuming statistical independence, we have $a(t,k) = \prod_{l=1}^{k} \alpha(t-l)$. It can then be shown that the filter gain $K(t)$ obeys the following recursive updating rule:

$$K(t) = K(t-1)/[K(t-1) + \alpha(t)] \quad . \tag{3}$$

We evaluate $\alpha(t)$ depending on the difference $\Delta(t) = y(t) - x(t-1)$. It is intuitively clear that continuity is less likely if $|\Delta(t)|$ is relatively large. The "continuity probability" should therefore be a monotonically decreasing function of $|\Delta(t)|$.

The recursive filter as described so far is adaptive to motion, because motion results in increased values for $|\Delta(t)|$, and hence decreased values of $\alpha(t)$. The filter gain given by eq. (3) then tends to larger values, thus reducing temporal integration in eq. (1). Blurring of moving objects is thus avoided. The downside of this mechanism is that the reduced filtration results in an increased noise level in the filtered data $x(t)$ *even when the moving object has already passed.* These increased noise levels in turn increase the observed values for the subsequent difference image $|y(t+1) - x(t)|$, thus pretending motion. It hence takes a few frames for the filter to increase integration again once a moving object has passed. The result is a noise tail following each moving object.

To reduce this "hysteresis"-like artefact, we have integrated into our filter an estimation algorithm for the standard deviation $s(t)$ of $\Delta(t)$. Clearly, $s(t)$ depends on the (signal dependent) noise power $\sigma^2(y(t))$ for the unfiltered data $y(t)$, which is known from the relationship of quantum noise power over signal intensity discussed in section 2. Additionally, $s(t)$ depends on the *filtered* noise level in $x(t-1)$, i.e. the filtration history. The influence of the filtration history on the noise level s(t) can for instance be captured by the previous filter gain $K(t-1)$.

The estimated noise standard deviation $s(t)$ is now used to *normalize* the differences $\Delta(t)$ before determining the "continuity probabilities" $\alpha(t)$, i.e. $\alpha(t)$ is now defined as a monotonically decreasing function of $|\Delta(t)|/s(t)$. This relation is illustrated in Fig. 2. The total effect of doing so is that the filter now "knows" the increased noise levels caused by motion adaptation, and examines the difference $y(t) - x(t-1)$ in this light. Combined with the calculation of the new filter gain $K(t)$ in eq. (3), this results in a much faster filter adaptation, and hence shorter noise tails.

A temporally filtered version of the image in Fig. 1 is shown in Fig. 2. As the full noise reduction performance can only be appreciated on a CRT monitor, we have also quantitatively evaluated the noise power which remains after filtering. This noise power is still signal dependent, and also depicted in Fig. 1, showing a noise reduction by more than 60%. Naturally, this result is motion dependent.

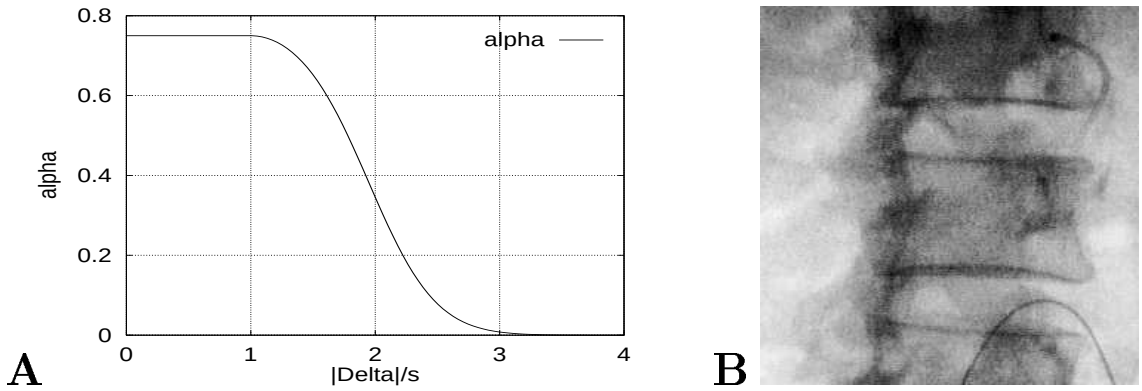

Figure 2: **A** "Continuity probability" α as a function of the normalized difference between new observation and last filter output. **B** Temporally filtered version of Fig. 1.

