

16.323 Principles of Optimal Control
Final Project

Frequency Domain LQR Control

Brett Shapiro

4 May 2010

1 Introduction

Frequency dependent cost functions have useful applications for solving optimal control problems, particularly if disturbances and noise are non-Gaussian, the performance criteria are frequency dependent, or there is simply a strong frequency dependence on noise terms.

The Linear Quadratic Regulator (LQR) method, while known for its idealized guaranteed stability and relative robustness, does not on its own permit the application of frequency dependent cost functions. However, frequency weighted variables can be worked into the LQR cost function by augmenting the system plant model.

Much of this work is based off the frequency domain LQR work done by Gupta in [2].

2 Frequency Domain LQR

This section develops the theory behind the frequency domain LQR method.

The system plant is specified in Eq. (1) and Eq. (2) with the usual LTI state space notation. The state vector is x , the actuator input is u , and the output is y .

$$\dot{x} = Ax + Bu \tag{1}$$

$$y = Cx + Du \tag{2}$$

Typically the infinite horizon LQR cost function is expressed as

$$J = \int_0^{\infty} [x^T Qx + u^T Ru] dt \tag{3}$$

This cost function can be converted to the frequency domain by invoking Parseval's Theorem.

$$J = \frac{1}{2} \int_{-\infty}^{\infty} [x(-jw)^T Qx(jw) + u(-jw)^T Ru(jw)] dw \tag{4}$$

In this form, Q and R can also be generalized as frequency dependent matrices. However, since we really want to define the cost in the frequency domain, we shall specify the cost function initially in the frequency domain and then use Parseval's theorem to convert it into

a time domain form that can be minimized by solving the Riccati equation given by the standard LQR method.

$$J = \frac{1}{2} \int_{-\infty}^{\infty} [\bar{x}^T(-jw)^T \bar{x}(jw) + \bar{u}(-jw)^T \bar{u}(jw)] dw \quad (5)$$

\bar{x} is a filtered version of x and \bar{u} is a filtered version of u . Weighting matrices do not need to be specified at this point because all relative weights are incorporated into the filtering functions defined in Eqs. (6) and (7) below. χ and μ represent intermediate states.

$$\dot{\chi} = F_x \chi + G_x x \quad \bar{x} = H_x \chi + S_x x \quad (6)$$

$$\dot{\mu} = F_u \mu + G_u u \quad \bar{u} = H_u \mu + S_u u \quad (7)$$

Invoking Parseval's theorem to return to the time domain and plugging in the second column of Eqs. (6) and (7), the cost function becomes

$$J = \int_0^{\infty} [x^T S_x^T S_x x + 2x^T S_x^T H_x \chi + \chi^T H_x^T H_x \chi + \mu^T H_u^T H_u \mu + 2\mu^T H_u^T S_u u + u^T S_u^T S_u u] dt \quad (8)$$

Simplifying the notation Eq. (8) can be rewritten in a more familiar form as

$$J = \int_0^{\infty} [z^T Q z + u^T R u + 2z^T N u] dt \quad (9)$$

The augmented state z and the weighting matrices Q , R , and N are

$$z = \begin{bmatrix} x \\ \chi \\ \mu \end{bmatrix} \quad (10)$$

$$Q = \begin{bmatrix} S_x^T S_x & S_x^T H_x & 0 \\ H_x^T S_x & H_x^T H_x & 0 \\ 0 & 0 & H_u^T H_u \end{bmatrix} \quad (11)$$

$$R = S_u^T S_u \quad (12)$$

$$N = \begin{bmatrix} 0 \\ 0 \\ H_u^T S_u \end{bmatrix} \quad (13)$$

The system to plug into the LQR algorithm is augmented with the filtering states

$$\dot{z} = \bar{A} z + \bar{B} u \quad (14)$$

where

$$\bar{A} = \begin{bmatrix} A & 0 & 0 \\ G_x & F_x & 0 \\ 0 & 0 & F_u \end{bmatrix} \quad (15)$$

$$\bar{B} = \begin{bmatrix} B \\ 0 \\ G_u \end{bmatrix} \quad (16) \quad 3$$

The control law is then found by solving the algebraic Riccati equation

$$\bar{A}^T P + P \bar{A} - (P \bar{B} + N) R^{-1} (P \bar{B} + N)^T + Q = 0 \quad (17)$$

and finally

$$u = R^{-1} (P \bar{B} + N)^T z \quad (18)$$

where the feedback gain is

$$K = R^{-1} (P \bar{B} + N)^T \quad (19)$$

The usual restrictions on the weighting matrices apply where $Q \geq 0$ and $R > 0$. These requirements have implications for the filtering functions chosen to apply the frequency weighting. The filtering function on the state x must have at least as many poles as zeros, meaning the transfer function is proper or strictly proper.

