

POLYPHASE FIR COMPLEMENTARY FILTERS FOR CONTROL SYSTEMS

Wensheng Hua, Dan B. DeBra, Corwin T. Hardham, Brian T. Lantz
Stanford University, Stanford, CA

Joseph A. Giaime

Louisiana State University, Department of Physics and Astronomy, Baton Rouge, LA

1 INTRODUCTION

A pair of filters are called complementary filters if their transfer functions sum to one at all frequencies in a complex sense, i.e. the phase is zero and the magnitude is one. Often, it is desirable to have one of the complementary filters be a high-pass filter and the other be a low-pass filter. A computational efficient way of implementing high performance complementary filters will be discussed in this paper.

In dynamic control systems, different sensors are often used to measure the same physical variable in different frequency bands. For example, in a navigation system, a low frequency GPS sensor can be used together with a high frequency inertial sensor to measure the position of vehicle. In that case, complementary filters can be used to combine the signals from those two sensors: filter the signal from the GPS sensor by the low-pass filter and filter the signal from the inertial sensor by the high-pass filter, and then sum the two filter outputs together to generated a “super sensor” signal. Ideally, the super sensor will have high sensitivity at all frequencies. This technique is called sensor blending.

A similar technique for combining sensors is called sensor correction. For example, in an active vibration isolation system, a relative position sensor can be used to measure the platform’s position with respect to the ground. From a different point of view, one can say that the relative position sensor is sensitive to both the platform motion and the ground motion. If an inertial sensor is set up on the ground to measure the ground motion, its signal can be used to cancel out the ground sensitivity of the position sensor. To do so, one has to use a sensor correction filter to match the gains of the inertial sensor and the position sensor in complex sense, i.e. both magnitude and phase. The sensor correction filter also has to take care of noise. For example, in general, the

inertial sensors are noisy at low frequencies. Hence, it is desirable to process the ground inertial signal with a high-pass sensor correction filter which can filter out the low frequency noises. In that case, the high-pass filter in the complementary filter pair can be used as the sensor correction filter. Details of complementary filter design for both sensor blending and sensor correction filters are going to be discussed in the following sections.

It is worth mentioning that a similar kind of filter pair, so-called “magnitude complementary filter pair”, is widely used in communication systems [1]. The difference between those filters and the filters discussed in this paper is that the sum of the transfer functions of a magnitude complementary filter pair does not have to have zero phase. In many communication systems, a little bit of time delay will not damage the system’s performance much, so some phase lag is allowed in the filters. However, the phase lag could be disastrous for dynamic control systems. Therefore, only strict complementary filters are studied in this paper.

2 FIR COMPLEMENTARY FILTER DESIGN

For simplicity, all sensors are assumed to have a nominal transfer function of one in this paper.

2.1 SENSOR BLENDING FILTERS

Robust sensor blending filters in feedback control systems should have following properties:

1. The filters’ transfer functions should be close to zero in stop band;
2. The filters’ transfer functions should be close to one in pass band;

3. The filters should be robust to sensor gain change, i.e., if sensor gains are changed within certain limits, the control system would not become unstable. This implies that the magnitude of the filters' transfer functions should be limited at all frequencies.

2.2 SENSOR CORRECTION FILTERS

The design of our optimal sensor correction filter is a tradeoff between gain match error and noise error. For example, if the filter gain is one, there is no gain match error but all the noise will go through the filter and generate a large noise transmission. On the other hand, if the filter gain is zero, no noise can go through the filter but the gain match error is huge. The sensor filter's noise error is determined by the magnitude of its transfer function, while the filter's gain match error is determined by the magnitude of the complement's transfer function. Hence, the sensor correction filter design problem is also a complementary filter design problem, even though only one filter is actually implemented. For this reason, it is not surprising that the optimal sensor correction filters have similar properties as the sensor blending filters:

1. The filter's transfer function should be close to zero in stop band;
2. The filter's transfer function should be close to one in pass band;
3. The magnitude of the filter's transfer functions should be limited in all frequencies.