4 Spectral-Domain Spatial Filters

The described temporal filter is particularly well suited for full frame rate fluoroscopy with 25 - 30 frames per second, and moderate amounts of motion. In so-called pulsed fluoroscopy, where the frame rate may be as low as one frame per second, or in the presence of high and erratic local motion, the effectiveness of temporal filtering can, however, be strongly reduced. In this section, we therefore describe a spatial filtering technique the performance of which is independent of frame rate and motion.

As the NPS of quantum noise is lowpass-shaped over spatial frequency, low frequency noise components contribute strongly to the total noise power. In order to apply noise reduction evenly to high *and* low spatial frequencies, we employ a spatial-frequency domain method which can be tailored easily to a coloured NPS [2, 3]. The central idea of our approach is to attenuate observed noisy frequency coefficients $Y(k)$ depending on their *instantaneous* signal-to-noise ratios (SNR) r_k by

$$X(k) = Y(k) \cdot h(r_k), \quad \text{where } r_k^2 = |Y(k)|^2 / \Phi_n(k) \ , \tag{4}$$

with $\Phi_n(k)$ denoting the (estimated) NPS, and $X(k)$ the estimated noise-reduced coefficient. The *attenuation function* $h(r)$ increases monotonically with the SNR r_k, and asymptotically reaches one for large SNR. Eq. (4) hence attenuates spectral coefficients which are likely to represent mainly noise. As $h(r)$ is real-valued, eq. (4) is a zero-phase filter affecting only the amplitude of $Y(k)$. In speech processing, this approach is therefore known as *spectral amplitude estimation* [4]. The frequency coefficients $Y(k)$ are obtained by a standard block transform like the Discrete Fourier Transform (DFT) or Discrete Cosine Transform (DCT). To permit adaptivity to local detail, processing is based on small overlapping image blocks.

Fig. 3 shows the block diagram of the noise filter. The box labelled "noise model" contains a normalized NPS which integrates to unity (NPS-*shape*). To determine the absolute height or *scale* of the NPS, the relation of noise power over intensity is also stored in a look-up table (LUT) within the noise model box. After decomposing the observed image into blocks and subjecting these to the DFT, the spatial frequency index k of each coefficient $Y(k)$ is used to read the corresponding coefficient of the normalized NPS. To

scale the normalized NPS to the mean noise level in the image block currently processed, the mean grey level $Y(0)$ of the block is used to read the noise power σ^2. Multiplication of the normalized NPS coefficient with σ^2 yields the NPS-coefficient $\Phi_n(k)$ [3]. The instantaneous SNR r_k as defined in eq. (4) can now be computed. The corresponding attenuation factor $h(r_k)$ is then read from a LUT labelled "attenuation function", and applied as a weighting factor to the observed coefficient $Y(k)$. After taking the inverse DFT, the noise-reduced image is reconstructed by an overlap-add operation.

The exact shape of $h(r)$ can be determined based on a mimimum mean square error (MMSE) estimation approach [3]. For each image block, we assume that the undistorted image signal appears in only a few frequency coefficients. With appropriate statistical models for "noise only" and "signal-plus-noise" coefficients, the following expression for $h(r)$ can be derived [3]:

$$h(r) = (1 + \lambda \exp(-\gamma r^2))^{-1} \quad , \tag{5}$$

where γ is an adjustable weight similar to the one used in generalized Wiener filters. The factor λ depends on signal and noise variances, but is seen here as an adjustable parameter, too. This attenuation function is depicted in Fig. 3.

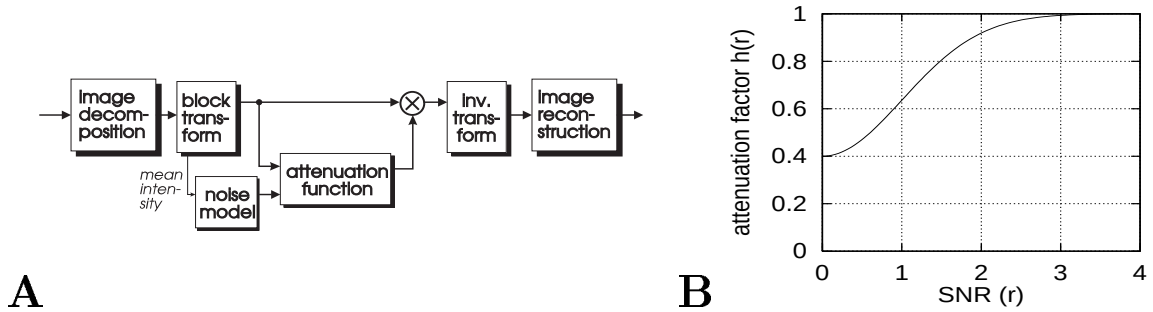

Figure 3: **A** Block diagram of spectral amplitude estimation. **B** Attenuation function $h(r)$ as a function of the instantaneous SNR r, plotted for $\lambda = 1.5$, $\gamma = 1$.

