Filtering of ICU Monitor Data to Reduce False Alarms and Enhance Clinical Decision Support

Apkar Salatian, Lawrence Oborkhale
School of Information Technology and Communications
American University of Nigeria,
Yola Bypass, PMB 2250,
Yola, Nigeria
{apkar.salatian, lawrence.oborkhale}@aun.edu.ng

Abstract

The monitors in the Intensive Care Unit generate alarms whenever a signal passes beyond a preset limit. Such an approach for alarming generates many false alarms because, in general, clinically insignificant events cause signals to go beyond these limits e.g taking a blood sample. Here the alarm has been caused by a clinically insignificant event, not a disturbance in the patient's physiology. If these limits are set to the maximum allowable physiological deviation from the normal or expected value, the monitor will alarm when the patient is already in a serious condition. Likewise, if the limits are adjusted to increase sensitivity, the monitor will be more prone to giving false alarms. There is, therefore, a strong need to reduce the number of false alarms. Our approach to reducing false alarms is to use filtering techniques which will not only remove clinically insignificant events but also allow medical staff to view noise free data to enhance clinical decision support. Using real monitor data, in this paper we review a number of filtering techniques and present our findings to determine which is most suitable for the Intensive Care Unit monitors.

Keywords: *filtering, ICU, monitoring*

1. Introduction

Patients are admitted to the ICU either due to major trauma, as a planned extension to their postoperative care, or as a result of a sudden deterioration. The purpose of intensive care is to maintain accurate, continuous observations of the patient's vital functions, and to treat or support a failing or failed biological system - this is achieved by attaching the patient to monitors to record their physiological signals such as the heart rate and blood pressure.

The ICU monitors generate a large volume of data for each patient. The challenge is to interpret these signals and to distinguish between clinically significant and insignificant events.

Clinically insignificant events can be considered as an interval of time where a specific predicate holds, for example line flush being taken, which does not reflect a disturbance in the patient's physiology.

Clinically insignificant events need to be processed for a number of reasons:

- Presently alarms in the ICU are based on a single monitored parameter passing beyond a preset limit e.g an upper threshold for a systolic pressure and a lower threshold on a diastolic pressure. Many clinically insignificant events may occur during the stay of an ICU patient - these will include the taking of blood samples and line flushes. Such events will set off unnecessary alarms because they typically reach values beyond the set thresholds; for example, when taking a line flush a flushing device is pressed which in turn opens a valve between the anti-coagulant solution and the pressure transducer. As a consequence, the pressure received by the pressure transducer is equivalent to the pressure of the pressure bag (usually 200mmHg) which, in turn, will be above the upper threshold for this signal. This alarm has been caused by a clinically insignificant event, not a disturbance in the patient's physiology. Such false alarms are a major problem - studies indicate that 40-75 percent of all alarms are false [1]. Indeed due to this high incidence, many alarms are ignored totally; while some persistent alarms may lead to medical staff to overlook other alarms, which are, perhaps, really important [2].

- Retaining the data obtained during clinically insignificant events and including them in further processing will give an inaccurate history of the patient when, for example, requesting the mean blood pressure.

Such data must be filtered out (ignored) by the clinical staff. One approach to remove clinically insignificant events from ICU monitor data is to use filters from the signal processing community.

In this paper we will study the application of standard signal processing filtering techniques to ICU monitor data and determine their ability to remove clinically insignificant events. By displaying the filtered data graphically on a computer we will provide a multi-media approach to allow medical staff to view noise free data to facilitate better interpretations to be made of the patient state.

The structure of this paper is as follows: section 2 reviews the various filtering techniques from the signal processing community, section 3 describes the results of the application of the filtering algorithms to a data set taken from a neonatal ICU; and final conclusions are in section 4.

2. Filtering

The signals in the ICU contain noise which can arise from a variety of sources. The purpose of the filtering operation is the effective elimination or attenuation of the noise that is corrupting the desired signal. A filter is a device in the form of a piece of physical hardware or software that is applied to a data set in order to de-noise it [3]. Filters are classified as either analog filters or digital filters.