2.3 COMPLEMENTARY FILTER DESIGN

Both sensor correction filters and sensor blending filters can be design by solving the following problem: find a pair of complementary filters that satisfy:

$$L(\omega) + H(\omega) = 1; \quad (1)$$

$$|L(\omega)| \leq \bar{L}(\omega); \quad (2)$$

$$|H(\omega)| \leq \bar{H}(\omega); \quad (3)$$

Here, ω denotes frequency. Complex functions $L(\omega)$ and $H(\omega)$ are used to denote the transfer functions of the low-pass filter and high-pass filter respectively. $\bar{L}(\omega)$ and $\bar{H}(\omega)$ are known real functions of frequency. In general, the desired complementary filters are designed by choosing appropriate values of $\bar{L}(\omega)$ and $\bar{H}(\omega)$ and then solving the problem given by equations 1 through 3.

2.4 FIR COMPLEMENTARY FILTER

For a finite impulse response (FIR) filter, the transfer function, $G(\omega)$, is the Fourier transform of its filter coef-

ficients in the time domain, $g(n)$:

$$G(\omega) = \sum_n g(n)e^{-i2n\pi\omega}. \quad (4)$$

Because $G(\omega)$ is a linear function of $g(n)$ for FIR filters, the problem given by equations 1 through 3 is a convex optimization problem which can be solved very efficiently using well developed convex optimization tools [2]. For example, a 512th order FIR filter can be designed using the SeDuMi [3] convex optimization tool. Its frequency and time domain responses are shown in Figures 1 and 2. The length of the filter's impulse response is about 300 seconds. The filters have very sharp frequency discrimination. For example, the high-pass filter can attenuate the signal proportional to frequency's fourth power in its stop band below 0.04 Hz, while its transfer function is almost exactly one in its pass band above 0.1 Hz.

This filter was designed for the horizontal ground mounted sensor correction seismometers in our isolation system, which need good gain-matching performance above 0.1 Hz, but are plagued by tilt noise below 0.04 Hz.




Figure 1: Frequency response of the FIR filter. This is a 512th order FIR filter with sampling frequency of 1.67 Hz.




Figure 2: Time impulse response of the high pass FIR filter. This is a 512th order FIR filter with a sampling frequency of 1.67 Hz. The total impulse response time is about 300 seconds.

2.5 IIR COMPLEMENTARY FILTER

For an infinite impulse response (IIR) filter, the transfer function, $K(\omega)$, can also be easily obtained from its time domain numerator coefficients $k_N(n)$ and denominator coefficients $k_D(n)$:

$$K(\omega) = \frac{\sum k_N(n)e^{-in2\pi\omega}}{\sum k_D(n)e^{-in2\pi\omega}}. \quad (5)$$

However, finding $k_N(n)$ and $k_D(n)$ that satisfy equations 1 through 3 is not a convex problem, and hence very difficult to solve.

The difference between FIR and IIR complementary filters is summarized in the following table.

	FIR	IIR
design	convex, easy	non-convex, difficult
implementation	complex	simple
computation	much	little
performance	great	ok

3 POLYPHASE FIR FILTER IMPLEMENTATION

For some precision applications the performance of the complementary filter exceeds the accuracy that can be achieved with simple IIR filter design tools. We have found that FIR filters can provide the necessary performance, but the complexity leads to computational loads that are impractical.

When implementing FIR filters in the traditional way, the amount of calculations required per second is given

by

$$\mathcal{N} = TF_s^2, \quad (6)$$

where T is the filter's response time and F_s is the sampling frequency. In order to obtain the desired sharp transition between stop-band and pass-band, the filter has to have high frequency resolution, which implies a long time response. However, the signal sampling frequency cannot be too low, otherwise the anti-aliasing and interpolation filters used before and after the FIR filter will introduce enough phase error to destroy the overall performance of the complementary filters. For these reasons, the calculation speed required by the traditional high performance FIR complementary filter is often too high to be practical.

