

Filtering Motion Capture Data for Real-Time Applications

Ståle A. Skogstad, Kristian Nymoen,
Mats Hovin and Sverre Holm
University of Oslo, Department of Informatics
{savskogs,krisny,matsh,sverre}@ifi.uio.no

Alexander Refsum Jensenius
University of Oslo, Department of Musicology
a.r.jensenius@imv.uio.no

ABSTRACT

In this paper we present some custom designed filters for real-time motion capture applications. Our target application is *motion controllers*, i.e. systems that interpret hand motion for musical interaction. In earlier research we found effective methods to design nearly optimal filters for real-time applications. However, to be able to design suitable filters for our target application, it is necessary to establish the typical frequency content of the motion capture data we want to filter. This will again allow us to determine a reasonable *cutoff frequency* for the filters. We have therefore conducted an experiment in which we recorded the hand motion of 20 subjects. The frequency spectra of these data together with a method similar to the *residual analysis* method were then used to determine reasonable cutoff frequencies. Based on this experiment, we propose three cutoff frequencies for different scenarios and filtering needs: 5, 10 and 15 Hz, which correspond to *heavy*, *medium* and *light* filtering, respectively. Finally, we propose a range of *real-time* filters applicable to motion controllers. In particular, *low-pass filters* and *low-pass differentiators* of degrees one and two, which in our experience are the most useful filters for our target application.

1. INTRODUCTION

Motion capture (MoCap) and sensor technologies are often used for real-time interactive musical applications, e.g. game controllers like Wii Remote, PlayStation Move, Kinect, and other controllers like mobile phones and novel interfaces for desktop computers. The increased availability of new and improved MoCap technologies together with algorithms that interpret user motion as control data, make it increasingly affordable and feasible to use it for musical interaction. We refer to such interfaces as *motion controllers* (also known as *gesture controllers*) [6]. However, many MoCap and sensor technologies give noisy results, therefore making it necessary to apply noise removal filters [13, 18].

Low latency is a prerequisite for achieving intimate control in musical interactive applications [15]. And, as one might expect, there will always be a corresponding delay penalty when employing a *digital filter*. More specifically, this delay performance is given as the *group delay* and is measured in samples, or sampling periods. This further implies that the given time delay of a filter is proportional to




Figure 1: There is an intrinsic delay penalty when employing digital filters on MoCap data.

the sampling rate of the MoCap system in use [7]. Since most MoCap systems have a relatively low sampling rate, normally between 30 and 200 Hz, this implies that the given group delay of the filter is critical for the total amount of delay. The goal of the current paper has been to develop filters that are optimized for motion controllers and that also minimize the latency they add to the musical applications (Figure 1).

In our previous work we found methods to design nearly optimal *digital filters* with low group delay [11]. However, to be able to design application specific filters, it is necessary to determine the frequency properties of the data to be filtered. We have therefore conducted an experiment to determine these properties for musical application based on *free-hand motion in the air*.

In the next section we give a brief introduction to digital filters. Then, in section 3, we present the experiment and how to determine reasonable frequency properties of human MoCap data. Based on these results, a range of nearly optimal filters for the target application is presented, together with some evaluations in section 4, before the results are discussed in section 5.

2. BACKGROUND - DIGITAL FILTERS

Our main goal when applying *filters* is to smooth data or to restore signals that have been distorted with noise. There exist several methods, and they can roughly be divided into two categories; *curve fitting techniques* and *digital filters*. Curve fitting can intuitively be explained as trying to graphically fit a smooth curve to noisy data. The most common methods are *polynomial fit* and *spline methods* [18]. However, curve fitting noisy MoCap data is known to be suboptimal since human motion does not follow polynomial curves [9]. Digital filters are seen as the most general method for noise smoothing and is the technique we are going to adapt in this paper, since we want a *causal filter* with good real-time properties. Causal here indicates that the filter output depends only on past and present inputs, i.e. a mandatory property for real-time applications.

2.1 The filter objectives

Formally, the goal of a noise filter is to extract the desired signal from some noisy data. Typically this is done by designing a filter, with the purpose of removing the noise com-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME'13, May 27 – 30, 2013, KAIST, Daejeon, Korea.
Copyright remains with the author(s).