Analog filters operate on continuous time signals such as the heart rate and blood pressure.

However, if these measurements are taken at every fixed time interval then they are considered as discrete signals which digital filters operate on. When digital filters are to be used for an analog input signal it has to be represented by discrete samples obtained in time moments proportional to the sampling interval which must correspond to the Nyquist criteria i.e the sampling frequency (f_s) is at least two times higher than the highest frequency in the signal spectrum [4]. In discrete signal analysis, the frequency (f), as a rule, is represented as a normalized frequency (f*) i.e

$$f^* = \frac{f}{f_s}$$

In our study, f is the frequency of the ICU monitor signal before it is sampled. Since our data is digitized we will use digital filters. Digital filters are, in turn, classified as either linear filters or nonlinear filters. A digital filter is said to be linear if the filtered quantity at the output of the device is a linear function of the observations applied to the filter input, otherwise it is a non-linear digital filter.

A non-linear digital filter locates and removes data that is considered as noise. A non-linear algorithm looks at each data point and decides if that data is noise or a valid signal. If the point is considered as noise then it is simply removed and replaced by an estimate based on the surrounding data points. If the data is not considered noise then the data is not modified at all. Linear filters lack such a decision capability and therefore will modify all of the data. Nonlinear filters are useful for removing very short duration features that have high amplitudes from data [4,5] - such a filter can be thought of as a *noise spike-rejection* filter, and can be effective for removing short duration features, such as those seen in the ICU e.g blood samples.

We have decided to investigate two linear (low-pass and high-pass) and two non-linear (median and average) filters to determine which is most suitable in reducing the number of false alarms generated by the ICU monitors. In applying these digital filters we are representing our input and output signals and characteristics of the filter as discrete instants of time. Moreover, a continuous-time signal may always be represented by a sequence of digital samples that are derived by observing the signal at uniformly spaced instants of time.

When we will perform filtering, all of the techniques will involve a moving window. For historical data, the window is centered on the point x_n i.e if the window is of size 2k+1 the window contains the points x_{n-k} to x_{n+k}. Likewise, for real-time data the window is from the current data point x_n i.e if the window is of size 2k+1 the window contains the points x_n to $x_{n-(2k+1)}$. In what follows only (historical) centered windows will be considered - note that in the case of real-time data the results would be the same but shifted by k values.

Median Filter

The median filter algorithm simply takes the unweighted median of all values within the window, i.e the value which has as many values which are greater as are less than it is taken to be the median.

$$y_n = \text{median}(x_{n-k},\ldots, x_n,\ldots, x_{n+k})$$

Average Filter

An average filter simply takes the mean of all the values within the window.

$$y_n = \text{average}(x_{n-k},\ldots, x_n,\ldots, x_{n+k})$$

Low-Pass Filter

Low-pass finite and symmetric digital filter takes the form [6]:

$$y_n = \sum_{k=-N}^{k=N} c_k \, x_{n-k} \qquad (c_k = c_{-k})$$

where

$$c_k = \frac{\sin(0.4\pi k)}{\pi k} \frac{N}{\pi k} \sin\left(\frac{\pi k}{N}\right) \quad k = 1 \text{ to } N$$

$$c_0 = 0.4$$

The sampling frequency of our data, f_s, was 86400 samples/day (i.e one sample every second or 1 Hertz) which is the typical rate for ICU monitors [7]. The highest frequency that can be defined by the Nyquist frequency is half the sampling frequency i.e 43200 cycles/day. Therefore the frequency range of interest for this application of ICU monitoring data is 0 to 43200 cycles/day.

