Available online at www.sciencedirect.com
ScienceDirect
Journal of Electrocardiology 43 (2010) 486–496
JOURNAL OF Electrocardiology
www.jecgonline.com

A review of electrocardiogram filtering

Shen Luo, PhD,* Paul Johnston
Cardiac Science Corporation, Deerfield, Wisconsin, USA
Received 12 May 2010

Abstract Analog filtering and digital signal processing algorithms in the preprocessing modules of an electrocardiographic device play a pivotal role in providing high-quality electrocardiogram (ECG) signals for analysis, interpretation, and presentation (display, printout, and storage). In this article, issues relating to inaccuracy of ECG preprocessing filters are investigated in the context of facilitating efficient ECG interpretation and diagnosis. The discussion covers 4 specific ECG preprocessing applications: anti-aliasing and upper-frequency cutoff, baseline wander suppression and lower-frequency cutoff, line frequency rejection, and muscle artifact reduction. Issues discussed include linear phase, aliasing, distortion, ringing, and attenuation of desired ECG signals. Due to the overlapping power spectrum of signal and noise in acquired ECG data, frequency selective filters must seek a delicate balance between noise removal and deformation of the desired signal. Most importantly, the filtering output should not adversely impact subsequent diagnosis and interpretation. Based on these discussions, several suggestions are made to improve and update existing ECG data preprocessing standards and guidelines.

Keywords: ECG filter; Frequency domain; Aliasing; Linear phase; Magnitude distortion; Ringing artifact; ECG standards

Introduction

Electrocardiographic devices, including interpretative electrocardiogram (ECG) units, stress ECG systems, bedside ECG monitors, ambulatory recorders, and others, are primary tools used by routine clinical practices for conducting cardiovascular diagnostic and monitoring procedures.[1] The accurate recording and precise analysis of the ECG signals are crucial due to their extensive applicability and also the high-performance expectations of medical professionals. Accurate interpretations of ECGs have always relied heavily on state-of-the-art signal processing.

The preprocessing modules for an electrocardiographic device contain multiple levels of signal manipulation and detection routines, which start by converting analog signals into digital data that is used for analysis, interpretation, and presentation (display, printout, and storage). This article aims to discuss a key ECG processing topic: ECG filtering and data accuracy in a modern electrocardiographic device.

There are 4 typical filter processes in an ECG device: (*a*) anti-aliasing and upper-frequency cutoff, (*b*) baseline wander suppression and lower-frequency cutoff, (*c*) line-frequency rejection, and (*d*) muscle artifact reduction. Use of additional proprietary filter algorithms (such as cubic spline technique,[2] time-varying muscle artifact filter,[3] source consistency filter,[4] etc) to manipulate waveform data is not included in this discussion. In this review paper, we first introduce the basis of filtering in the time and frequency domains because these concepts are fundamental to the discussion of data accuracy. For each filter process, we then pick 1 or more important topics regarding inaccuracy issues to summarize a compliance review and a discussion of the implications.

Reviewing the basis of ECG filtering

Filtering, magnitude (or amplitude), and phase distortions

A given signal in the time domain can mathematically be represented in terms of its magnitude and phase responses in the frequency domain, whereas a given filter with its impulse response in the time domain can be characterized by its magnitude and phase responses in the frequency domain (see Fig. 1).

Except for some very special situations (such as an all-pass filter), a filter is generally designed to attenuate or remove some frequencies from the input data. We expect that a filter only removes the noise without changing the desired signal. In the real world, noise and desired signals often

* Corresponding author. Cardiac Science Corporation, 500 Burdick Parkway, Deerfield, WI 53531, USA.
E-mail address: sluo@cardiacscience.com

0022-0736/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jelectrocard.2010.07.007