We have developed a technique called polyphase FIR filtering, as shown in Figure 3, to reduce the the calculation load. The incoming data stream is divided into m sub-streams by a cyclic distributor according to the data samples' arrival time. Hence, each sub-stream has the sample rate of F_s/m (in Figure 3, $m = 3$). The FIR filters on each sub-stream are identical traditional FIR filters which work at frequency $F_f = F_s/m$. At each time step, the cyclic collector picks up an output sample from the FIR filter that is working in that phase and sends it to the output. Hence, the output sample rate is F_s , which is the same as the input sample rate.




Figure 3: Polyphase FIR filter implementation. The three FIR filters have identical filter coefficients and each operates at frequency F_f , but are offset in phase. The cyclic distributor distributes the incoming data samples to different FIR filters according to the samples' arriving time phase. At each time step, the cyclic collector picks up an output sample from the FIR filter with the appropriate phase and sends it to the output. Hence, the overall sampling frequency is $F_s = m \cdot F_f$, where m is the number of cyclic phases.

The total amount of calculation that is required by the polyphase FIR filter is given by the total of m parallel FIR filters:

$$\mathcal{N}_p = mTF_f^2 = TF_sF_f = \frac{\mathcal{N}}{m}, \quad (7)$$

which is m times smaller than what the traditional FIR filter requires.

The output of the polyphase FIR filter is given by

$$y(n) = \sum_{i=0}^{N_f} x(n - im)f(i), \quad (8)$$

where $x(n)$ denotes the n th input signal, $y(n)$ denotes the n th output signal, N_f is the number of coefficients of each FIR filter and $f(i)$ denotes the i th filter coefficient of the FIR filter.

If the transfer function of each FIR filter is $F(\omega)$, $\omega \in [0, F_f]$, the transfer function of the polyphase FIR filter is given by:

$$P(\omega) = F(\omega - kF_f), \omega \in [0, F_s], \quad (9)$$

where $k \in I$, and $k = \text{floor}(\frac{\omega}{F_f})$.

At low frequencies, when $\omega \in [0, F_f]$, $k = 0$ and $P(\omega)$ is exactly the same as $F(\omega)$. However, it is obvious that $P(\omega)$ is not perfect at higher frequencies. In fact, $P(\omega)$ is just a periodic repetition of the function of $F(\omega)$ every F_f .

Figure 4 shows the transfer function of the polyphase high-pass filter based on the high-pass FIR filter shown in Figure 1 and 2. $P(\omega)$ is the same as $F(\omega)$ at low frequencies. However, $P(\omega)$ is not unity as desired for $\omega > F_f$.




Figure 4: Transfer function of the high pass polyphase FIR filter based on the high-pass FIR filter shown in Figure 1 and 2.

To address this problem for the high-pass filter, another pair of complementary filters, L_a and H_a , are added to the polyphase high-pass FIR filter, as shown in Figure 5.




Figure 5: Complemented polyphase FIR filter. Complementary filter pair H_a and L_a are used to work with the original polyphase FIR filter to correct its high frequency transfer function.

The overall transfer function of the final filter is given by

$$H_{cp}(\omega) = H_p(\omega)L_a(\omega) + H_a(\omega) \quad (10)$$

$$= H_p(\omega) + (1 - H_p(\omega))H_a(\omega) \quad (11)$$

$$= 1 - (1 - H_p(\omega))L_a(\omega). \quad (12)$$

$H_a(\omega)$ is selected such that

$$|(1 - H_p(\omega))H_a(\omega)| \ll |H_p(\omega)| \quad (13)$$

in H_p 's stop band. According to Equation 11,

$$|H_{cp}(\omega)| \cong |H_p(\omega)|. \quad (14)$$

So, the new filter and the original polyphase FIR filter have similar low frequency properties.