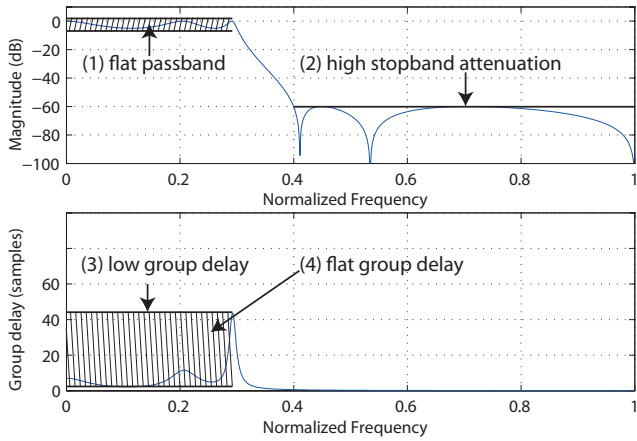


Figure 2: The frequency domain plot of an IIR low-pass filter. The filter objectives are highlighted.

ponent while leaving the desired signal unchanged. In other words, the main two filter objectives are:

- *Maximize noise attenuation.* That is, reduce the amount of noise to maximize the *signal-to-noise ratio* (SNR).
- *Minimize the signal distortion.* That is, avoid altering the desired signal.

There exists much theory regarding the two objectives above [17]. However, in this paper we are especially interested in the following additional objective:

- *Minimize the filter delay.* That is, to minimize the time it takes for the signal to pass the filter.

The most common way to design a digital filter is in the *frequency domain* [17]. Here the aim is to determine the localization of the signal and the noise in the frequency domain, and then designing an appropriate filter based on these properties. The *passband* refers to the frequencies that are passed, i.e. wanted, while the *stopband* refers to the frequencies we want to filter out. This technique works particularly well if the signal and the noise can be effectively separated in the frequency domain. However, this is not necessarily the case for MoCap data. For instance, so-called *white noise* is a common property for sensors [16], and is evenly distributed in the whole frequency band. In other words, not even an ideal low pass filter can suppress all the noise since there will also be noise in the passband [18]. In these cases we need to compromise between noise attenuation and signal distortion. We return to this challenge in section 3.

In Figure 2 we have plotted the frequency properties of a typical *low-pass* filter, which is the type we are going to work with since human motion mainly consists of low frequencies [18]. The figure highlights also the objectives of filter design. Simultaneously, we want: (1) *flat passband*, i.e. low signal distortion, (2) *high stopband attenuation*, i.e. high noise suppression, (3) *low group delay*, i.e. low latency, and (4) *flat group delay*, i.e. that all frequency components of the wanted signal are similarly delayed, also known as linear phase [7]. Let us now consider the different digital filter types.

2.2 Digital filter types (FIR and IIR)

There exist two main digital filter types, *finite impulse response* (FIR) filters and *infinite impulse response* (IIR) filters. *Moving average* is probably the most simple and intuitive realization of a FIR filter [14]. While the moving

average filter have low-pass filter properties, the frequency domain properties are solely specified by its length, i.e. the order of the filter. In most cases there will exist more optimal FIR filter solutions [14], but moving average filters are frequently used because they are intuitive and simple to implement.

IIR filters, as the name suggests, have an infinite impulse response that is the result of their recursive nature. While a FIR filter only bases its output on the input signal, an IIR filter bases its output on former output values as well. In essence, IIR filters offer an effective way of achieving a long *impulse response*, without having to use long FIR filters. Therefore, if the goal is to minimize the group delay, the use of IIR filters seems reasonable, since they can have dramatically lower order than symmetric FIR filters with similar performance [7]. Our results in [11] support this claim as well.

There is one main advantage to so-called *symmetric* FIR filters compared to causal IIR filters, being that they have *linear phase* which implies a constant group delay [17], i.e. all frequencies are delayed by the same amount. Symmetric FIR filters have additionally a fixed group delay of $n/2$ samples where n is the given filter order. In other words, their constant group delay comes at the expense of a fairly high filter delay compared to IIR filters with similar performance [11]. Furthermore, it is not certain that an IIR filter with a moderate amount of *group delay error* is a big concern for our target applications.