The filtering function for the control input is less flexible. Note that $R = S_u^T S_u$ in Eq. (12). In general, the only way to have a nonzero S_u matrix is if the filtering function has the same number of zeros as poles. Thus, it appears we can only use proper transfer functions, not even strictly proper transfer functions. This restriction can be relaxed if the vector \bar{u} contains the control effort in addition to the filtered control effort. In other words, the S_u matrix is forced to have nonzero elements even when the number of zeros is less than the number of poles. More or less the end result is the similar, because in either case the cost function contains terms relating to both μ and u by the requirement of $R > 0$. This means no matter what we do there is always a broadband weight and a filtered weight on the control effort. However, by allowing for strictly proper weighting functions on the control effort we get more freedom in the design of the frequency dependent part of the cost function (which also tends to help prevent 'ill-conditioning' warnings in Matlab's LQR command), and we can directly chose how much we want to weight the broadband control effort term.

3 Application to a Double Pendulum Interferometer for Measuring Gravitational Waves

This control technique is illustrated with the length control of a $4km$ Fabry-Perot interferometer used to Measure gravitational waves (GWs). See Figure 1.

The existence of GWs is predicted by Einstein's Theory of General Relativity. They are produced from all accelerating mass. However, like classical gravity, GWs interact very weakly in the universe compared to other forces such as electromagnetism. Consequently, the only sources that are practical for observation are compact, massive objects with high acceleration such as merging black holes and neutron stars, supernovae, pulsars, or remnant radiation from the Big Bang. The Laser Interferometers Gravitational-Wave Observatory (LIGO, www.ligo.caltech.edu) is working to measure these waves for the first time and bring the into them realm of modern astronomy.

GWs produce a strain-like effect in space. For example, when they pass through an object they will distort its length by a small amount proportional to the original length of

the object. The object's length will expand and contract at the frequency of the wave. The amplitude of the strain is proportional to the amplitude of the GW. The strongest GWs can be expected to produce a maximum displacement of about $10^{-18}m/\sqrt{Hz}$ for a $4km$ cavity.

To measure GWs the length between the two mirrors is measured very precisely interferometrically. In order for the interferometer to remain at its operating point the optical cavity must be held near resonance which means the length of the cavity must be maintained near a constant value (absolute length is less critical). This length control is done by using feedback control to force one mirror, the active mirror, to track the other, the passive mirror. The mirrors are suspended from wires in the form of identical double pendulums. The pendulums were chosen to have multiple stages to enhance isolation from ground vibrations (in fact this is a simplification of the real LIGO pendulums that have 4 stages).

The system has a single measurement, which is the change in length of the cavity. There are two control inputs, one at each stage of the active pendulum pushing in a direction parallel to the beam axis. Only motion of the pendulum stages parallel to the beam axis is modeled. The other degrees of freedom are ignored for simplicity (in practice this is actually a reasonable assumption in most cases).

Since actuation is applied to the active pendulum, the passive mirror motion is modeled simply as a disturbance. Another disturbance, d_2 is introduced from ground vibrations through the platform the active pendulum hangs from. Drifts in the interferometer's laser frequency enter the system as sensor noise, v . There is also process noise, w_1 and w_2 , known as Barkhausen noise generated by the actuators that contaminates the control force applied to the active pendulum. This Barkhausen noise is non-Gaussian and is generated by discrete flips in the magnetic domains of the permanent magnets used in the electromagnetic actuators. The rate of domain flips is proportional to the force applied to the magnet. Thus, the magnitude of the process noise entering the loop is amplified by the control effort.

Since the GWs alter the length of space between the mirrors it turns out that they enter the system in the same way as sensor noise. Thus, increasing the gain on the feedback loop suppresses the effect of the GW in the same way the error signal is suppressed. This is not a problem however because the loop gain of the system is known. Thus, the recorded data can be calibrated into an effective open loop measurement (where the GW effect would not have been distorted) by multiplying it with the denominator of the closed loop transfer function, i.e. $1 + \text{loop gain}$. Calibrating the data to amplify the GW signal also amplifies the disturbances and noises in the same way. Thus, by design the impact of all these disturbances and noises must be less than that of the GWs in the frequency band of interest. Figure 2 has a block diagram and illustration of how this calibration is done.