Similarly, choose L_a so that

$$|(1 - H_p(\omega))L_a(\omega)| \ll 1 \quad (15)$$

above the original FIR filter's pass-band. To realize this, L_a might have notches at frequency $F_f, 2F_f, \dots$ so that the peaks of H_p are compensated. This way, the overall filter transfer function is close to one in all the frequencies above the original FIR filter's pass-band. As shown in Figure 6, the high-pass filter transfer function is very close to unity in its pass-band above 0.1 Hz. Finally, we get the high-pass filter that we wanted!

The low-pass filter of the complementary filter can be simply implemented as $L_{cp} = 1 - H_{cp}$.




Figure 6: Filter performance.

There is an intuitive way of thinking about this hybrid polyphase FIR filter implementation. The original polyphase FIR filter behaves very well at low frequencies, but not so well at high frequencies; on the other hand, a simple transfer function of unity behaves very well as a high-pass filter at high frequencies, but not so well at low frequencies. So, it makes sense to use a complementary filter pair, L_a and H_a , to combine these two transfer functions and form a “super” transfer function. For this reason, the filter shown in Figure 5 is called complemented polyphase FIR filter. There are a few remarks that we would like to make about this complemented polyphase FIR filter implementation.

1. The added complementary filter pair does not have to have as sharp a stop-band to pass-band transition as the original FIR complementary pair does. Hence, they are easier to implement, so they can be simple IIR filters. Of course FIR filters work just as well.
2. The added low-pass filter can be implemented as all or part of the decimation and interpolation filters to save calculation power.
3. In multi-rate FIR filter systems [1], there is a multi-phase decimation FIR filter implementation, which combines the decimation process and the anti-aliasing filtering process together to save calculation loads. There are two differences between that multi-phase FIR filter and the polyphase filter in this paper. First, the parallel FIR filters are not identical to each other in the multi-phase filters. Second, the multi-phase FIR filter sums all the FIR filter outputs together, rather than using a cyclic collector. Interestingly, there is a duality relationship between these two kinds of filters: the multi-phase decimation filter decimates the signal,

while the polyphase filter in this paper decimates the filter coefficients.

4 SENSOR CORRECTION USED FOR THE LIGO ACTIVE VIBRATION ISOLATION AND ALIGNMENT SYSTEM

The Laser Interferometer Gravitational-wave Observatory (LIGO) is designed to measure gravitational-waves from astrophysical sources by observing a spatial strain change between suspended test masses using a laser interferometer [4]. At low frequencies, seismic vibration is one of the major noise sources for LIGO. An active vibration isolation system is proposed to reduce the platform vibration for Advanced LIGO [5]. Particularly, the ground seismic motion has a peak, called the micro-seismic peak, at about 0.15 Hz, whose amplitude spectral density (ASD) is about $10^{-6} \text{meter}/\sqrt{\text{Hz}}$ at the LIGO experiment sites. For Advanced LIGO to reach its design sensitivity, the isolation system needs to reduce the ASD motion of the platform to $2 \cdot 10^{-7} \text{meter}/\sqrt{\text{Hz}}$ at 0.15 Hz [6].




Figure 7: Picture of the active vibration isolation and alignment system prototype for Advanced LIGO.

The micro-seismic peak isolation requirement is achieved by an active vibration isolation and alignment system, a prototype of which is currently at Stanford. As

shown in Figure 7, the prototype is a double-stage active isolation system. Capacitive position sensors, L-4C geophones and Streckeisen STS-2 seismometers are used to measure the motion of the stages and the ground. Voice-coils are used as actuators. A dSPACE realtime computer is used to implement the control algorithms whose sampling frequency is 2.5 kHz.

For frequencies below 1 Hz, the system’s control algorithm is shown in Figure 8. A high gain feedback control loop (the loop gain is more than 1000 below 1 Hz) is closed based on the signal from the position sensor. A Streck-eisen STS-2 seismometer is used to measure the ground motion with respect to inertial space to provide the sensor correction signal. The basic idea of the control design is very simple: use the high-gain feedback loop to accurately control the motion of the platform and then command the platform to move just opposite to the ground motion by following the sensor correction signal.