2.3 Low-pass differentiators (LPD)

Differentiators are a filter type that are commonly used to extract velocity and acceleration data from position data [13]. When differentiating MoCap data, it is normal to experience an increase of noise in the differentiated data. This is due to the fact that differentiation acts as a high pass filter. Accordingly, the low frequency motion data in the passband will be attenuated while the white noise in the higher frequencies will be amplified. As a result, we end up with a lower SNR value for the differentiated data, which increases the need for filtering [18, 2]. This is why it is reasonable to use so-called *low-pass differentiators*, since they avoid the undesirable amplification of noise in the higher frequency band. They also provide better total filter solutions than to use a low-pass filter in cascade with a differentiator operator, as we have shown in [10]. Similarly, it is better to use one low-pass differentiator of degree two, than to use two of degree one in cascade

2.4 Filter design methods

The design of symmetric FIR filters is a linear problem and there exist different general solutions for most FIR design problems, e.g. the *least square method* and the *Parks-McClellan* method [8, 4]. The design of IIR filters is, on the other hand, a nonlinear problem, and there are no general optimal design methods. There are however different construction methods, which can give optimal solutions for some special cases. The most known classical IIR filter methods are Butterworth, Chebyshev and elliptical (Cauer) [17]. They are very useful for standard filter types as long as there is little restriction on the group delay responses [5, 11]. It is therefore necessary to use alternative design methods if we need more control over the group delay specifications. In our earlier research we presented a successful method for designing nearly optimal IIR filters with arbitrary specifications, including low-pass filters with minimal group delay [11] and IIR low-pass differentiators [10]. In that work we regarded filter design as a multi-objective optimization problem, which was solved using an unbiased metaheuristic

search algorithm. Using this method we are able to custom design nearly optimal IIR filters with the desired trade-off between group delay and the other filter objectives given above. For more details about this method see [10] and [11]. However, before we can design filters for our applications, we need to determine the typical frequency properties of the MoCap data we want to filter.

3. FREQUENCY PROPERTIES OF MOTION

As we show below, it is possible to determine reasonable cutoff frequencies from recorded MoCap data. The best method would be to determine the cutoff frequency before filtering a given set of data. However, this is impossible for real-time applications since the cutoff frequency needs to be specified beforehand. In practice, we are forced to use predetermined filters, and therefore need to estimate generic frequency properties for free-hand motion. Let us start by presenting our analysis methods before we continue with presenting the experiment in section 3.2.

3.1 Analysis methods

Before we can begin the discussion on how to estimate a reasonable generic cutoff frequency, we need to make some assumptions about the noise distribution of the relevant MoCap technologies. There can be many sources of noise in a MoCap system: it can be sensor noise, wobbling markers, electrical interference, quantization noise and more, dependent on the MoCap system used [19]. As already mentioned, sensors are known to have white noise properties [16, 19]. Some MoCap technologies may have a different noise distribution. However, for simplicity, in this paper we assume that the MoCap system has a white noise distribution. Consequently, our goal is to attenuate as much as possible of the frequency band that is not part of the signal band. If it is mandatory not to distort signal, we need to choose a cutoff frequency that is just outside the signal band. However, if we need higher noise suppression than is possible with this conservative choice, we need to compromise signal distortion by lowering the cutoff frequency inside the signal band [18]. The determination of the optimal cutoff frequency will then be based on the noise attenuation needed and how much we can lower the cutoff frequency inside the signal band without distorting the desired signals too much. To be able to determine the latter, we used the following two methods.

3.1.1 Power spectral density (PSD) estimation

The most common method to determine the frequency content of a digital signal is to analyze the *frequency spectrum*, which can be derived in different ways with the Fourier transform. A non smoothed spectrum estimation with the *Periodogram*, a classic non-parametric technique, will normally be too noisy to clearly show the trend in the data [3]. We therefore ended up using the *Welch's method* with a *Hann* window of length 100 (sampling frequency of 100 Hz). This is a much used method which reduces the noise in the spectral density estimation in exchange for reduced resolution in the frequency domain. However, other spectrum estimators and windows will give similar results [3].