Because of this calibration step the goal of the control loop is effectively insensitive to noise and disturbances in the usual linear sense since the noise must be less than the GW anyway. However, there is an issue when noise is amplified in a nonlinear sense. The example of this is the non-Gaussian Barkhausen process noise terms w_1 and w_2 . This control effort dependent process noise is one of the motivations for using a frequency dependent formulation of LQR. The influence of the Barkhausen noise can be minimized by applying small high frequency forces at the mirror itself, while the larger low frequency forces are applied to the upper stage where the amplified process noise is mechanically filtered by the stages below. Additionally, it is known that the ground vibrations, d_1 and d_2 are larger at low frequencies, so it is beneficial to place more cost on the mirror position at low frequencies to optimize the length control within the finite strength of the actuators.

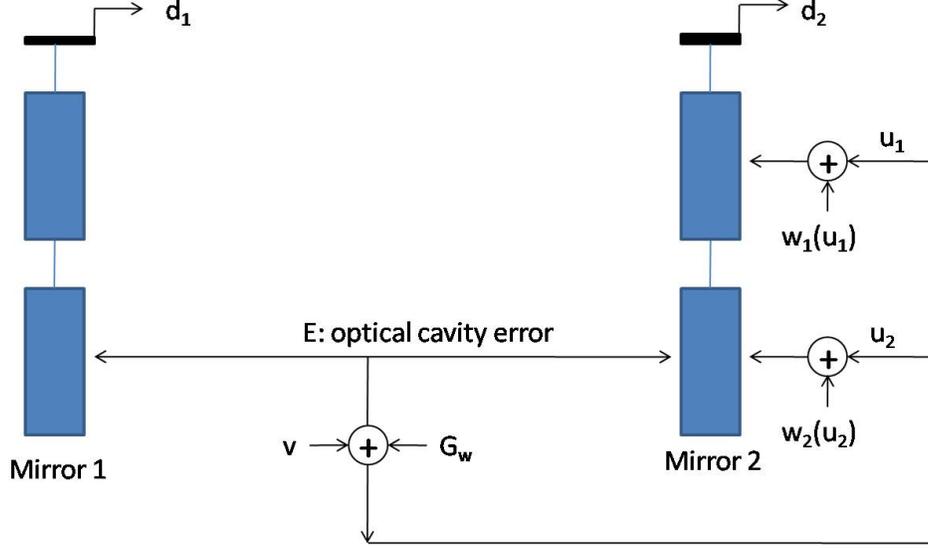


Figure 1: The Fabry-Perot interferometer resonant between two double pendulum suspended mirrors. The length of the optical cavity, nominally $4km$, is controlled by pushing on the stages of the second pendulum. The sensor noise, v , is generated by drifts in the laser frequency. d_1 and d_2 represent ground motion. The process noises w_1 and w_2 come from the actuators and is dependent on the magnitude of the control effort. The GW, G_w , enters in the same way as the sensor noise.

4 Modeling of the System

The system can be modeled in state space form as shown in Eq. (20) to (24).

$$\dot{x} = Ax + B(u + w(u)) + B_d d_2 \quad y = Cx + d_1 + G_w \quad (20)$$

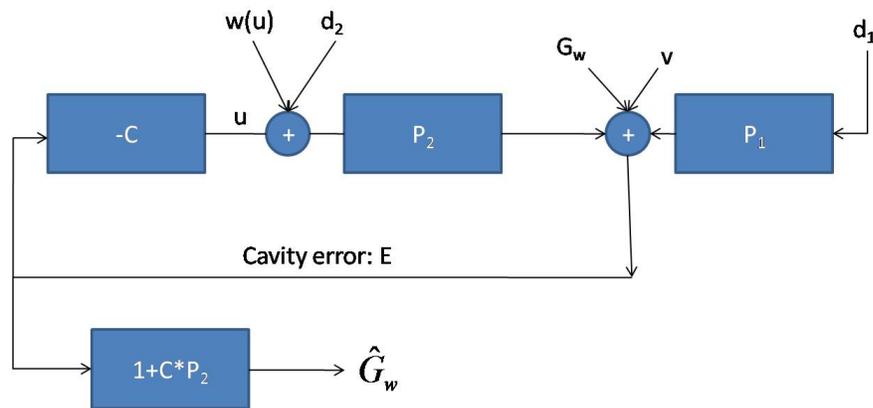
where

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -73.83 & 16.04 & -0.001 & 0 \\ 16.36 & -16.36 & 0 & -0.001 \end{bmatrix} \quad (21)$$

$$B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0.02476 & 0 \\ 0 & 0.02525 \end{bmatrix} \quad (22)$$

$$B_d = \begin{bmatrix} 0 \\ 0 \\ 57.8 \\ 0 \end{bmatrix} \quad (23)$$

$$C = [0 \ 1 \ 0 \ 0] \quad (24)$$



- Closed loop error

$$E = \frac{1}{1 + C * P_2} \{G_w + v + P_2 * [d_2 + w(u)] + P_1 * d_1\}$$

- Estimate of the GW in the frequency band of interest

$$\hat{G}_w = E * [1 + C * P_2] = G_w + v + P_2 * [d_2 + w(u)] + P_1 * d_1 \approx G_w$$