For demonstration purposes, our low pass filter was designed to let through frequencies f* < 0.2 where the dimensionless frequency f* is defined in terms of the sampling frequency

$$f^* = \frac{f}{f_s}$$

High-Pass Filter

For demonstration purposes, our high pass digital filter was designed to let through frequencies in the range 0.2 < f* < 0.5. The high pass filter coefficients, c'_k, can be derived from the low pass filter coefficients, c_k :

$c'_k = -c_k$ for $k \neq 0$

$c'_k = 1 - c_k$ for $k = 0$

$$c_k = \frac{\sin(0.4\pi k)\cos(0.7\pi k)}{\pi k}\,\frac{N}{\pi k}\sin\left(\frac{\pi k}{N}\right) \qquad k = 1 \text{ to } N$$

$$c_0 = 0.6$$

3. Results

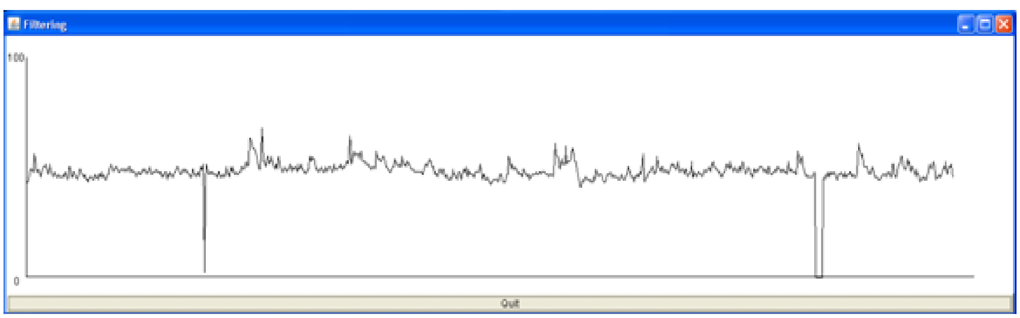

Figure 1. Original Data

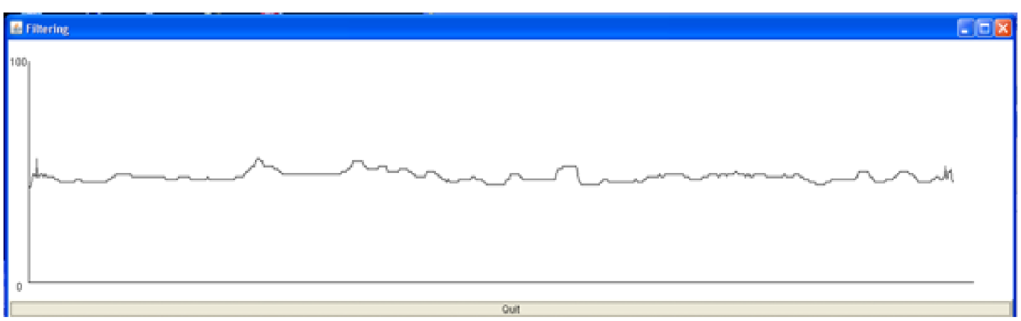

Figure 2. Application of the Median Filter

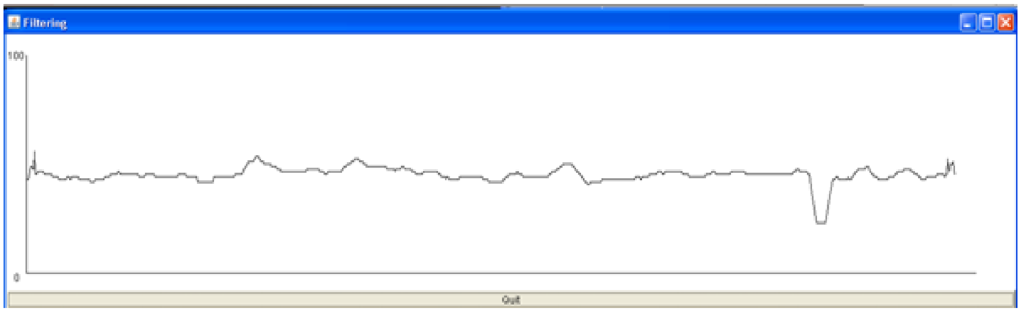

Figure 3. Application of the Average Filter

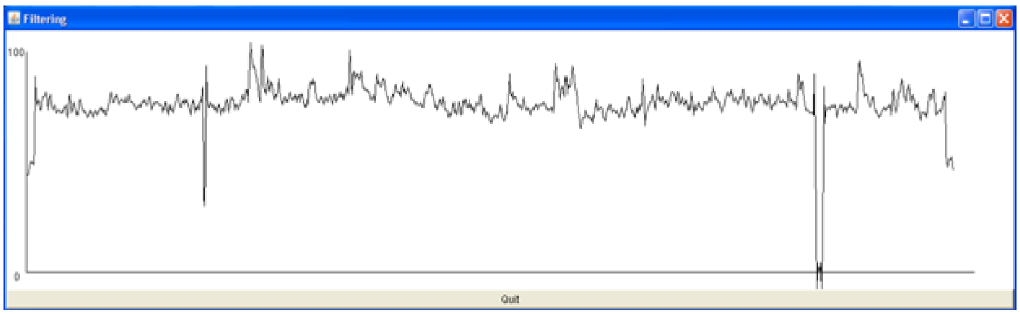

Figure 4. Application of the Low-Pass Filter

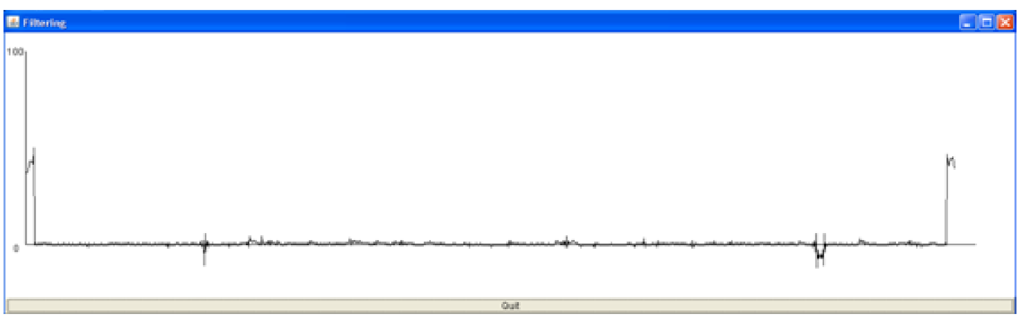

Figure 5. Application of the High-Pass Filter

The results of applying the filters to historical ICU monitor data are now presented - note that in the case of real-time data the filters would be applied at the leading edge and results would be the same but shifted by k values. In order to analyse the results an interesting case was chosen which contains clinically insignificant events. The data set is a blood pressure trace taken from the monitors in a neonatal ICU in the United Kingdom. The frequency of the data was one value every second which is the typical rate for ICU monitors [7]. The original data set is shown in Figure 1. In Figure 1 it can be seen that there are 2 transient features where the blood pressure drops to 0 – these are caused by medical staff taking blood samples (and considered as clinically insignificant events) and would cause false alarms. Figures 2 through to 5 are the graphical results of applying the various filters centered on a point with k = 10. Note that the first and last k values cannot be filtered. We will now look at the results of each filter in turn.

Figure 2 shows the result of applying a median filter – it can be seen that both transient features have been removed and, in turn, would have prevented 2 false alarms. By always choosing the median value in the window as the filtered value, the median filter will remove transient features lasting shorter than k without distortion of the base line signal; features lasting more than that will remain.

In Figure 3 it can be seen that the average filter smoothes the data to a greater extent than the median filter. It can also been seen that the spikes in the data have been attenuated rather than removed.

Figure 4 shows the result of applying a low-pass filter. It can be seen that the low-frequency variations are allowed to "pass through" the filter. In a low-pass filter, the low frequency (long-period) waves are barely affected by the smoothing. The output from the low pass filter is still quite noisy compared with the either the average or median filter because the filter has added more distortion due to the modification of all of the data. It can also be seen that the high frequency data values (such as the large spikes) from the signal are not filtered out.