3.1.2 Residual analysis

While the above mentioned method offers a good basis for making a conservative determination of the passband edge, it does not necessarily provide us with a good basis to determine a reasonable cutoff frequency. For a more hands on approach, it is possible to visually inspect the MoCap data when filtered with different cutoff frequencies. We can then choose the cutoff that provides a good balance between noise reduction and signal distortion. A more systematic version

of this technique is known as *residual analysis*, which is a common method used for this task in the field of biomechanics [18]. The method consists of low-pass filtering the data with different cutoff frequencies and calculating the *residual*, i.e. what is left over when we subtract the filtered data from the raw data. As long as the filter is only attenuating noise, the residual should be rather small. However, when the filter starts to attenuate the desired signal, the residual will become larger. By performing this analysis for several cutoff frequencies, and plotting the resulting residuals, we get an overall picture of their impact. This plot can then serve as the basis for determining a reasonable cutoff frequency [18].

When computing the residual plots, care should be taken to make sure that the applied filters have constant group delay and are consistent with each other. This will ensure that the change in residual is not due to difference in the filter characteristics other than the cutoff frequency. It is common to use the actual intended filters which are supposed to be used in the final application [18]. However, our goal is not to find the optimal filter for a given set of data, but to find the main frequency trend of free-hand motion among several recordings. We ended up using the *window method* [7] to design the needed filters with an order of 200. This symmetric FIR design method has a broad cutoff frequency range and gives consistent filter characteristics for different cut-off frequencies [1].

3.2 The experiment

3.2.1 Setup and recordings

The experiment consisted of recording the hand motion of 20 subjects, 4 females and 16 males in the age range of 22-47. We used an optical infrared marker based MoCap system, OptiTrack, to record the subjects's hand motion at 100 Hz. The MoCap setup consisted of eight OptiTrack *V100:R2* cameras that were attached to tripods in a room measuring about 7x8 meters. One 16 mm reflective spherical marker was attached to the subject's dominant hand, close to the index finger, see Figure 3. Care was taken to minimize wobbling of the marker, which can introduce additional noise to the MoCap data. For the same reason, we also spent time calibrating the OptiTrack system. We did not want to perform post processing of the recorded data, e.g. for gap filling, which could potentially have distorted our results. Recordings with invalid or missing data were therefore omitted. The subject's hand motion were further recorded in the following two takes, both 20 seconds long.

- *Take 1:* The subjects were asked to move their dominant hand as rapid as possible in an arbitrary pattern. The intention of these recordings was to find an upper frequency limit for hand motion.
- *Take 2:* The subjects were asked to simulate that they were controlling some application with more *articulated* and *controlled* motion. Here we wanted to ex-




Figure 3: Placement of the marker (left) and an illustration of the experiment (right).

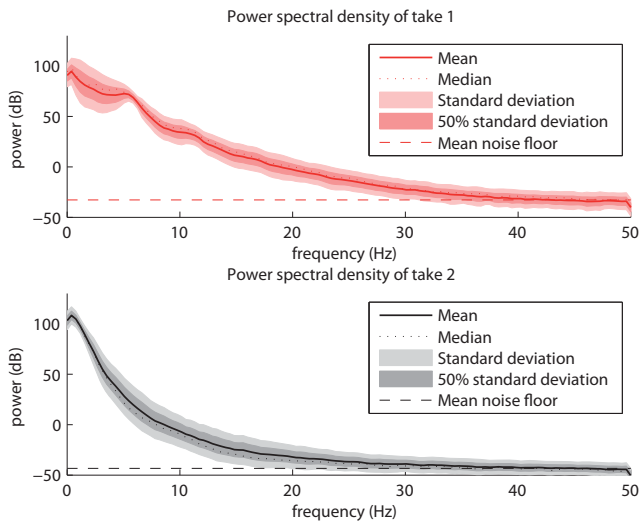


Figure 4: PSD estimation of the recorded data using Welch's method. The data is shown as statistical results of all 20 recordings, with results from both Take 1 (rapid) and Take 2.

amine the typical frequency content of the motion we anticipate to see most of in our target application.

We expected the latter to result in the need for a lower cut-off frequency than the former, which makes it possible to remove more noise. During all recordings, the subjects were asked to not clap their hands or make other limb collisions. We wanted to avoid collisions since they can be problematic to study, e.g. contain high frequency components that require higher sampling rates, and added noise problems with wobbling markers.