$$\text{if } G_w \gg v + P_2 * [d_2 + w(u)] + P_1 * d_1$$

Figure 2: A block diagram of the closed loop system. Since the GW is suppressed in the same way as the noises and disturbances, the effect of the GW must be extracted by calibrating the measured error signal with the measured loop gain. For this reason the expected GW effect must fundamentally be greater than all the noises and disturbances in the frequency band of interest. The equations below the block diagram show mathematically how the calibration is done.

The sensor noise is ignored since it is by design made to be small, and in fact negligible at most frequencies (in the real system shot noise becomes significant at high frequencies, but this is above the bandwidth of the controller). Here d_1 now represents exactly the motion of the passive mirror instead of the motion of its support table, where it is less than the GW at all frequencies within the band of interest, 10 Hz to 8 kHz. d_2 is still the motion of the active pendulum's support table. The Barkhausen process noise w is a function of the control effort u . The effect of this process noise is to be minimized in the design of the controller.

w is modeled as shown in Figure 3. This model is known to not represent the true effect, but it is enough to shed light on how Barkhausen noise contaminates the GW measurement when large amounts of control effort are applied directly at the mirror.

The support table motion d_2 is modeled as shown in Figure 4. d_1 is obtained by filtering this spectrum by the transfer function of the pendulum model.

The laser frequency in the interferometer has $1064nm$ wavelength, so the optical cavity must be controlled within a boundary layer much less than this. In fact, to keep it in the linear range, and to prevent higher order effects from introducing more noise terms it must be controlled within an average displacement of $10^{-12}m_{RMS}$. Additionally the top mass actuator will saturate at $1mN$ and the bottom mass (mirror) actuator will saturate at $0.1mN$. Gravitational waves are expected to be observable at $10^{-18}m/\sqrt{Hz}$, and our goal is to try to measure them between $10Hz$ and $8kHz$ (the sampling rate is $16384Hz$). Consequently, the required calibrated error signal is stated to be less than $10^{-18}m/\sqrt{Hz}$ at 10 Hz and roll down by 20 dB per decade. The roll off is included in the requirement because our noise should roll off by at least this much, so why not get better measurements of GWs at higher frequencies.

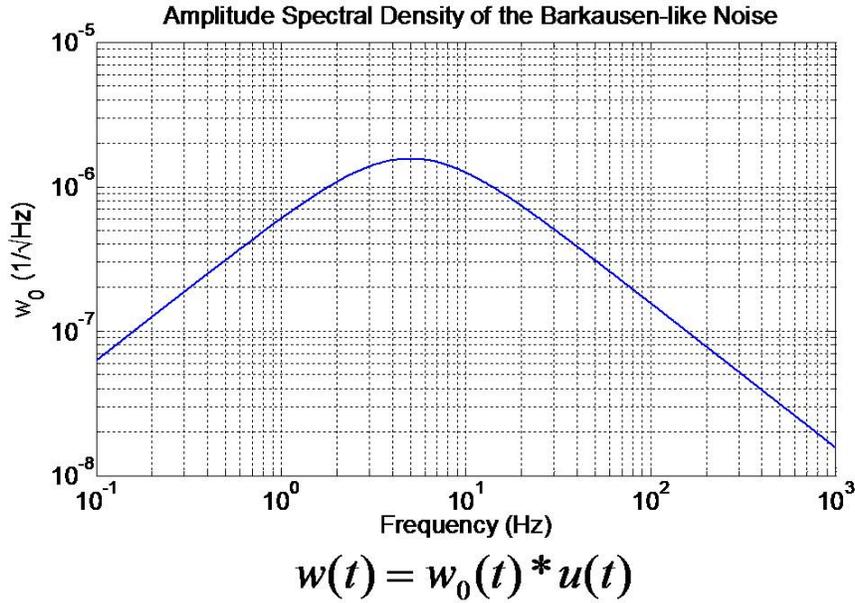


Figure 3: The model of the Barkhausen noise used in this project. The magnitude scales with the magnitude of the control effort at each time step. It rolls off at 20 dB per decade in the frequency band of interest since it is known the noise has this approximate shape. The exact behavior of Barkhausen noise is not well understood, particularly at low frequencies, so I chose to roll the gain down below 10 Hz since ground motion is known to dominates there anyway.