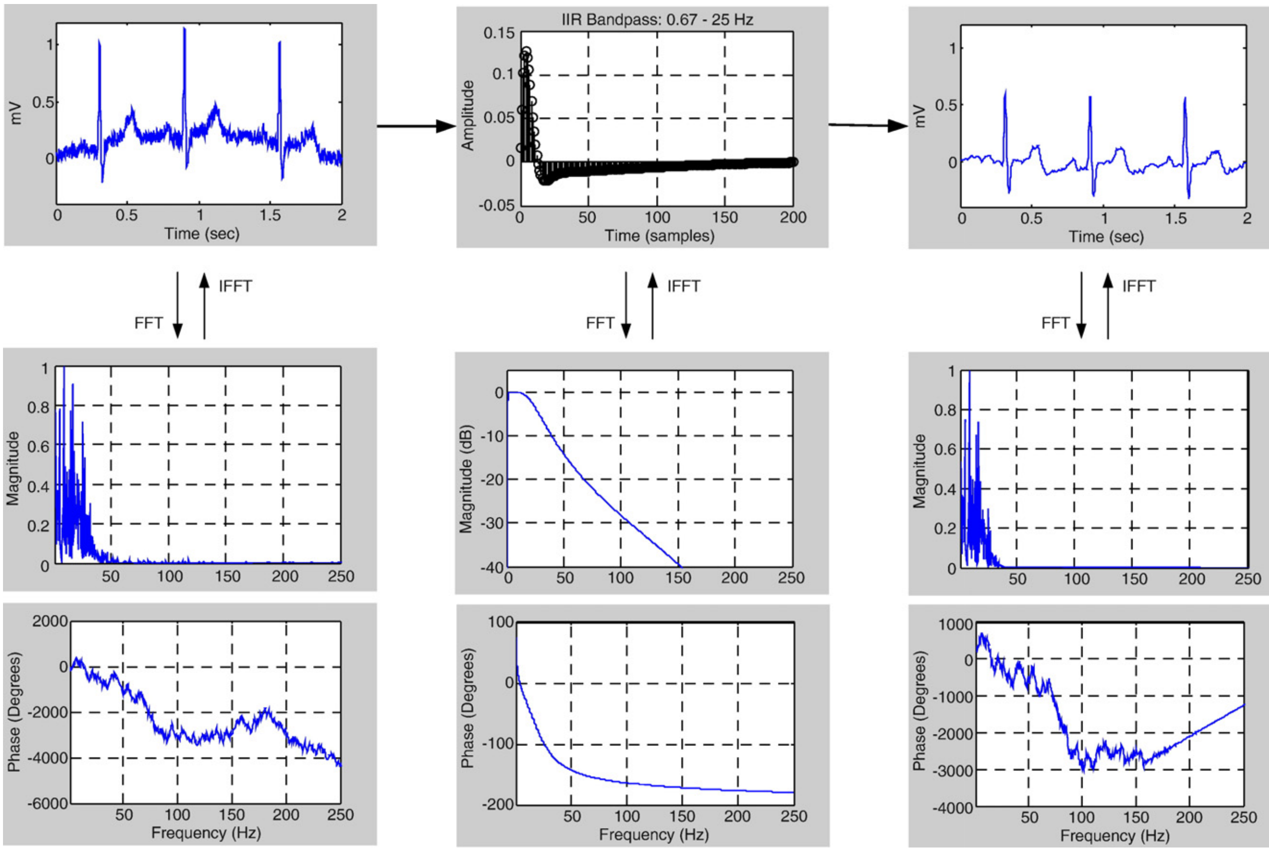

Fig. 1. Both signals and filters can be mathematically converted (mapped) from the time domain to frequency domain, or vice versa. When an ECG signal (left) goes through a filter (middle), the filtered ECG trace (right) could be altered, although noise is reduced. The frequency response, including both magnitude and phase responses, can help us to evaluate and comprehensively design a better filtering system to minimize the distortions and maximize the signal to noise ratio.

overlap in the frequency domain. As a result, when a filter attenuates the frequency components of noise, the overlapping desired signal can also be impacted, causing magnitude distortion of the desired signal.

In addition to the magnitude response, the phase response is another property of a filter. Phase distortion occurs when a filter's phase response is not a linear function of frequency so that the phase shift is not directly proportional to the frequency. Phase distortion introduced by a filter could produce a significant impact on data accuracy (examples of phase distortion can be found later in Figs. 3 and 4). As a simple expression, a digital infinite impulse response (IIR) filter has nonlinear-phase response, whereas a digital finite impulse response (FIR) filter can be designed to have a linear-phase characteristic over the frequency range of interest.[5] Regarding other more general comparisons between IIR and FIR filters, please refer to signal processing and filter design textbooks (such as those of Oppenheim and Schafer[5] and Jackson[6]). Also, we will discuss some strengths and weaknesses, especially for the FIR filter, in the individual ECG filtering process sections that follow.

Data sampling, Nyquist frequency, and anti-aliasing

A modern ECG device is basically a digital system. After preliminary processing by the front-end module, the analog ECG signal is immediately converted into a digital form (Analog to Digital Conversion [A/D]) at a particular sampling rate or frequency for further usage.

The Nyquist frequency, also named the folding frequency, is half the sampling rate. Per the sampling theorem, the bandwidth of the input signal should not be greater than the Nyquist frequency. Signal frequencies higher than the Nyquist frequency will encounter a "folding" about the Nyquist frequency and map false components back into lower frequencies. An example is a frequency component at 10 Hz above the Nyquist frequency, which is folded backward to 10 Hz below the Nyquist frequency. This effect is called aliasing.

The Nyquist frequency is a key concept in the initial data sampling (with A/D conversion) and resampling processing (such as further down-sampling).[5] To prevent aliasing interference, the signal must be band-limited. The anti-aliasing process is to use a low-pass filter (LPF) to reject the unwanted frequencies (equal to and greater than the particular Nyquist frequency) of the input signal before sampling or resampling.

Oversampling and down-sampling

Oversampling technique is used in many ECG devices. Oversampling simply refers to an initial A/D conversion sampling rate f_{os}, which is many times higher (eg, 8000 Hz) than the final data resolution target sampling rate f_s (eg, 500