3.2.2 Results and interpretations

The results of the experiment are shown in Figures 4 and 5. As we can see from the spectral density estimates of Take 2, the mean value starts to move away from the noise floor between 20 and 30 Hz. For Take 1, the mean value starts to move away between 25 and 35 Hz. Furthermore, the main frequency content for Take 2 reaches roughly up to about 5–10 Hz, while Take 1 has a wider frequency distribution.

The residual plots in Figure 5 are somewhat easier to interpret since deviation in mm is more comprehensible than power in dB. When filtering hand motion, which normally has a displacement in the range of 200–1000 mm, a deviation of 1 mm is normally not significant. We have further seen a general trend for what the residual values indicates. When it was below 1 mm, the filters did not severely distort the MoCap data. But when the value increased above 5 - 10 mm, the filters started to clearly distort some high frequency parts of the MoCap data.

By using the above indicators and the statistical residual results in Figure 5, it seems reasonable to set the lower cutoff frequency for Take 2 to about 5 Hz, since the standard deviation is below 5 mm at this cutoff value. A reasonable upper frequency cutoff for Take 1, can further be set to be between 15 and 20 Hz, since the mean value goes below 1 mm in this region. A sensible trade off between these two outer cutoffs is in our opinion 10 Hz, since Take 2 is below 1 mm and Take 1 is below 5 mm for this cutoff value. Examples of how these cutoff frequencies perform can be seen in Figure 6. Based on this experiment, we propose the following three frequency cutoffs for filtering free-hand motion:

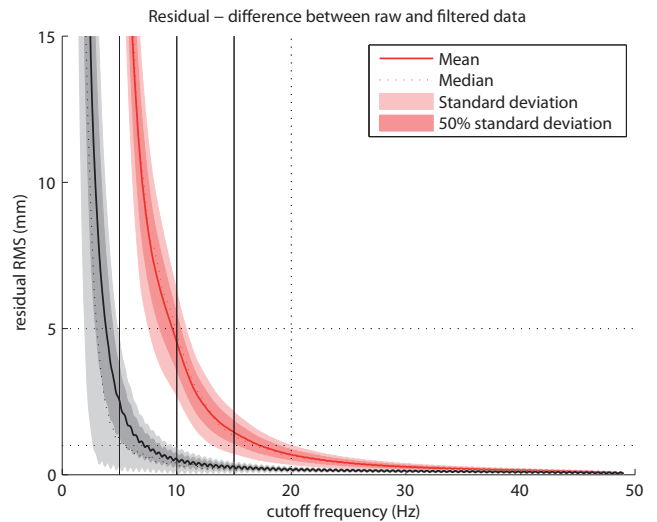


Figure 5: Statistical results of the residual analysis of the recorded data. Take 1 (rapid motion) is in red while take 2 is given in black.

5 Hz *Heavy filtering*: Fast and rapid motion may be heavily smoothed out. However, the filtered data will contain the main features of normal controlled hand motion.

10 Hz *Medium filtering*: Most features of normal and medium rapid motion will be kept in the filtered data. However, some of the higher frequencies will be partially distorted.

15 Hz *Light filtering*: All main features of both rapid and normal motion are kept. Only the most extreme parts of the data may be partially blurred.

We could have added a cutoff frequency at 20 Hz, since the residual plot shows that the mean value of Take 1 decreases below 1 mm at about 20 Hz. But we have omitted this cutoff since we are not sure if the content that is blurred away with the 15 Hz cutoff, is due to noise or actual motion. The residual difference with the 20 Hz cutoff, is also minimal. However, a cutoff frequency of 20 Hz can be used if it is