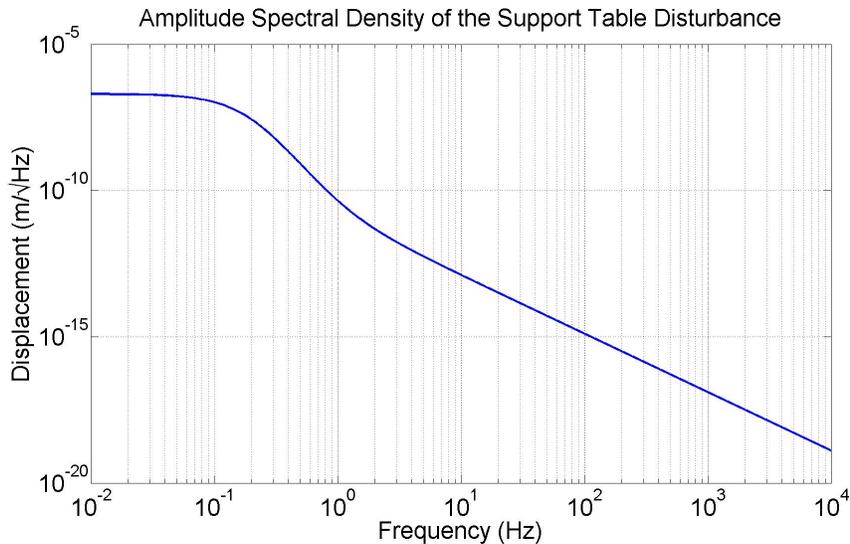


Figure 4: The approximate amplitude spectrum of the motion of the support table. The multiple stages in the pendulum were designed to filter this noise sufficiently to measure GWs above 10 Hz. This curve was achieved with previous stages of vibration isolation.

Since only a single state is measured an estimator must be employed to reconstruct the rest. A Kalman filter is not the best approach here since the noise and disturbance terms are only important relative to the GWs, not the state x . The estimator is designed using the Matlab 'place' function where the poles are all placed at least as high as desired bandwidth. In this case the place function was the most useful since the poles of the estimator could be put where they give the best loop gain characteristics. Here they were placed at $100Hz$ in an approximate butterworth pattern. With LQR design not all the poles of the estimator increased arbitrarily with decreasing cost on the measurement. Thus, placing them manually at relatively high frequencies ensures the estimator responds fast enough to produce the desired loop shape asked for by the LQR frequency dependent weighting functions used in the controller design.

The frequency domain cost functions are plotted in Figure 5. The only cost placed on the state was that of the mirror motion, x_2 . It was chosen to roll up at low frequencies because that is where the disturbance is concentrated. This choice yields a controller that mimics integral control. The weights of the control effort terms were chosen to encourage top mass actuation at low frequencies and bottom mass actuation at high frequencies. In this way, small high frequency forces are applied directly at the mirror, while the larger Barkhausen noise contaminated forces are applied at the top mass where they can be filtered by the pendulum below. Thus, the top mass control effort was weighted to have high cost at high frequencies and low cost at low frequencies while the bottom mass was weighted in the opposite way. The broadband control effort term, required to ensure $R > 0$, was simply chosen to be small enough that the shape of the frequency dependent weights would appear in the overall loop gain. Practically this meant that it is no greater than the largest frequency dependent weight within the loop bandwidth.

The resulting loop gain is given in Figure 6. Note that the overall loop gain is in fact given by the top mass at low frequencies and by the bottom mass at high frequencies.

For comparison a similar plot is given in Figure 7 with a controller that was designed with the standard LQR method. Here the same estimator is used. The LQR weights were chosen to give about the same amplitude error signal, and place as much control at the top mass as possible without saturating the actuator.