To improve the performance of the above spectral amplitude estimation filter particularly with respect to perceptually important oriented patterns formed by lines and edges, we have extended our algorithm by integrating sensitivity to local orientation. When using the DFT, presence of an oriented pattern like a catheter within an image block results in a concentration of spectral energy along the line perpendicular to the spatial orientation and passing through the origin. For each transformed block, we estimate local orientation by evaluating an "inertia" matrix [5]. The eigenvectors of this 2×2-matrix determine in a least-squares sense the directions along which concentration of energy is strongest (local orientation) and least, respectively. From the corresponding eigenvalues we calculate a measure indicating how distinctive the detected local orientation actually is.

On the one hand, we use this orientation information to reduce the attenuation of spectral coefficients situated along the direction of local orientation, with this adaptivity being the more pronounced, the more distinctive the detected orientation. This behaviour enables achieving a strong reduction of noise without losing perceptually important detail. On the other hand, availability of orientation information also allows selective enhancement of oriented structures. Details of this approach can be found in [5].

Fig. 4 shows the processing result for the image depicted in Fig. 1, based on $\lambda = 1.5$ and $\gamma = 3$. Orientation information was not exploited in this case. Also shown is a comparison of the original noise power and the noise power after filtering. The achieved noise power reduction is about 60%.

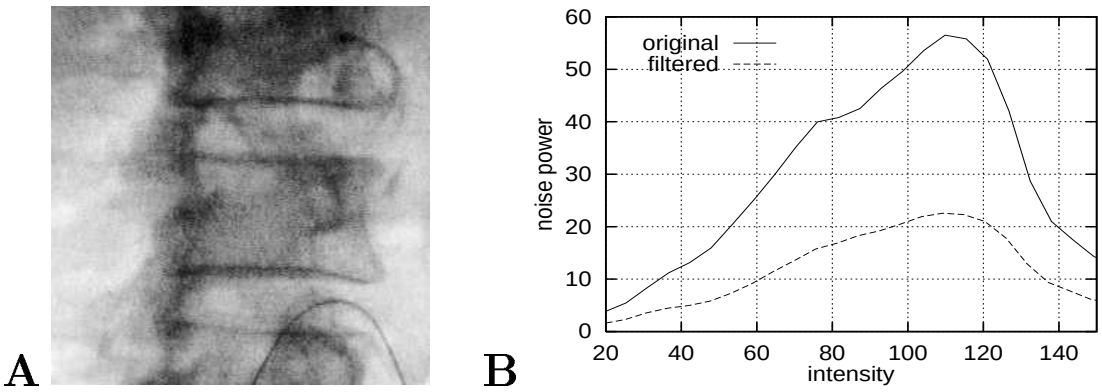

Figure 4: **A** Spatially filtered version of the frame in Fig. 1. **B** Noise variance versus intensity for original and processed image.

5 Conclusions

The temporal filter in section 3 was developed based on a *motion adaptive* mean square error criterion. Its advantages are a simple recursive structure combined with a fast adaptivity to even small moving objects. The strength of the spatial filter is that it takes the known statistical properties of quantum noise into account, whereas only very weak assumptions about the unknown signal are needed. This is particularly important in the context of medical images, for which appropriate modelling is difficult if not impossible. Both the described temporal and spatial filtering algorithms allow a good noise reduction to be achieved, without generating unacceptably strong artefacts.

References

[1] Spekowius G et. al.: Simulation of the imaging performance of X-ray image intensifier/ TV camera chains. Medical Imaging 1995 (SPIE Vol. 2432):12–23, 1995.

[2] Aach T, Schiebel U, Spekowius G: Digital image acquisition and processing for medical X-ray imaging applications. Proc. Intl. Symp. Electr. Photog. (ISEP), Cologne, Sept. 21–22, 1996 (in print).

[3] Aach T, Kunz D: Spectral estimation filters for noise reduction in X-ray fluoroscopy imaging. Proc. EUSIPCO-96 (Ramponi G et. al. (eds.), Edizioni LINT), Trieste, Sept. 10–13, 1996, 571–574.

[4] Vary P: Noise suppression by spectral magnitude estimation – mechanism and limits. Signal Processing 8(4):387–400, 1986.

[5] Aach T, Kunz D: Anisotropic spectral magnitude estimation filters for noise reduction and image enhancement. Proc. ICIP, Lausanne, Sept. 16–19, 1996, 335–338.