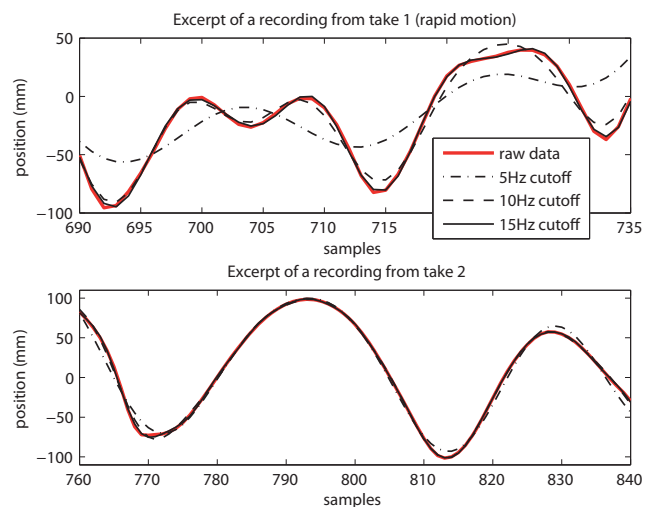


Figure 6: Excerpts from Take 1 and Take 2. While a 5 Hz filter cutoff works well for the Take 2 below, the rapid motion needs a 10 Hz or a 15 Hz cutoff frequency to follow the details in the recording.

important to keep all details in the recordings, and noise suppression is secondary.

3.2.3 Discussion

With this experiment we wanted to determine a generic trend in frequency content of free-hand motion. However, it was not straightforward to give instructions to the subjects. We hesitated to give them specific tasks, since this could lead them to do certain motion which could have influenced our results. We therefore ended up giving them quite general and open tasks, which resulted in a range of different interpretations and motion. However, as the results show, there is a quite clear trend among the recordings.

We considered testing expert subjects trained in moving at high frequencies, e.g. drummers. However, their motion is normally an effect of collisions and special techniques to be able to achieve high frequency. These motion were not part of our scope. Furthermore, inspection of the recorded data revealed that some contained *position jumps* that could not have been due to human motion. The errors clearly distorted the PSD data and raised the overall noise floor. It is therefore important to remove these errors if one wants valid PSD data. However, these errors had minimal impact on the residual plots, which shows that the residual method is a somewhat more robust analysis method.

4. PROPOSED IIR FILTERS

In our previous work we have based our *sound excitation* on three main types of MoCap data: *position*, *velocity* and *acceleration* [12]. We found these motion features to be the most useful for controlling sonic and musical features. We have therefore chosen to focus on the filter types that extracts these motion features from raw positional MoCap data, respectively *low-pass filters* and *low-pass differentiators* of degree 1 and 2.

4.1 Proposed IIR vs. symmetric FIR filters

We have already shown in our previous work that our IIR design method can produce better low delay filters than currently available methods [10, 11]. As we can see from Table 1 and Figure 8, the proposed IIR filters are significantly better than symmetric FIR filters if low delay and high noise attenuation are of priority, giving a potential noise suppression gain between 5-16 dB for the relevant filter types. The presented IIR filters have a group delay of 2 samples or less. This group delay amount was found to give a well balanced trade-off between the different filter objectives. For a more thorough low-delay comparison between different filter types, see [11]. The specification of the proposed IIR filters is given on our project web page together with a MAX/MSP implementation [1], and a subset of these filters is given in Table 2. (To convert *normalized frequency* to hertz, multiply by half the sample frequency.)

4.2 Filter evaluation

We have tested the proposed IIR filters and confirmed their performance in MAX/MSP. It is not trivial to evaluate the filters for general NIME use as it depends strongly on the end application. While some applications may want to minimize noise to get the most robust performance, some applications may benefit artistically from MoCap noise as it can add a desirable texture to the resulting sound synthesis. Over-smoothing, i.e. deliberately distorting the signal, can also be appropriate for some applications. However, it is important to use a cutoff frequency that satisfies the need for the given task, as the following example shows. By identifying high peaks in the acceleration data, we are able to detect




Figure 7: The effects of using different *cutoff frequencies* when extracting acceleration of a *hand clap*. The collision is more easily detected if the *cutoff frequency* is above 10 Hz (acceleration in m/s^2).

sudden motion and limb collisions, which we have used to trigger sonic and musical features [12]. The effect of using a too low cutoff frequency when extracting the acceleration data is shown in Figure 7. Not only does it attenuate more of the white noise, it also attenuates the acceleration peak. This is an expected effect, since a collision can be seen as an impulse which has a flat frequency response, i.e. the energy is spread out in the whole frequency band. The more of the frequency band that is included when differentiating, the more the collision power will be seen in the acceleration data.