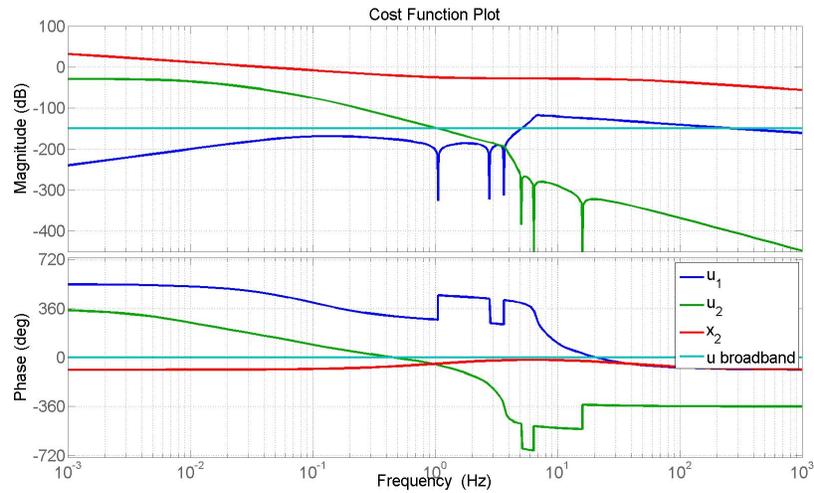


Figure 5: Bode plots of the cost functions for the frequency domain LQR method. The red curve represents the transfer function from the mirror motion x_2 to \bar{x} , the blue curve the transfer function from the top mass input u_1 to \bar{u}_1 , the green curve the transfer function from the bottom mass input u_2 to \bar{u}_2 , and the light blue curve the broadband component of the control effort cost. The significant plot is the magnitude plot on top. The phase is irrelevant to the cost.

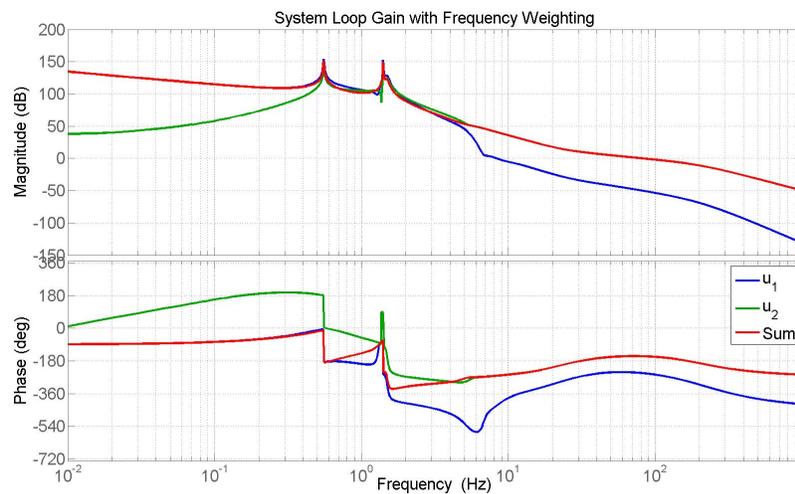


Figure 6: A bode plot of the system loop gain. The blue curve is the loop gain from the error signal to the top mass input u_1 , the green curve is the loop gain from the error signal to the bottom mass input u_2 . The red curve is the overall loop gain, which is the sum of the green and blue curves.

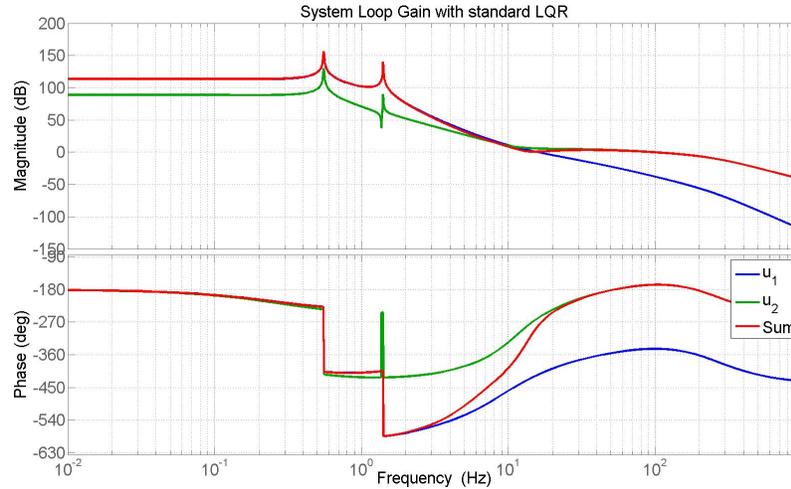


Figure 7: A bode plot of the system loop gain. The blue curve is the loop gain from the error signal to the top mass input u_1 , the green curve is the loop gain from the error signal to the bottom mass input u_2 . The red curve is the overall loop gain, which is the sum of the green and blue curves.

6 Control Law Performance

This section analyses the performance of the two LQR methods for this application. All the results were obtained from a simulation in Simulink.

Figure 8 shows that the error signal achieved with both methods is similar. Both have peak-peak values below the RMS requirement. Similarly, Figure 9 shows that the top mass force applied in both cases is comparable. In fact, the top mass forces are nearly identical. They approach the saturation limit in order to minimize the amount of force needed at the bottom mass to reduce the Barkhausen noise contamination.

A difference between the two methods appears in Figure 10. Here the standard LQR method yields about ten times more force at the bottom mass than the frequency domain LQR method. In fact the standard LQR bottom mass force is really just a scaled version of the top mass force. This is where the frequency domain cost functions really show their influence. A very low cost was placed on high frequency forces at the bottom mass, and this is what was given accordingly.