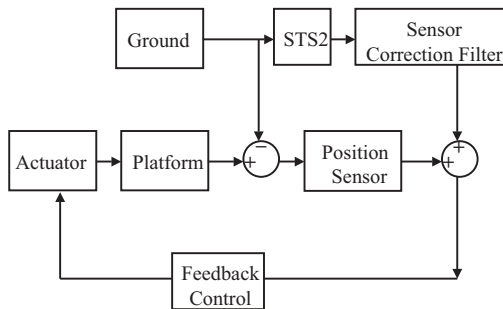


Figure 8: The block diagram of the low frequency control algorithm. The feedback control has an open loop gain of more than 1000 at frequencies below 1 Hz. STS-2 seismometer is used to measure the ground motion. Active vibration isolation is achieved by cancelling out the position sensor’s ground motion sensitivity using the sensor correction signal.

The sensor correction filter design is critical. In the horizontal direction, the noise level of the STS-2 seismometer grows very rapidly at low frequencies: its ASD value in units of $meter/\sqrt{Hz}$ grows proportionally to inverse frequency cubed below 0.04 Hz. On the other hand, to isolate the micro-seismic peak at 0.15 Hz, the platform should start to have isolation above 0.1 Hz. We have found that the FIR filter shown in Figure 6 can provide the necessary performance. It provides a potential isolation factor of 20 above 0.1 Hz and dramatically reduces the low frequency noise below 0.04 Hz with a

transfer function that is proportional to frequency to the fourth power. With this sensor correction filter, the vibration isolation system’s performance is shown in Figure 9. An isolation factor of 10 is achieved in the horizontal direction without introducing unnecessary low frequency noises. The noise level of the platform reaches $2 \cdot 10^{-7} meter/\sqrt{Hz}$ at 0.15 Hz which meets the requires for Advanced LIGO.

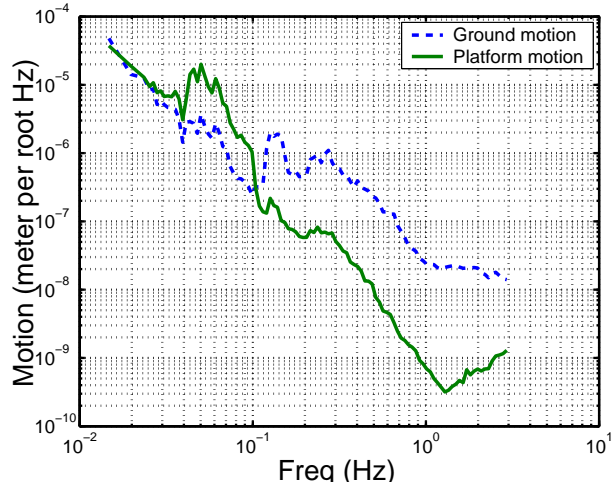


Figure 9: Horizontal vibration isolation performance. The micro-seismic peak is reduced by a factor of 10 at 0.15 Hz.

References

- [1] P. Vaidyanathan, "Multirate systems and filter banks," Prentice Hall, 1st edition, 1992.
- [2] S. Boyd and L. Vandenberghe, "Convex Optimization," Cambridge University Press.
- [3] "SeDuMi," <http://fewcal.kub.nl/sturm/software/sedumi.html>
- [4] Abramovici et. al., "LIGO - The Laser Interferometer Gravitational-wave Observatory," Science 256, (1992), pp. 325-333.
- [5] R. Abbott et. al., "Seismic Isolation Enhancements for Initial and Advanced LIGO," Classical and Quantum Gravity, volume 21, issue 5, pages S915 - S921.
- [6] D. Shoemaker, D. Coyne, "LIGO II Seismic Isolation Design Requirements Document," LIGO internal document, LIGO-E990303-02-D, Nov 4, 1999, <http://www.ligo.caltech.edu/docs/E/E990303-02.pdf>.