Another important issue is what impact a moderate amount of *group delay error* can have on our target application. In our experience, there does not appear to be any dramatic negative distortion effect if the upper frequency range has some group delay error, as long as the main content (up to 5–10 Hz) has a fairly constant group delay. The optimized IIR filters are further superior if high noise attenuation, combined with low passband distortion and low group delay are desired. In our findings, it is possible achieve up to one-third the delay by using optimized IIR filters, as compared to symmetric FIR filters with similar performance. A delay of two samples, as opposed to six, yields a delay reduction of 40 ms for a MoCap system with a sampling frequency of 100 Hz, which should be a favorable reduction for a typical MoCap setup used for musical interaction [13]. In short, the optimized IIR filters have much better *low delay potential* than symmetric FIR filters for our target application, at the expense of a more complicated design.

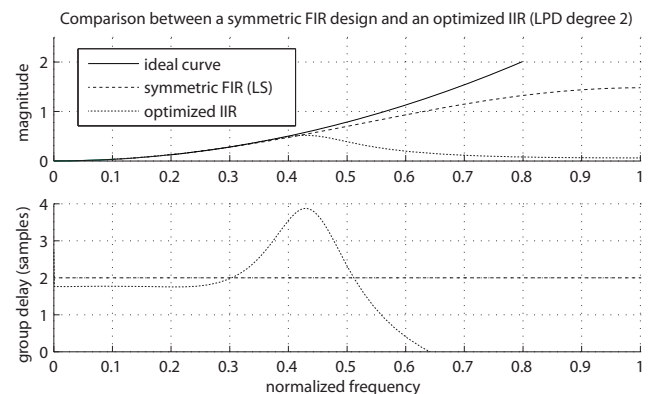


Figure 8: Comparison between 4th order low-pass differentiators (LPD) of degree 2 with a normalized cutoff frequency of 0.3. If low passband distortion is desired, the optimized IIR differentiator of degree 2, gives a noise suppression improvement of about 13 dB (~ 4.5 times more noise attenuation) with similar or better performance for the other filter objectives.

Table 1: Potential noise attenuation gain in dB of the proposed IIR filters compared to optimal symmetric FIR designs, all of order 4. While the symmetric FIR filters have a fixed group delay of 2 samples, the proposed IIR filters have a group delay of 2 samples or less. For some of the proposed filters we have tolerated a moderate amount of group delay error.

normalized cutoff	0.1	0.2	0.3	0.4	0.5
low-pass filters	8	8	8	6	5
low-pass diff. of degree 1	10	10	9	7	6
low-pass diff. of degree 2	16	15	13	12	10

5. DISCUSSION AND CONCLUSION

In this paper we have addressed the challenge of using digital filters for real-time applications, focusing on filtering free-hand motion. To be able to design filters for such motion data, we conducted an experiment to determine the generic frequency properties of free-hand motion. Based on this experiment, we propose 3 different filter cutoffs; 5, 10 and 15 Hz. The 5 Hz, and partly the 10 Hz, cutoff will attenuate some of the high frequency parts of rapid free-hand motion. However, this may be necessary to get the needed noise suppression.

Although the experiment has only considered the frequency content of free-hand motion, our review of previous frequency studies in biomechanics suggests that most human motion is reported to be close to our found cutoff values, or more specifically between 3-26 Hz [9, 19, 20]. Our proposed frequency cutoffs should therefore work for most parts of the body, with some reasonable generalizations and adjustments, by regarding the kinematics of the used limb. Our proposed analysis method can be used if more certain knowledge is needed [1].

Finally, we propose a set of filters for our target applications, which has lower delay than what is achievable by established filter design methods. The main purpose of these filters has been to present some IIR filters designed with low group delay in mind, which is an important feature for intimate control for musical interactions. Compared to optimal symmetric FIR filters, they give a noise attenuation increase between 5-16 dB with similar delay, or up to 2-3 times the delay reduction for similar magnitude properties. These filters and some tools are published on our project page together with a Max/MSP implementation [1]. Since the optimal filter depends heavily on application specific details (e.g. sampling frequency, intended use), it is not possible to present a complete list of filters for all different applications and scenarios. However, our proposed set of filters should demonstrate the potential of using our filter design approach.