The real success of the frequency domain method is summarized in Figure 11. The simple result is that from about 10Hz to 100Hz the standard LQR method does not meet the requirement where the frequency domain LQR method does meet the requirement. At 10Hz the difference is about a factor of 30.

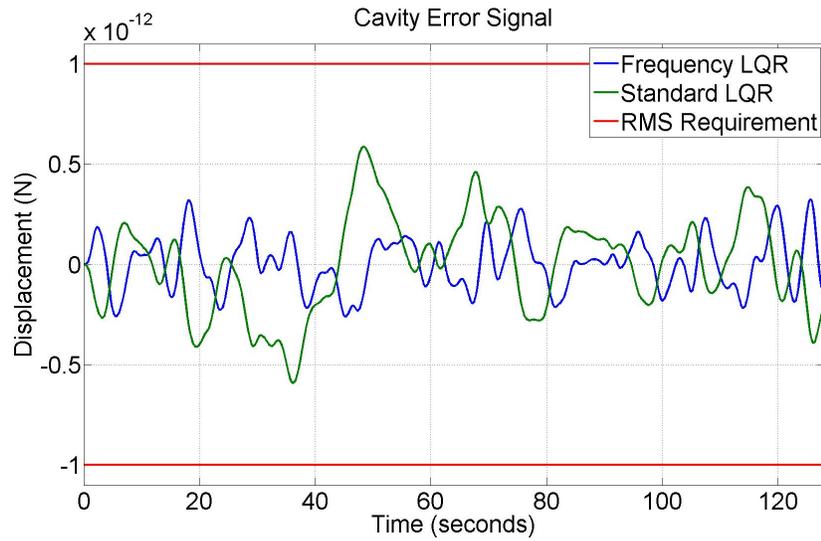


Figure 8: The error signal given by both the frequency domain and standard LQR methods. They were designed to give approximately the same amplitude.

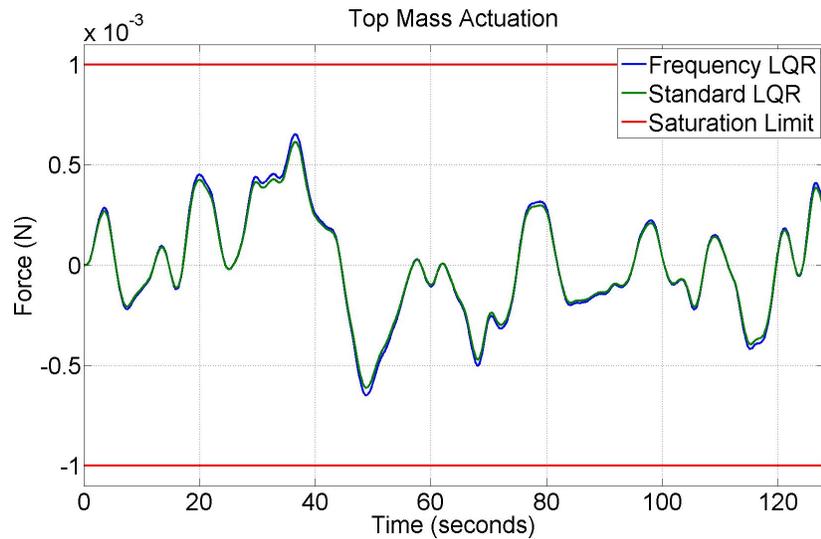


Figure 9: The actuation force applied to the top mass with both the frequency domain and standard LQR methods. Each controller was designed to nearly saturate the top mass actuation to reduce the amount of Barkhausen noise injected into the bottom mass.