In contrast Figure 5 shows that the high pass filter has eliminated the low-frequency variations and the high frequency components remain or are attenuated. It can be seen that the spikes (high frequency components) have been attenuated rather than removed. The output signal is around 0 because the high pass filter has removed all the low frequency components of the data.

From the results it is concluded that the median filter with k=10 is the best filter to use as it removes all the very short duration spikes from the data whilst revealing the short duration trends hidden in the raw data. This is because an average filter smoothes transient features, a low-pass filter attenuates noise (noise may have some low-frequency components) and a high-pass filter eliminates low frequency variations and trends leaving only the higher frequency components.

The Median filter is a nonlinear process useful in reducing random noise [8]. It is a non-linear digital filter that is used to remove the impulsive noise from a signal [5, 9]. Median filtering is ideal for smoothing and de-noising applications. Furthermore, it is a more robust method than the other filtering techniques, because it preserves the sharp edges [3]. For properties of the median filter, the reader is advised to read [10].

4. Conclusions

There is a need to remove clinically insignificant events in ICU monitor data. One way to do this is to filter the data which has advantages from a real-time and historical perspective.

From a real-time perspective, there are a number of deficiencies with the monitoring alarms currently used in the ICU. If the monitoring limits are set to the maximum allowable physiological deviation from the normal or expected value, a monitor alarms when the patient already is in a serious condition or on a clinically insignificant event (false alarm). However, if the limits are adjusted to increase sensitivity, the monitor is prone to giving false alarms. Filtering is a way to address this problem.

From a historical perspective, one has to remove clinically insignificant events which would otherwise give an inaccurate history of the patient. Filtered past data can be used to give summaries and patient state assessments to medical staff to allow them to make more informed clinical decisions.

Work is ongoing in two main areas: to find the ideal window size and using hybrid filters. These features, we feel, can be incorporated into a bedside workstation, acting as a background multi-media tool to allow medical staff to view noise free data to enhance clinical decision support.

References

[1] Mylrea, K. C., Orr, J. A., Westenskow, D. R., Integration of Monitoring for Intelligent Alarms in Anaesthesia: Neural networks - Can they help ?, Journal of Clinical Monitoring 9, pages 305-311, 1993.

[2] Mora, F. A., Passariello, G., Carrault, G., Le Pichon, J-P., Intelligent Patient Monitoring and Management Systems: A Review, IEEE, Engineering in Medicine and Biology, pages 23-33, December 1993.

[3] Baxes, G.A., Digital Image Processing, Principles & Applications, Wiley & Sons, 1994.

[4] Hauck, S., The Roles of FPGAs in Reprogrammable Systems, Proc. of IEEE, 86(4), pages 615-638, 1998.

[5] Haralick, R.M., Shapiro. L.G., Computer and Robot Vision, Addison-Wesley, vol. 1, 1992.

[6] Hamming, R. W., Digital Filters, 3rd Edition, Dover Publications, New York, 1989.

[7] Borowski, M, Siebig, S., Wrede, C., Imhoff, M., Reducing False Alarms of Intensive Care Online-Monitoring Systems: An Evaluation of Two Signal Extraction Algorithms, Computational and Mathematical Methods Medicine, Volume 2011, Article ID 143480, pages 1-11, 2011

[8] Guoa, K., Zhanga, X., Lia, H., Meng, G., Application of EMD method to friction signal processing, Mechanical Systems and Signal Processing, Volume 22, Issue 1, Elsevier Ltd, pages 248-259, January 2008.

[9] Yin, L., Yang, R., Gabbouj, M., Neuvo, Y., Weighted Median Filters: A Tutorial, IEEE Transactions on Circuits and Systems, 43(3), pages 157-192, 1996.

[10] Gallagher Jr., N.C., Median Filters: A Tutorial, Proceedings of ISCAS 1988, pages 1737-44, 1988.