6. REFERENCES

- [1] Project web page: fourms.uio.no/projects/sma/subprojects/mocapfilters.
- [2] G. Giakas and V. Baltzopoulos. Optimal digital filtering requires a different cut-off frequency strategy for the determination of the higher derivatives. *Journal of biomechanics*, 30(8):851–855, 1997.
- [3] M. Hayes. *Statistical digital signal processing and modeling*. John Wiley & Sons, 1996.
- [4] X. Lai. Optimal Design of Nonlinear-Phase FIR Filters With Prescribed Phase Error. *Signal Processing, IEEE Trans.*, (9):3399–3410, sept. 2009.
- [5] M. Lang. Least-squares design of IIR filters with

Table 2: Transfer functions for a subset of the proposed IIR filters with a maximum group delay of 2 samples and a normalized frequency cutoff of 0.2. The table presents a low-pass filter (LF) and low-pass differentiators of degree 1 and 2 (LPD1 and LPD2). More filters are presented on our project web page [1].

type	b_1 a_1	b_2 a_2	b_3 b_3	b_4 b_4	b_5 b_5
LF	0.1227	-0.064575	0.044457	0.01949	0.019725
	1	-2.4965	2.8553	-1.5848	0.36631
LPD1	0.21077	-0.171566	-0.0552011	0.0182798	-0.00228255
	1	-1.71529	1.48777	-0.658224	0.118839
LPD2	-0.0797369	0.117985	-0.0144855	-0.00603732	-0.0177256
	1	-1.3405	1.19009	-0.531258	0.0931822

prescribed magnitude and phase responses and a pole radius constraint. *Signal Processing, IEEE Transactions on*, 48(11):3109–3121, nov 2000.

- [6] E. R. Miranda and M. Wanderley. *New Digital Musical Instruments: Control And Interaction Beyond the Keyboard*. A-R Editions, Inc., 2006.
- [7] S. Mitra. *Digital signal processing: a computer based approach*. McGraw-Hill H. E., 2005.
- [8] T. W. Parks and C. Burrus. *Digital filter design. Topics in digital signal processing*. Wiley, 1987.
- [9] D. Robertson. *Research Methods in Biomechanics*. Human Kinetics, 2004.
- [10] S. A. Skogstad, S. Holm, and M. Hovin. Designing Digital IIR Low-Pass Differentiators With Multi-Objective Optimization. pages 10 – 15. ICSP, IEEE Computer Society, 2012.
- [11] S. A. Skogstad, S. Holm, and M. Hovin. Digital IIR Filters With Minimal Group Delay for Real-Time Applications. pages 1 – 6. ICET, IEEE Computer Society, 2012.
- [12] S. A. Skogstad, K. Nymoen, Y. de Quay, and A. R. Jensenius. Developing the Dance Jockey system for musical interaction with the Xsens MVN suit. In *Proc. of NIME*, pages 226 – 229, Ann Arbor, Michigan, 2012.
- [13] S. A. Skogstad, K. Nymoen, and M. Hovin. Comparing inertial and optical mocap technologies for synthesis control. In *Proc. Sound and Music Computing*, pages 421 – 426, Padova, 2011.
- [14] S. Smith. *The Scientist and Engineer’s Guide to Digital Signal Processing*. California Tech. Pub., 1997.
- [15] D. Wessel and M. Wright. Problems and prospects for intimate musical control of computers. In *Proc. NIME*, Seattle, USA, 2001.
- [16] J. S. Wilson. *Sensor technology handbook*, volume 1. Newnes, 2005.
- [17] S. Winder. *Analog and digital filter design*. EDN series for design engineers. Newnes, 2002.
- [18] D. Winter. *Biomechanics and Motor Control of Human Movement*. John Wiley & Sons, 2009.
- [19] G. Wood. Data smoothing and differentiation procedures in biomechanics. *Exercise and sport sciences reviews*, 10(1):308, 1982.
- [20] J. Wosk and A. Voloshin. Wave attenuation in skeletons of young healthy persons. *Journal of Biomechanics*, 14(4):261–263, 1981.