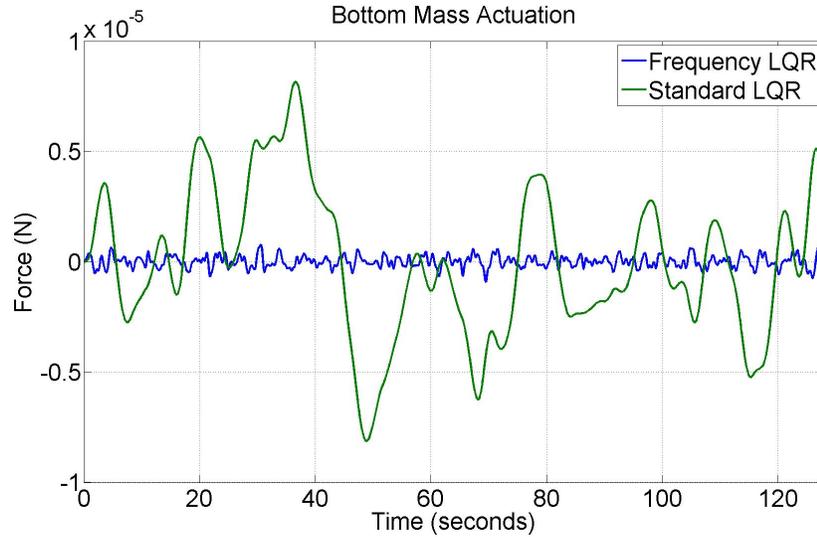


Figure 10: The actuation force applied to the bottom mass with both the frequency domain and standard LQR methods. The saturation limit is not shown on this plot because both controllers are well below the limit at this stage. The blue curve is about ten times less than the green curve.

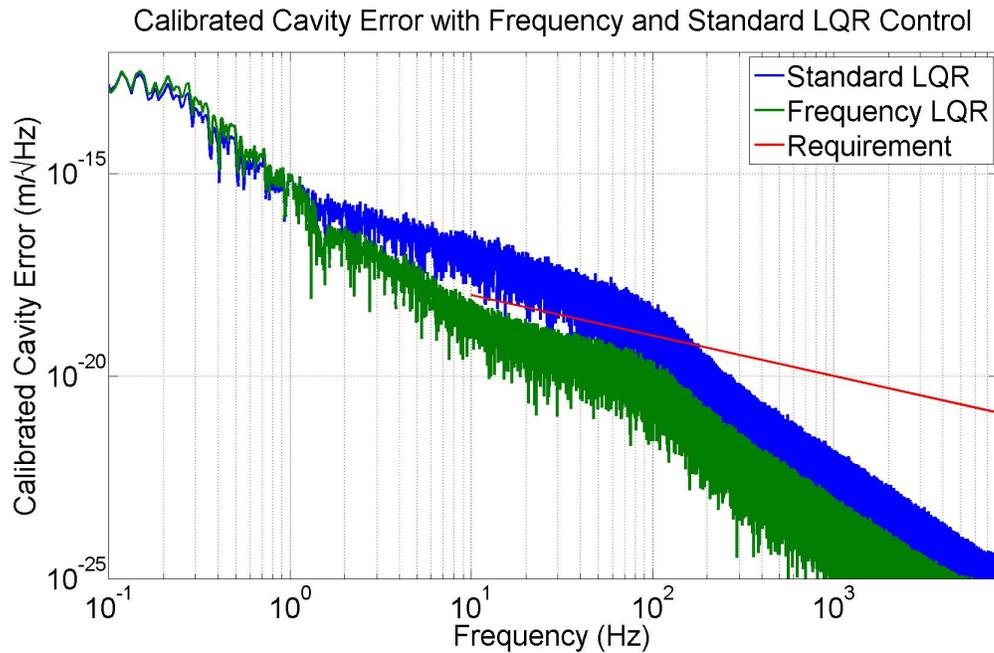


Figure 11: The amplitude spectrum of the calibrated error signal. The frequency domain LQR method (green) clearly meets the requirement (red) for measuring GWs where the standard LQR method (blue) does not.

For certain applications, including this one, the frequency domain LQR method provides a more flexible way to optimize the performance of a compensator. It combines some of the flexibility of loop shaping, where one tunes a loop by placing poles and zeros, with the guaranteed stability of the LQR method. In this case it was useful because of the non-Gaussian Barkhausen process noise and the frequency domain requirement. Another useful application would be active vibration cancellation. An example of this is in the work done by Gupta et al. where a large narrow band helicopter vibration was reduced by placing a large cost on the state at that frequency [1].

Care must be taken in the design however, especially where estimators are employed. Clearly the estimator model must be accurate. Additionally, if the estimator has poles below the desired bandwidth the loop gain may not follow what one might expect from the cost functions. In general it will be lower, and may not follow the shape of the cost functions.

Also, for reasons I do not understand, when the Matlab 'place' command is used in the estimator design the combined estimator and controller compensator may have unstable poles. This issue holds true for both LQR methods. Designing the estimator with the 'lqr' command does not appear to do this, however not all of the poles of the estimator reach a point beyond the desired bandwidth since the Hamiltonian matrix has zeros that absorb poles for small values of cost on the measurement.

References

- [1] N. Gupta and R. Du Val. A New Approach for Active Control of Rotocraft Vibration. *Journal of Guidance, Control, and Dynamics*, 5(2):143–150, 1982.
- [2] Narendra Gupta. Frequency-Shaped Cost Functionals: Extension of Linear-Quadratic-Gaussian Design Methods. *Journal of Guidance, Control, and Dynamics*, 3(6):529–535, 1980.